

Received November 7, 2021, accepted November 23, 2021, date of publication December 6, 2021, date of current version December 20, 2021.

Digital Object Identifier 10.1109/ACCESS.2021.3133370

# Reachability Analysis of Neural Feedback Loops

MICHAEL EVERETT<sup>1</sup>, (Member, IEEE), GOLNAZ HABIBI, (Member, IEEE),  
CHUANGCHUANG SUN, (Member, IEEE), AND JONATHAN P. HOW<sup>1</sup>, (Fellow, IEEE)

Aerospace Controls Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

Corresponding author: Michael Everett (mfe@mit.edu)

This work was supported in part by Ford Motor Company, in part by ARL DCIST under Cooperative Agreement W911NF-17-2-0181, and in part by Scientific Systems Company, Inc., under Research Agreement SC-1661-04.

Code: [https://github.com/mit-acl/nn\\_robustness\\_analysis](https://github.com/mit-acl/nn_robustness_analysis).

**ABSTRACT** Neural Networks (NNs) can provide major empirical performance improvements for closed-loop systems, but they also introduce challenges in formally analyzing those systems' safety properties. In particular, this work focuses on estimating the forward reachable set of *neural feedback loops* (closed-loop systems with NN controllers). Recent work provides bounds on these reachable sets, but the computationally tractable approaches yield overly conservative bounds (thus cannot be used to verify useful properties), and the methods that yield tighter bounds are too intensive for online computation. This work bridges the gap by formulating a convex optimization problem for the reachability analysis of closed-loop systems with NN controllers. While the solutions are less tight than previous (semidefinite program-based) methods, they are substantially faster to compute, and some of those computational time savings can be used to refine the bounds through new input set partitioning techniques, which is shown to dramatically reduce the tightness gap. The new framework is developed for systems with uncertainty (e.g., measurement and process noise) and nonlinearities (e.g., polynomial dynamics), and thus is shown to be applicable to real-world systems. To inform the design of an initial state set when only the target state set is known/specified, a novel algorithm for backward reachability analysis is also provided, which computes the set of states that are guaranteed to lead to the target set. The numerical experiments show that our approach (based on linear relaxations and partitioning) gives a  $5\times$  reduction in conservatism in  $150\times$  less computation time compared to the state-of-the-art. Furthermore, experiments on quadrotor, 270-state, and polynomial systems demonstrate the method's ability to handle uncertainty sources, high dimensionality, and nonlinear dynamics, respectively.


**INDEX TERMS** Reachability analysis, neural networks, deep learning, safe learning, verification.

## I. INTRODUCTION

Neural Networks (NNs) are pervasive in many fields because of their ability to express highly general input-output relationships, such as for perception, planning, and control tasks in robotics. However, before deploying NNs on safety-critical systems, there must be techniques to guarantee that the closed-loop behavior of systems with NNs, which we call *neural feedback loops* (NFLs), will meet desired specifications. The goal of this paper is to develop a framework for guaranteeing that systems with NN controllers will reach their goal states while avoiding undesirable regions of the state space, as in Fig. 1.

Despite the importance of analyzing closed-loop behavior, much of the recent work on formal NN analysis

The associate editor coordinating the review of this manuscript and approving it for publication was Ángel F. García-Fernández<sup>1</sup>.



**FIGURE 1. Forward reachability analysis.** The objective is to compute blue sets  $\mathcal{R}_t(\mathcal{X}_0)$ , to ensure a system starting in  $\mathcal{X}_0$  (yellow) ends in  $\mathcal{G}$  (green) and avoids  $\mathcal{A}_0, \mathcal{A}_1$  (red). This is challenging for systems with NN control policies.

has focused on NNs in isolation (e.g., for image classification) [1]–[6], with an emphasis on efficiently

relaxing NN nonlinearities [7]–[13]. On the other hand, closed-loop system reachability has been studied for decades, but traditional approaches, such as Hamilton-Jacobi methods [14], [15], do not consider NNs in the loop.

Several recent works [16]–[21] propose methods that compute forward reachable sets of NFLs. A key challenge is in maintaining computational efficiency while still providing tight bounds on the reachable sets. In addition, large initial state sets present challenges as the underlying NN relaxations become loose, leading to loose reachable set estimates. The literature also typically assumes perfect knowledge of system dynamics, with no stochasticity. Furthermore, the literature focuses on forward reachability analysis (what states could the system end in, starting from a given  $\mathcal{X}_0$ ?), but backward reachability analysis (what states can the system start from, to end in a given  $\mathcal{X}_T$ ?) is an equally important problem in safety analysis. While forward and backward reachability differ only by a change of variables in classical systems [15], [22], both the nonlinearities and matrix dimensions of NN controllers raise fundamental challenges that prevent propagating sets backward through a NFL.

To address these challenges, this work's contributions include:

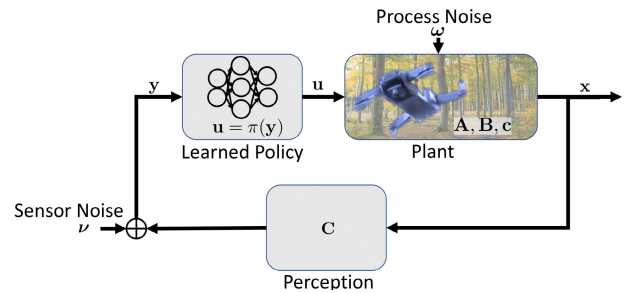
- A linear programming-based formulation of forward reachability analysis for NFLs, which provides a computationally efficient method for verifying safety properties,
- The use of input set partitioning techniques to provide tight bounds on the reachable sets despite large initial state sets,
- The consideration of measurement noise, process noise, and nonlinear dynamics, which improves the applicability to real systems with uncertainty, and
- A framework for backward reachability analysis of NFLs, which enables estimating which starting states will lead to a target set despite non-invertible NN weight matrices and activations.

Numerical experiments show that the new method provides  $5\times$  better accuracy in  $150\times$  less computation time compared to [21] and that the method can handle uncertainty sources, high dimensionality, and nonlinear dynamics via applications on quadrotor, 270-state, and polynomial systems.

This article extends our conference paper [23] by providing an improved problem formulation, notation, and algorithmic details. It also includes:

- A closed-form solution to provide an order-of-magnitude computational speedup for the same bounds in §IV-E,
- An extension of state-of-the-art partitioning techniques to the NFL setting in §V,
- Formulations that handle nonlinearities (control limits and polynomial terms) in the dynamics in §VI,
- A framework for backward reachability analysis in §VII.

Open-source software implementations of this paper's algorithms and results can be found at [https://github.com/mit-acl/nn\\_robustness\\_analysis](https://github.com/mit-acl/nn_robustness_analysis).



**FIGURE 2. Neural Feedback Loop.** Many real world deployments of NNs are part of complex, closed-loop systems (e.g., robots with learned policies take actions that modify the world state and influence future observations).

## II. RELATED WORK

Related work on reachability analysis can be categorized into works on NNs in isolation, closed-loop systems without NNs, and closed-loop systems with NNs. For instance, machine learning literature includes many methods to verify properties of NNs, often motivated by defending against adversarial examples [24]. These methods broadly range from exact [2] to tight [13] to efficient [10] to fast [7]. Although these tools are not designed for closed-loop systems, the NN relaxations from [10] provide a key foundation here.

For closed-loop systems, reachability analysis is a standard component of safety verification. Modern methods include Hamilton-Jacobi Reachability methods [14], [15], SpaceX [25], Flow\* [26], CORA [27], and C2E2 [28], [29], but these do not account for NN control policies. Orthogonal approaches that do not explicitly estimate the system's forward reachable set, but provide other notions of safety, include Lyapunov function search [30] and control barrier functions (CBFs) [31].

Recent reachability analysis approaches that do account for NN control policies face a tradeoff between computation time and conservatism. [16]–[18] use polynomial approximations of NNs to make the analysis tractable. Most works consider NNs with ReLU approximations, whereas [19] considers sigmoidal activations. [20], [32] introduce conservatism by assuming the NN controller could output its extreme values at every state. Most recently, [21] formulated the problem as a SDP, called Reach-SDP. This work builds on both [20], [21] and makes the latter more scalable by re-formulating the SDP as a linear program, introduces sources of uncertainty in the closed-loop dynamics, and shows further improvements by partitioning the input set.

## III. PRELIMINARIES

### A. CLOSED-LOOP SYSTEM DYNAMICS

Consider a discrete-time linear time-varying system,

$$\begin{aligned} \mathbf{x}_{t+1} &= \mathbf{A}_t \mathbf{x}_t + \mathbf{B}_t \mathbf{u}_t + \mathbf{c}_t + \boldsymbol{\omega}_t \\ \mathbf{y}_t &= \mathbf{C}_t^T \mathbf{x}_t + \mathbf{v}_t, \end{aligned} \quad (1)$$

where  $\mathbf{x}_t \in \mathbb{R}^{n_x}$ ,  $\mathbf{u}_t \in \mathbb{R}^{n_u}$ ,  $\mathbf{y}_t \in \mathbb{R}^{n_y}$  are state, control, and output vectors,  $\mathbf{A}_t$ ,  $\mathbf{B}_t$ ,  $\mathbf{C}_t$  are known system matrices,  $\mathbf{c}_t \in \mathbb{R}^{n_x}$  is a known exogenous input, and  $\boldsymbol{\omega}_t \sim \Omega$  and  $\mathbf{v}_t \sim N$

are process and measurement noises sampled at each timestep from unknown distributions with known, finite support (i.e.,  $\omega_t \in [\underline{\omega}_t, \bar{\omega}_t]$ ,  $\mathbf{v}_t \in [\underline{\mathbf{v}}_t, \bar{\mathbf{v}}_t]$  element-wise).

We assume an output-feedback controller  $\mathbf{u}_t = \pi(\mathbf{y}_t)$  parameterized by an  $m$ -layer feed-forward NN, optionally subject to control constraints,  $\mathbf{u}_t \in \mathcal{U}_t$ . We denote the closed-loop system with dynamics (1) and control policy  $\pi$  as

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t; \pi). \quad (2)$$

This NFL is visualized in Fig. 2. At each timestep, the state  $\mathbf{x}$  enters the perception block and is perturbed by sensor noise  $\mathbf{v}$  to create observation  $\mathbf{y}$ . The observation is fed into the learned policy, which is a NN that selects the control input  $\mathbf{u}$ . The control input enters the plant, which is defined by  $\mathbf{A}$ ,  $\mathbf{B}$ ,  $\mathbf{c}$  matrices and subject to process noise  $\omega$ . The output of the plant is the state vector at the next timestep.

### B. REACHABLE SETS

For the closed-loop system (2), we denote  $\mathcal{R}_t(\mathcal{X}_0)$  the forward reachable set at time  $t$  from a given set of initial conditions  $\mathcal{X}_0 \subseteq \mathbb{R}^{n_x}$ , which is defined by the recursion

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

### C. FINITE-TIME REACH-AVOID VERIFICATION PROBLEM

The finite-time reach-avoid property verification problem is defined as follows: Given a goal set  $\mathcal{G} \subseteq \mathbb{R}^{n_x}$ , a sequence of avoid sets  $\mathcal{A}_t \subseteq \mathbb{R}^{n_x}$ , and a sequence of reachable set estimates  $\mathcal{R}_t \subseteq \mathbb{R}^{n_x}$ , determining that every state in the final estimated reachable set will be in the goal set and any state in the estimated reachable sets will not enter an avoid set requires computing set intersections,  $\text{VERIFIED}(\mathcal{G}, \mathcal{A}_{0:N}, \mathcal{R}_{0:N}) \equiv \mathcal{R}_N \subseteq \mathcal{G} \ \& \ \mathcal{R}_t \cap \mathcal{A}_t = \emptyset, \forall t \in \{0, \dots, N\}$ .

In the case of our nonlinear closed-loop system (2), where computing the reachable sets exactly is computationally intractable, we can instead compute outer-approximations of the reachable sets,  $\bar{\mathcal{R}}_t(\mathcal{X}_0) \supseteq \mathcal{R}_t(\mathcal{X}_0)$ . This is useful if the finite-time reach-avoid properties of the system as described by outer-approximations of the reachable sets are verified, because that implies the finite-time reach-avoid properties of the *exact* closed loop system are verified as well. Tight outer-approximations of the reachable sets are desirable, as they enable verification of tight goal and avoid set specifications, and they reduce the chances of verification being unsuccessful even if the exact system meets the specifications.

### D. CONTROL POLICY NEURAL NETWORK STRUCTURE

Using notation from [10], for the  $m$ -layer neural network used in the control policy, the number of neurons in each layer is  $n_l, \forall l \in [m]$ , where  $[a]$  denotes the set  $\{1, 2, \dots, a\}$ . Let the  $l$ -th layer weight matrix be  $\mathbf{W}^{(l)} \in \mathbb{R}^{n_l \times n_{l-1}}$  and bias vector be  $\mathbf{b}^{(l)} \in \mathbb{R}^{n_l}$ , and let  $\mathbf{F}_l : \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_l}$  be the operator mapping from network input (measured output vector  $\mathbf{y}_t$ ) to layer  $l$ . We have  $\mathbf{F}_l(\mathbf{y}_t) = \sigma(\mathbf{W}^{(l)}\mathbf{F}_{l-1}(\mathbf{y}_t) + \mathbf{b}^{(l)}), \forall l \in [m-1]$ , where  $\sigma(\cdot)$  is the coordinate-wise activation function. The

framework applies to general activations, including ReLU,  $\sigma(\mathbf{z}) = \max(0, \mathbf{z})$ . The network input  $\mathbf{F}_0(\mathbf{y}_t) = \mathbf{y}_t$  produces the (unclipped) control input,

$$\mathbf{u}_t = \pi(\mathbf{y}_t) = \mathbf{F}_m(\mathbf{y}_t) = \mathbf{W}^{(m)}\mathbf{F}_{m-1}(\mathbf{y}_t) + \mathbf{b}^{(m)}. \quad (4)$$

### E. NEURAL NETWORK ROBUSTNESS VERIFICATION

A key step in quickly computing reachable sets of the NFL (2) is to relax the nonlinear constraints induced by the NN's nonlinear activation functions. The relaxation converts each nonlinearity into a linear upper and lower bound, where each bound holds within the known range of inputs to the activation.

To represent the range of inputs to the activation functions, we first define the  $\epsilon$ -ball. Denote the  $\epsilon$ -ball (also called the  $\ell_p$ -ball) under the  $\ell_p$  norm, centered at  $\mathbf{x}$ , with scalar radius,  $\epsilon$ ,

$$\mathcal{B}_p(\mathbf{x}, \epsilon) = \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}\|_p \leq \epsilon\}. \quad (5)$$

The  $\epsilon$ -ball is extended to the case of vector  $\epsilon \in \mathbb{R}_{\geq 0}^n$  (i.e.,  $\epsilon$ -ball), defined as

$$\mathcal{B}_p(\mathbf{x}, \epsilon) = \{\mathbf{x} \mid \lim_{\epsilon' \rightarrow \epsilon^+} \|(\mathbf{x} - \mathbf{x}) \oslash \epsilon'\|_p \leq 1\}, \quad (6)$$

where  $\oslash$  denotes element-wise division.

*Theorem 3.1 (From [10], Convex Relaxation of NN):* Given an  $m$ -layer neural network control policy  $\pi : \mathbb{R}^{n_y} \rightarrow \mathbb{R}^{n_u}$ , there exist two explicit functions  $\pi_j^L : \mathbb{R}^{n_y} \rightarrow \mathbb{R}^{n_u}$  and  $\pi_j^U : \mathbb{R}^{n_y} \rightarrow \mathbb{R}^{n_u}$  such that  $\forall j \in [n_m], \forall \mathbf{y} \in \mathcal{B}_p(\mathbf{y}, \epsilon)$ , the inequality  $\pi_j^L(\mathbf{y}) \leq \pi_j(\mathbf{y}) \leq \pi_j^U(\mathbf{y})$  holds true, where

$$\pi_j^U(\mathbf{y}) = \Psi_{j,:} \mathbf{y} + \alpha_j \quad (7)$$

$$\pi_j^L(\mathbf{y}) = \Phi_{j,:} \mathbf{y} + \beta_j, \quad (8)$$

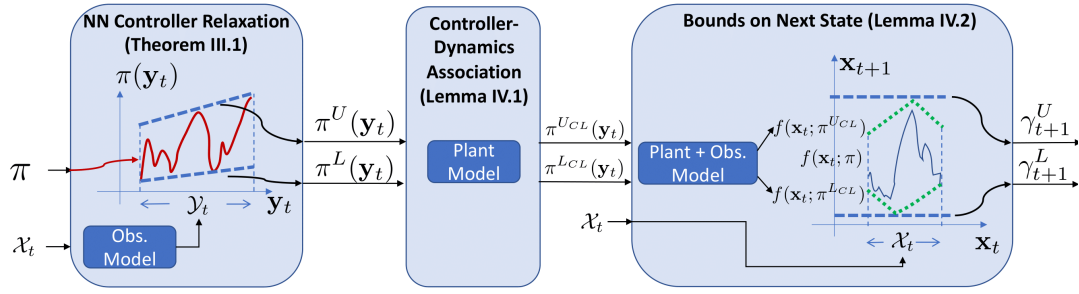
where  $\Psi, \Phi \in \mathbb{R}^{n_u \times n_y}$  and  $\alpha, \beta \in \mathbb{R}^{n_u}$  are defined recursively using NN weights, biases, and activations (e.g., ReLU, sigmoid, tanh), as detailed in [10].

In a closed-loop system, Theorem 3.1 bounds the control output for a *particular* measurement  $\mathbf{y}$ . Moreover, if all that is known is  $\mathbf{y} \in \mathcal{B}_p(\mathbf{y}, \epsilon)$ , Theorem 3.1 provides affine relationships between  $\mathbf{y}$  and  $\mathbf{u}$  (i.e., bounds valid within the known set of possible  $\mathbf{y}$ ). These relationships enable efficient calculation of NN output bounds, using Corollary 3.3 of [10].

We could leverage [10] to compute reachable sets by first bounding the possible controls, then bounding the next state set by applying the extreme controls from each state. This is roughly the approach in [20], [32], for example. However, this introduces excessive conservatism, because both extremes of control would not be applied at every state (barring pathological examples). To produce tight bounds on the reachable sets, we leverage the relationship between measured output and control in § IV.

### IV. FORWARD REACHABILITY ANALYSIS

Recall that our goal is to find the set of all possible next states,  $\mathbf{x}_{t+1} \in \mathcal{X}_{t+1}$ , given that the current state lies within a known set,  $\mathbf{x}_t \in \mathcal{X}_t$ . This will allow us to compute reachable sets recursively starting from an initial set  $\mathcal{X}_t = \mathcal{X}_0$ .



**FIGURE 3.** Approach overview for simple 1D system. Theorem 3.1 relaxes the NN to give affine relationships between observation  $\mathbf{y}_t$  and control:  $\pi^U, \pi^L$ . Theorem 4.1 uses the system dynamics to associate  $\pi^U, \pi^L$  with the next state set. Theorem 4.2 optimizes the closed-loop dynamics over all states  $\mathbf{x}_t \in \mathcal{X}_t$  to compute bounds on the next state,  $\gamma_{t+1}^U, \gamma_{t+1}^L$ .

The approach follows the architecture in Fig. 3. After first relaxing the NN controller using Theorem 3.1, we then associate linearized extreme controllers with extreme next states in § IV-B. Then, using the linearized extreme controller, we optimize over all states in the input set to find extreme next states in § IV-C. We extend the formulation to handle control limits in § VI-A, then describe how to convert the solutions of the optimization problems into reachable set descriptions in § IV-D.

#### A. ASSUMPTIONS & PROBLEM STATEMENT

This work assumes that  $\mathcal{X}_t$  is described by either:

- an  $\ell_p$ -ball for some norm  $p \in [1, \infty]$ , radius  $\epsilon$ , and centroid  $\mathbf{x}$ , s.t.  $\mathcal{X}_t = \mathcal{B}_p(\mathbf{x}, \epsilon)$ ; or
- a polytope, for some  $\mathbf{A}^{\text{in}} \in \mathbb{R}^{m_{\text{in}} \times n_x}$ ,  $\mathbf{b}^{\text{in}} \in \mathbb{R}^{m_{\text{in}}}$ , s.t.  $\mathcal{X}_t = \{\mathbf{x}_t \mid \mathbf{A}^{\text{in}} \mathbf{x}_t \leq \mathbf{b}^{\text{in}}\}$ ,

and shows how to compute  $\mathcal{X}_{t+1}$  as described by either:

- an  $\ell_\infty$ -ball with radius  $\epsilon$  and centroid  $\mathbf{x}$ , s.t.  $\mathcal{X}_{t+1} = \mathcal{B}_\infty(\mathbf{x}, \epsilon)$ ; or
- a polytope for a specified  $\mathbf{A}^{\text{out}} \in \mathbb{R}^{m_{\text{out}} \times n_x}$ , meaning we will compute  $\mathbf{b}^{\text{out}} \in \mathbb{R}^{m_{\text{out}}}$  s.t.  $\mathcal{X}_{t+1} = \{\mathbf{x}_t \in \mathbb{R}^{n_x} \mid \mathbf{A}^{\text{out}} \mathbf{x}_t \leq \mathbf{b}^{\text{out}}\}$ .

We assume that either  $\mathbf{A}^{\text{out}}$  is provided (in the case of polytope output bounds), or that  $\mathbf{A}^{\text{out}} = \mathbf{I}_{n_x}$  (in the case of  $\ell_\infty$  output bounds). Note that we use  $i$  to index polytope facets,  $j$  to index the control vectors, and  $k$  to index the state vectors. In this section, we assume that  $\mathcal{U}_t = \mathbb{R}^{n_u}$  (no control input constraints) for cleaner notation; this assumption is relaxed in § VI-A.

The exact 1-step closed-loop reachability problem is as follows. For each row  $i$  in  $\mathbf{A}^{\text{out}}$ , solve the following optimization problem,

$$\begin{aligned} (\mathbf{b}_i^{\text{out}})^* &= \max_{\mathbf{x}_t \in \mathcal{X}_t} \mathbf{A}_{i,:}^{\text{out}} \mathbf{x}_{t+1} \\ &\text{s.t. } \mathbf{x}_{t+1} = f(\mathbf{x}_t; \pi), \end{aligned} \quad (9)$$

where  $(\mathbf{b}_i^{\text{out}})^*$  defines the tightest description of  $\mathcal{X}_{t+1}$  associated with  $\mathbf{A}^{\text{out}}$ . However, the nonlinearities in  $\pi$  make solving this problem intractable in practice, so this section describes a relaxation that provides bounds on (9).

#### B. BOUNDS ON $\mathbf{x}_{t+1}$ FROM A PARTICULAR $\mathbf{x}_t$

*Lemma 4.1:* Given an  $m$ -layer NN control policy  $\pi : \mathbb{R}^{n_y} \rightarrow \mathbb{R}^{n_u}$ , closed-loop dynamics  $f : \mathbb{R}^{n_x} \times \Pi \rightarrow \mathbb{R}^{n_x}$  as in Eqs. (1) and (2), and specification matrix  $\mathbf{A}^{\text{out}} \in \mathbb{R}^{m_{\text{out}} \times n_x}$ , for each  $i \in [m_{\text{out}}]$ , there exist two explicit functions  $\pi_{:,i}^{\text{LCL}} : \mathbb{R}^{n_y} \rightarrow \mathbb{R}^{n_u}$  and  $\pi_{:,i}^{\text{UCL}} : \mathbb{R}^{n_y} \rightarrow \mathbb{R}^{n_u}$  such that  $\forall j \in [n_m], \forall \mathbf{x}_t \in \mathcal{B}_p(\mathbf{x}_t, 0, \epsilon)$  and  $\forall \mathbf{y}_t \in \mathcal{B}_\infty(\mathbf{C}_t^T \mathbf{x}_t + \frac{\bar{v}_t + \underline{v}_t}{2}, \frac{\bar{v}_t - \underline{v}_t}{2})$ , the inequality  $\mathbf{A}_{i,:}^{\text{out}} f(\mathbf{x}_t, \pi_{:,i}^{\text{LCL}}) \leq \mathbf{A}_{i,:}^{\text{out}} f(\mathbf{x}_t, \pi) \leq \mathbf{A}_{i,:}^{\text{out}} f(\mathbf{x}_t, \pi_{:,i}^{\text{UCL}})$  holds true, where

$$\pi_{:,i}^{\text{UCL}}(\mathbf{y}_t) = \mathbf{\Upsilon}_{i,:} \mathbf{y}_t + \mathbf{\Gamma}_{:,i} \quad (10)$$

$$\pi_{:,i}^{\text{LCL}}(\mathbf{y}_t) = \mathbf{\Xi}_{i,:} \mathbf{y}_t + \mathbf{\Delta}_{:,i}, \quad (11)$$

letting  $\mathbf{\Upsilon}, \mathbf{\Xi} \in \mathbb{R}^{m_{\text{out}} \times n_u \times n_y}$  and  $\mathbf{\Gamma}, \mathbf{\Delta} \in \mathbb{R}^{n_u \times m_{\text{out}}}$ ,

$$\mathbf{\Upsilon}_{i,:} = s(\mathbf{A}_{i,:}^{\text{out}} \mathbf{B}_t, \mathbf{\Psi}, \mathbf{\Phi}) \quad (12)$$

$$\mathbf{\Xi}_{i,:} = s(\mathbf{A}_{i,:}^{\text{out}} \mathbf{B}_t, \mathbf{\Phi}, \mathbf{\Psi}) \quad (13)$$

$$\mathbf{\Gamma}_{:,i} = s(\mathbf{A}_{i,:}^{\text{out}} \mathbf{B}_t, \mathbf{\alpha}, \mathbf{\beta}) \quad (14)$$

$$\mathbf{\Delta}_{:,i} = s(\mathbf{A}_{i,:}^{\text{out}} \mathbf{B}_t, \mathbf{\beta}, \mathbf{\alpha}), \quad (15)$$

using selector function (similar to `torch.where`)  $s : \mathbb{R}^m \times \mathbb{R}^{m \times n} \times \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$ , where the matrix returned by  $s$  is defined element-wise  $\forall a \in [m], b \in [n]$ ,

$$[s(\mathbf{z}, \mathbf{A}, \mathbf{B})]_{a,b} = \begin{cases} \mathbf{A}_{a,b}, & \text{if } \mathbf{z}_a \geq 0 \\ \mathbf{B}_{a,b}, & \text{otherwise} \end{cases} \quad (16)$$

and where  $\mathbf{\Psi}, \mathbf{\Phi}, \mathbf{\alpha}, \mathbf{\beta}$  are computed from Theorem 3.1 with  $\mathbf{y}_0 = \mathbf{C}_t^T (\mathbf{x}_t, 0 + \frac{\bar{v}_t + \underline{v}_t}{2})$ , and  $\epsilon = \epsilon + \frac{\bar{v}_t - \underline{v}_t}{2}$ .

*Proof:* For any particular measurement  $\mathbf{y}_t$ , after relaxing the NN according to Theorem 3.1, let  $\Pi(\mathbf{y}_t) = \{\pi \mid \pi_j^L(\mathbf{y}_t) \leq \pi_j(\mathbf{y}_t) \leq \pi_j^U(\mathbf{y}_t) \forall j \in [n_u]\}$  denote the set of possible effective control policies. Denote the control policy  $\pi_{:,i}^{\text{UCL}} \in \Pi(\mathbf{y}_t)$  as one that induces the least upper bound on the  $i$ -th facet of the next state polytope,

$$\begin{aligned} &\mathbf{A}_{i,:}^{\text{out}} f(\mathbf{x}_t; \pi_{:,i}^{\text{UCL}}) \\ &= \max_{\pi \in \Pi(\mathbf{y}_t)} \mathbf{A}_{i,:}^{\text{out}} f(\mathbf{x}_t; \pi) \\ &= \max_{\pi \in \Pi(\mathbf{y}_t)} \mathbf{A}_{i,:}^{\text{out}} [\mathbf{A}_t \mathbf{x}_t + \mathbf{B}_t \pi(\mathbf{y}_t) + \mathbf{c}_t + \boldsymbol{\omega}_t] \\ &= \left[ \max_{\pi \in \Pi(\mathbf{y}_t)} \mathbf{A}_{i,:}^{\text{out}} \mathbf{B}_t \pi(\mathbf{y}_t) \right] + \mathbf{A}_{i,:}^{\text{out}} [\mathbf{A}_t \mathbf{x}_t + \mathbf{c}_t + \boldsymbol{\omega}_t], \end{aligned} \quad (17)$$



Thus for  $\mathbf{y}_t$ ,

$$\pi_{:,i}^{UCL} = \operatorname{argmax}_{\pi \in \Pi(\mathbf{y}_t)} \mathbf{A}_{i,:}^{\text{out}} \mathbf{B}_t \pi(\mathbf{y}_t). \quad (18)$$

The resulting control input  $\forall i \in [m_{t+1}], j \in [n_u]$  is,

$$\pi_{j,i}^{UCL}(\mathbf{y}_t) = \begin{cases} \pi_j^U(\mathbf{y}_t), & \text{if } \mathbf{A}_{i,:}^{\text{out}} \mathbf{B}_{t,:j} \geq 0 \\ \pi_j^L(\mathbf{y}_t), & \text{otherwise} \end{cases}. \quad (19)$$

Writing (19) in matrix form results in (10). The proof of the lower bound follows similarly.  $\square$

### C. BOUNDS ON $\mathbf{x}_{t+1}$ FROM ANY $\mathbf{x}_t \in \mathcal{X}_t$

Now that we can bound each facet of the next state polytope given a particular current state and observation, we can form bounds on the next state polytope facet given a set of possible current states. This is necessary to handle initial state set constraints and to compute “ $t > 1$ ”-step reachable sets recursively as in (3). We assume  $\mathbf{x}_t \in \mathcal{X}_t$ .

*Lemma 4.2:* Given an  $m$ -layer NN control policy  $\pi : \mathbb{R}^{n_y} \rightarrow \mathbb{R}^{n_u}$ , closed-loop dynamics  $f : \mathbb{R}^{n_x} \times \Pi \rightarrow \mathbb{R}^{n_x}$  as in Eqs. (1) and (2), and specification matrix  $\mathbf{A}^{\text{out}} \in \mathbb{R}^{m_{\text{out}} \times n_x}$ , for each  $i \in [m_{\text{out}}]$ , there exist two fixed values  $\boldsymbol{\gamma}_{t+1,i}^U$  and  $\boldsymbol{\gamma}_{t+1,i}^L$  such that  $\forall \mathbf{x}_t \in \mathcal{X}_t$ , the inequality  $\boldsymbol{\gamma}_{t+1,i}^L \leq \mathbf{A}_{i,:}^{\text{out}} f(\mathbf{x}_t; \pi) \leq \boldsymbol{\gamma}_{t+1,i}^U$  holds true, where

$$\boldsymbol{\gamma}_{t+1,i}^U = \max_{\mathbf{x}_t \in \mathcal{X}_t} \mathbf{M}_{i,:}^U \mathbf{x}_t + \mathbf{n}_i^U \quad (20)$$

$$\boldsymbol{\gamma}_{t+1,i}^L = \min_{\mathbf{x}_t \in \mathcal{X}_t} \mathbf{M}_{i,:}^L \mathbf{x}_t + \mathbf{n}_i^L, \quad (21)$$

with  $\mathbf{M}^U \in \mathbb{R}^{n_x \times n_x}$ ,  $\mathbf{n}^U \in \mathbb{R}^{n_x}$  defined as

$$\mathbf{M}_{i,:}^U = \left( \mathbf{A}_{i,:}^{\text{out}} \left( \mathbf{A}_t + \mathbf{B}_t \boldsymbol{\Upsilon}_{i,:} \mathbf{C}_t^T \right) \right) \quad (22)$$

$$\mathbf{M}_{i,:}^L = \left( \mathbf{A}_{i,:}^{\text{out}} \left( \mathbf{A}_t + \mathbf{B}_t \boldsymbol{\Xi}_{i,:} \mathbf{C}_t^T \right) \right) \quad (23)$$

$$\mathbf{n}_i^U = \mathbf{A}_{i,:}^{\text{out}} \left( \mathbf{B}_t \left( \boldsymbol{\Upsilon}_{i,:} s(\mathbf{A}_{i,:}^{\text{out}} \mathbf{B}_t \boldsymbol{\Upsilon}_{i,:} \mathbf{C}_t^T, \bar{\mathbf{v}}_t, \underline{\mathbf{v}}_t) + \boldsymbol{\Gamma}_{:,i} \right) + \mathbf{c}_t + s(\mathbf{A}_{i,:}^{\text{out}} \mathbf{B}_t, \bar{\boldsymbol{\omega}}_t, \underline{\boldsymbol{\omega}}_t) \right) \quad (24)$$

$$\mathbf{n}_i^L = \mathbf{A}_{i,:}^{\text{out}} \left( \mathbf{B}_t \left( \boldsymbol{\Xi}_{i,:} s(\mathbf{A}_{i,:}^{\text{out}} \mathbf{B}_t \boldsymbol{\Xi}_{i,:} \mathbf{C}_t^T, \underline{\mathbf{v}}_t, \bar{\mathbf{v}}_t) + \boldsymbol{\Delta}_{:,i} \right) + \mathbf{c}_t + s(\mathbf{A}_{i,:}^{\text{out}} \mathbf{B}_t, \underline{\boldsymbol{\omega}}_t, \bar{\boldsymbol{\omega}}_t) \right), \quad (25)$$

with  $\boldsymbol{\Upsilon}, \boldsymbol{\Xi}, \boldsymbol{\Gamma}, \boldsymbol{\Delta}$  computed from Theorem 4.1.

*Proof:* Bound the next state polytope’s  $i$ -th facet above,

$$\mathbf{A}_{i,:}^{\text{out}} \mathbf{x}_{t+1} = \mathbf{A}_{i,:}^{\text{out}} f(\mathbf{x}_t; \pi) \quad (26)$$

$$\leq \mathbf{A}_{i,:}^{\text{out}} f(\mathbf{x}_t; \pi_{:,i}^{UCL}) \quad (27)$$

$$\leq \max_{\mathbf{x}_t \in \mathcal{X}_t} \mathbf{A}_{i,:}^{\text{out}} f(\mathbf{x}_t; \pi_{:,i}^{UCL}) := \boldsymbol{\gamma}_{t+1,i}^U \quad (28)$$

$$= \max_{\mathbf{x}_t \in \mathcal{X}_t} \mathbf{A}_{i,:}^{\text{out}} \left[ \mathbf{A}_t \mathbf{x}_t + \mathbf{B}_t \pi_{:,i}^{UCL}(\mathbf{y}_t) + \mathbf{c}_t + \boldsymbol{\omega}_t \right] \quad (29)$$

$$= \max_{\mathbf{x}_t \in \mathcal{X}_t} \mathbf{A}_{i,:}^{\text{out}} \left[ \mathbf{A}_t \mathbf{x}_t + \mathbf{B}_t \left( \boldsymbol{\Upsilon}_{i,:} \mathbf{y}_t + \boldsymbol{\Gamma}_{:,i} \right) + \mathbf{c}_t + \boldsymbol{\omega}_t \right] \quad (30)$$

$$= \max_{\mathbf{x}_t \in \mathcal{X}_t} \mathbf{A}_{i,:}^{\text{out}} \left[ \mathbf{A}_t \mathbf{x}_t + \mathbf{B}_t \left( \boldsymbol{\Upsilon}_{i,:} \left( \mathbf{C}_t^T \mathbf{x}_t + \mathbf{v}_t \right) + \boldsymbol{\Gamma}_{:,i} \right) + \mathbf{c}_t + \boldsymbol{\omega}_t \right] \quad (31)$$

$$= \max_{\mathbf{x}_t \in \mathcal{X}_t} \left( \mathbf{A}_{i,:}^{\text{out}} \left( \mathbf{A}_t + \mathbf{B}_t \boldsymbol{\Upsilon}_{i,:} \mathbf{C}_t^T \right) \right) \mathbf{x}_t + \mathbf{A}_{i,:}^{\text{out}} \left( \mathbf{B}_t \left( \boldsymbol{\Upsilon}_{i,:} \mathbf{v}_t + \boldsymbol{\Gamma}_{:,i} \right) + \mathbf{c}_t + \boldsymbol{\omega}_t \right) \quad (32)$$

$$= \max_{\mathbf{x}_t \in \mathcal{X}_t} \left( \mathbf{A}_{i,:}^{\text{out}} \left( \mathbf{A}_t + \mathbf{B}_t \boldsymbol{\Upsilon}_{i,:} \mathbf{C}_t^T \right) \right) \mathbf{x}_t + \mathbf{A}_{i,:}^{\text{out}} \left( \mathbf{B}_t \left( \boldsymbol{\Upsilon}_{i,:} s(\mathbf{A}_{i,:}^{\text{out}} \mathbf{B}_t \boldsymbol{\Upsilon}_{i,:} \mathbf{C}_t^T, \bar{\mathbf{v}}_t, \underline{\mathbf{v}}_t) + \boldsymbol{\Gamma}_{:,i} \right) + \mathbf{c}_t + s(\mathbf{A}_{i,:}^{\text{out}} \mathbf{B}_t, \bar{\boldsymbol{\omega}}_t, \underline{\boldsymbol{\omega}}_t) \right), \quad (33)$$

where (30) substitutes the definition of  $\pi_{:,i}^{UCL}$  from Theorem 4.1, (31) substitutes the observation from (2), (32) separates terms that depend on  $\mathbf{x}_t$ , and (33) introduces the worst-case realizations of process and measurement noise. Substituting  $\mathbf{M}_{i,:}^U, \mathbf{n}_i^U$  results in (20). The proof of the lower bound follows similarly.  $\square$

The optimization problems in Eqs. (20) and (21) have convex cost with convex constraints  $\mathbf{x}_t \in \mathcal{X}_t$  (e.g., polytope  $\mathcal{X}_t$ ). We solve the linear programs (LPs) with `cvxpy` [33],

$$\boldsymbol{\gamma}_{t+1,i}^U = \text{LP}(\mathbf{M}_{i,:}^U \mathbf{x}_t, \mathbf{A}^{\text{in}}, \mathbf{b}^{\text{in}}) + \mathbf{n}_i^U \quad (34)$$

$$\boldsymbol{\gamma}_{t+1,i}^L = \text{LP}(-\mathbf{M}_{i,:}^L \mathbf{x}_t, \mathbf{A}^{\text{in}}, \mathbf{b}^{\text{in}}) + \mathbf{n}_i^L. \quad (35)$$

### D. CONVERTING STATE CONSTRAINTS INTO REACHABLE SETS

#### 1) REACHABLE SETS AS $\ell_\infty$ -BALLS

Assume  $\mathcal{X}_0$  is an  $\ell_p$ -ball. Define  $\{p, \boldsymbol{\epsilon}, \mathbf{x}\}$  s.t.  $\mathcal{X}_0 \subseteq \mathcal{B}_p(\mathbf{x}, \boldsymbol{\epsilon})$  and let  $\bar{\mathcal{R}}_0(\mathcal{X}_0) = \mathcal{X}_0$ . Using the results of the previous section, use  $\mathbf{x}_{t=0} \in \mathcal{B}_p(\mathbf{x}, \boldsymbol{\epsilon})$  to compute  $(\boldsymbol{\gamma}_{1,k}^U, \boldsymbol{\gamma}_{1,k}^L)$  for each index of the state vector  $k \in [n_x]$ , specifying  $\mathbf{A}^{\text{out}} = \mathbf{I}_{n_x}$ . Recursively compute

$$\bar{\mathcal{R}}_{t+1}(\mathcal{X}_0) = \mathcal{B}_\infty \left( \frac{\boldsymbol{\gamma}_{t+1}^U + \boldsymbol{\gamma}_{t+1}^L}{2}, \frac{\boldsymbol{\gamma}_{t+1}^U - \boldsymbol{\gamma}_{t+1}^L}{2} \right). \quad (36)$$

#### 2) REACHABLE SETS AS POLYTOPES

Assume  $\mathcal{X}_0$  is an  $\ell_p$ -ball or polytope. Either define  $\{p, \boldsymbol{\epsilon}, \mathbf{x}\}$  s.t.  $\mathcal{X}_0 \subseteq \mathcal{B}_p(\mathbf{x}, \boldsymbol{\epsilon})$  or define  $\{\mathbf{A}^{\text{in}}, \mathbf{b}^{\text{in}}\}$  s.t.  $\mathcal{X}_t = \{\mathbf{x}_t | \mathbf{A}^{\text{in}} \mathbf{x}_t \leq \mathbf{b}^{\text{in}}\}$  and let  $\bar{\mathcal{R}}_0(\mathcal{X}_0) = \mathcal{X}_0$ . Using the results of the previous section, use  $\mathbf{x}_{t=0} \in \mathcal{B}_p(\mathbf{x}, \boldsymbol{\epsilon})$  or  $\{\mathbf{A}^{\text{in}}, \mathbf{b}^{\text{in}}\}$  to compute  $(\boldsymbol{\gamma}_{1,i}^U, \boldsymbol{\gamma}_{1,i}^L)$  for each index of output polytope facets  $i \in [m_{\text{out}}]$ , giving

$$\bar{\mathcal{R}}_{t+1}(\mathcal{X}_0) = \left\{ \mathbf{x}_t \mid \begin{bmatrix} \mathbf{A}^{\text{out}} \\ -\mathbf{A}^{\text{out}} \end{bmatrix} \mathbf{x}_t \leq \begin{bmatrix} \boldsymbol{\gamma}_{t+1}^U \\ -\boldsymbol{\gamma}_{t+1}^L \end{bmatrix} \right\}. \quad (37)$$

In both cases,  $\bar{\mathcal{R}}_t(\mathcal{X}_0) \supseteq \mathcal{R}_t(\mathcal{X}_0) \forall t \geq 0$ , so these  $\bar{\mathcal{R}}_t$  can be used to verify the original closed loop system (2).

### E. CLOSED-FORM SOLUTION

Rather than employing an LP solver as in (34), the optimization problem in (20) can be solved in closed-form when  $\mathcal{X}_t$  is described by an  $\ell_p$ -ball.

*Lemma 4.3:* In the special case of Theorem 4.2 where  $\mathcal{X}_t = \mathcal{B}_p(\mathbf{x}_t, \boldsymbol{\epsilon})$  for some  $p \in [1, \infty)$ , the following closed-form expressions are equivalent to Eqs. (20) and (21), respectively,

$$\boldsymbol{\gamma}_{t+1,k}^U = \|\boldsymbol{\epsilon} \odot \mathbf{M}_{i,:}^U\|_q + \mathbf{M}_{i,:}^U \mathbf{x}_t + \mathbf{n}_i^U \quad (38)$$

$$\mathbf{y}_{t+1,k}^L = -\|\epsilon \odot \mathbf{M}_{i,:}^L\|_q + \mathbf{M}_{i,:}^L \mathbf{x}_t + \mathbf{n}_i^L, \quad (39)$$

where  $1/p + 1/q = 1$  (e.g.,  $p = \infty, q = 1$ ).

*Proof:* The proof proceeds as in [34] and [8].

$$\mathbf{y}_{t+1,k}^U = \max_{\mathbf{x}_t \in \mathcal{B}_p(\mathbf{x}_t, \epsilon)} \mathbf{M}_{i,:}^U \mathbf{x}_t + \mathbf{n}_i^U \quad (40)$$

$$= \left( \max_{\zeta \in \mathcal{B}_p(\mathbf{0}, 1)} \mathbf{M}_{i,:}^U (\zeta \odot \epsilon) \right) + \mathbf{M}_{i,:}^U \mathbf{x}_t + \mathbf{n}_i^U \quad (41)$$

$$= \left( \max_{\zeta \in \mathcal{B}_p(\mathbf{0}, 1)} (\epsilon \odot \mathbf{M}_{i,:}^U \zeta) \right) + \mathbf{M}_{i,:}^U \mathbf{x}_t + \mathbf{n}_i^U \quad (42)$$

$$= \|\epsilon \odot \mathbf{M}_{i,:}^U\|_q + \mathbf{M}_{i,:}^U \mathbf{x}_t + \mathbf{n}_i^U, \quad (43)$$

with  $\odot$  denoting element-wise multiplication (Hadamard product). Recall that  $\mathbf{n}_i^U$  does not depend on  $\mathbf{x}_t$ . From (40) to (41), we substitute  $\mathbf{x}_t := \zeta \odot \epsilon + \mathbf{x}_t$ , to shift and re-scale the observation to within the unit ball around zero,  $\zeta \in \mathcal{B}_p(\mathbf{0}, 1)$ . The commutative property of the Hadamard product allows (41) to (42). The maximization in (42) is equivalent to a  $\ell_q$ -norm in (43) by the definition of the dual norm  $\|\mathbf{a}\|_q = \{\sup \mathbf{a}^T \mathbf{b} : \|\mathbf{b}\|_p \leq 1\}$  and the fact that the  $\ell_q$  norm is the dual of the  $\ell_p$  norm for  $p, q \in [1, \infty)$  (with  $1/p + 1/q = 1$ ). The proof of the lower bound follows similarly.  $\square$

---

#### Algorithm 1 Closed-Loop CROWN Propagator

---


**Input:** initial state set  $\mathcal{X}_0$ , trained NN control policy  $\pi$ , dynamics  $f$ , time horizon  $\tau$

**Output:** forward reachable set approximations  $\tilde{\mathcal{R}}_{1:\tau}(\mathcal{X}_0)$

- 1:  $\tilde{\mathcal{R}}_0(\mathcal{X}_0) \leftarrow \mathcal{X}_0$
  - 2: **for**  $t \in \{1, 2, \dots, \tau\}$  **do**
  - 3:  $\mathcal{Y}_{t-1} \leftarrow \text{possibleObs}(\tilde{\mathcal{R}}_{t-1}(\mathcal{X}_0), f)$
  - 4:  $\Psi, \Phi, \alpha, \beta \leftarrow \text{CROWN}(\pi, \mathcal{Y}_{t-1})$
  - 5: **for**  $i \in [m_{\text{out}}]$  or  $k \in [n_x]$  **do**
  - 6:  $\mathbf{y}_{t+1,i}^U \leftarrow (34)$  or (38)
  - 7:  $\mathbf{y}_{t+1,i}^L \leftarrow (35)$  or (39)
  - 8: **end for**
  - 9:  $\tilde{\mathcal{R}}_t(\mathcal{X}_0) \leftarrow \left\{ \mathbf{x} \mid \begin{bmatrix} \mathbf{A}^{\text{out}} \\ -\mathbf{A}^{\text{out}} \end{bmatrix} \mathbf{x} \leq \begin{bmatrix} \mathbf{y}_{t+1}^U \\ -\mathbf{y}_{t+1}^L \end{bmatrix} \right\}$  from (37)
  - 10:  $\tilde{\mathcal{R}}_t(\mathcal{X}_0) \leftarrow \mathcal{B}_\infty \left( \frac{\mathbf{y}_{t+1}^U + \mathbf{y}_{t+1}^L}{2}, \frac{\mathbf{y}_{t+1}^U - \mathbf{y}_{t+1}^L}{2} \right)$  from (36)
  - 11: **end for**
  - 12: **return**  $\tilde{\mathcal{R}}_{1:\tau}(\mathcal{X}_0)$
- 

#### F. ALGORITHM FOR COMPUTING FORWARD REACHABLE SETS

The proposed procedure for estimating forward reachable sets from  $\mathcal{X}_0$  for a neural feedback loop is provided in Algorithm 1 (the term ‘‘propagator’’ will be explained in § V). After initializing the zeroth forward reachable set as the initial set (Line 1), we recursively compute  $\tilde{\mathcal{R}}_t(\mathcal{X}_0)$  forward in time (Line 2). For each timestep, we determine which observations the NN could receive, then pass those through CROWN [10]



**FIGURE 4.** System architecture. A closed-loop propagator uses the trained NN control policy and dynamics to estimate reachable sets, and a closed-loop partitioner decides how to split the initial state set into pieces. This is the closed-loop extension of the architecture from our prior work [36].

(or another NN relaxation method, provided it returns the corresponding terms). Depending on whether polytope facets or  $\ell_p$ -balls are used, the upper/lower bounds are computed facet-by-facet or state-by-state. Then, the  $\mathbf{y}$  terms are converted into set representations as in § IV-D. After following this procedure for each timestep, the algorithm returns a sequence of outer-approximations of the forward reachable sets (Line 12).

#### V. TIGHTER REACHABLE SETS BY PARTITIONING THE INITIAL STATE SET

NN relaxation methods can be improved by partitioning the input set [35], particularly when the input set is large but of low dimension. Here, we achieve tighter bounds by splitting  $\mathcal{X}_0$  into subsets, computing  $N$ -step reachable sets for each of the subsets separately, then returning the union of all reachable sets from each subset.

Recall that [36] introduced an architecture composed of a *partitioner* and *propagator* for analyzing NNs in isolation. The partitioner is an algorithm to split the NN input set in an intelligent way (e.g., uniform gridding [35], simulation guidance [20], greedy simulation guidance [36]), and the propagator is an algorithm to estimate bounds on the NN outputs given a NN input set (e.g., IBP [7], Fast-Lin [8], CROWN [10], SDP [13]). The partitioner can query the propagator repeatedly to refine the estimated output set bounds. The choice and design of partitioners and propagators has important implications on the bound tightness and computational runtime.

This work extends the framework from [36] to closed-loop systems, as visualized in Fig. 4. In particular, CL-CROWN and Reach-SDP represent *closed-loop propagators* (given an initial state set, they compute reachable sets for a trained NN control policy and closed-loop dynamics), and in this section we extend partitioners discussed in [36] to *closed-loop partitioners*. Altogether, we call the nested architecture a *closed-loop analyzer*.

#### A. CLOSED-LOOP PARTITIONERS

The simplest closed-loop partitioner splits the input set into a uniform grid, as described in Algorithm 2. After splitting  $\mathcal{X}_0$  into  $\prod_{k=0}^{n_x} \mathbf{r}_k$  cells (Line 2), each cell of  $\mathcal{X}_0$  is passed through a closed-loop propagator, CLProp (e.g., CL-CROWN (Line 1)

**Algorithm 2** Closed-Loop Uniform Partitioner

**Input:** initial state set  $\mathcal{X}_0$ , trained NN control policy  $\pi$ , dynamics  $f$ , time horizon  $\tau$ , number of partition cells  $\mathbf{r}$   
**Output:** forward reachable set approximations  $\bar{\mathcal{R}}_{1:\tau}(\mathcal{X}_0)$

- 1:  $\bar{\mathcal{R}}_t \leftarrow \emptyset \quad \forall t \in \{1, \dots, \tau\}$
- 2:  $\mathcal{S} \leftarrow \text{partition}(\mathcal{X}_0, \mathbf{r})$
- 3: **for**  $\mathcal{X}' \in \mathcal{S}$  **do**
- 4:  $\bar{\mathcal{R}}'_{1:\tau} \leftarrow \text{CLProp}(\mathcal{X}', f, \pi, \tau)$
- 5:  $\bar{\mathcal{R}}_t \leftarrow \bar{\mathcal{R}}_t \cup \bar{\mathcal{R}}'_t \quad \forall t \in \{1, \dots, \tau\}$
- 6: **end for**
- 7: **return**  $\bar{\mathcal{R}}_{1:\tau}(\mathcal{X}_0)$

or Reach-SDP [21]). For each timestep, the returned estimate of the reachable set for all of  $\mathcal{X}_0$  is thus the union of all reachable sets for cells of the  $\mathcal{X}_0$  partition.

There are several reasons to use an iterative partitioner rather than the uniform partitioner presented before, as discussed in [20], [36]. Thus, the closed-loop variant of the greedy simulation-guided partitioner from [36] is described in Algorithm 3. After acquiring  $N$  Monte Carlo samples from  $\mathcal{X}_0$  (Line 1), those sampled points are simulated forward in time according to the dynamics and control policy (Line 2). The extrema of these samples can be used to define the sampled reachable sets at each timestep (Line 3). Then, the full initial state set  $\mathcal{X}_0$  is passed through a closed-loop propagator (Line 4) and the input and reachable set estimates are added to a stack,  $M$ .

An iterative process is used to refine the reachable set estimates from the full  $\mathcal{X}_0$ . The element with largest distance from the sampled reachable set estimate is popped from  $M$  (e.g., use the same idea from Fig. 3a of [36] on the final timestep's reachable set). If the chosen element's reachable set is within the sampled reachable set for all timesteps, there is no need to further refine that cell and its reachable set can be added to set that will eventually be returned. Otherwise, the input set of that element is bisected (Line 11) and each of the two input sets are passed through CLProp and added to  $M$ . Termination conditions, such as an empty  $M$  or number of CLProp calls or a timeout can be implemented depending on the application. Finally, the remaining elements of the stack are added to the set that will be returned.

**VI. HANDLING NONLINEARITIES**

**A. ACCOUNTING FOR CONTROL LIMITS,  $\mathcal{u}_t$**

The key terms in Theorem 4.1 can be modified to account for control input constraints, as

$$\pi_{:,i}^{UCL}(\mathbf{y}_t) = \text{Proj}_{\mathcal{U}_t}(\mathbf{\Upsilon}_{i,:} \mathbf{y}_t + \mathbf{\Gamma}_{:,i}) \quad (44)$$

$$\pi_{:,i}^{LCL}(\mathbf{y}_t) = \text{Proj}_{\mathcal{U}_t}(\mathbf{\Xi}_{i,:} \mathbf{y}_t + \mathbf{\Delta}_{:,i}), \quad (45)$$

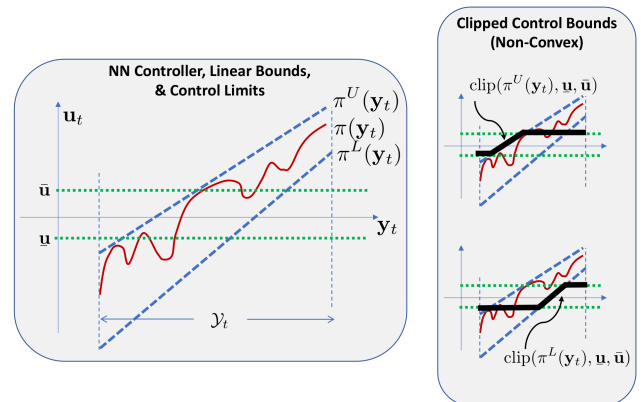
A common example is box control input constraints. The element-wise control input is,

$$\pi_{j,i}^{UCL}(\mathbf{y}_t) = \begin{cases} \text{clip}(\pi_j^U(\mathbf{y}_t), \underline{\mathbf{u}}_j, \bar{\mathbf{u}}_j), & \text{if } \mathbf{A}_{i,:}^{\text{out}} \mathbf{B}_{t,:j} \geq 0 \\ \text{clip}(\pi_j^L(\mathbf{y}_t), \underline{\mathbf{u}}_j, \bar{\mathbf{u}}_j), & \text{otherwise,} \end{cases} \quad (46)$$

**Algorithm 3** Closed-Loop Greedy Sim-Guided Partitioner

**Input:** initial state set  $\mathcal{X}_0$ , number of MC samples  $N$ , trained NN control policy  $\pi$ , dynamics  $f$ , time horizon  $\tau$   
**Output:** forward reachable set approximations  $\bar{\mathcal{R}}_{1:\tau}(\mathcal{X}_0)$

- 1:  $\mathcal{X}_0^{\text{MC}} \stackrel{N \text{ i.i.d.}}{\sim} \text{Unif}(\mathcal{X}_0)$
- 2:  $\mathcal{X}'_{t+1} \leftarrow f(\mathcal{X}_t^{\text{MC}}; \pi) \quad \forall t \in \{1, \dots, \tau\}$
- 3:  $\mathcal{R}'_t \leftarrow \text{extrema of } \mathcal{X}'_t \quad \forall t \in \{1, \dots, \tau\}$
- 4:  $\bar{\mathcal{R}}_{1:\tau} \leftarrow \text{CLProp}(\mathcal{X}_0, f, \pi, \tau)$
- 5:  $M \leftarrow \{(\mathcal{X}_0, \bar{\mathcal{R}}_{1:\tau})\}$
- 6: **while**  $M \neq \emptyset$  or termination condition **do**
- 7:  $(\mathcal{X}', \mathcal{R}'_{1:\tau}) \stackrel{M, \text{pop}}{\leftarrow} \text{argmax}_{(\mathcal{X}'', \mathcal{R}''_{1:\tau}) \in M} d(\mathcal{R}_{1:\tau}^{\text{MC}}, \mathcal{R}''_{1:\tau})$
- 8: **if**  $\mathcal{R}'_t \subseteq \bar{\mathcal{R}}_t \quad \forall t \in \{1, \dots, \tau\}$  **then**
- 9:  $\bar{\mathcal{R}}_t \leftarrow \bar{\mathcal{R}}_t \cup \mathcal{R}'_t \quad \forall t \in \{1, \dots, \tau\}$
- 10: **else**
- 11:  $\mathcal{X}^1, \mathcal{X}^2 \leftarrow \text{Bisect}(\mathcal{X}')$
- 12:  $\bar{\mathcal{R}}_{1:\tau}^1 \leftarrow \text{CLProp}(\mathcal{X}^1, f, \pi, \tau)$
- 13:  $\bar{\mathcal{R}}_{1:\tau}^2 \leftarrow \text{CLProp}(\mathcal{X}^2, f, \pi, \tau)$
- 14:  $M \leftarrow M \cup \{(\mathcal{X}^1, \bar{\mathcal{R}}_{1:\tau}^1), (\mathcal{X}^2, \bar{\mathcal{R}}_{1:\tau}^2)\}$
- 15: **end if**
- 16: **end while**
- 17:  $\bar{\mathcal{R}}_t \leftarrow \bar{\mathcal{R}}_t \cup (\cup_{(\mathcal{X}'', \mathcal{R}''_{1:\tau}) \in M} \mathcal{R}''_t) \quad \forall t \in \{1, \dots, \tau\}$
- 18: **return**  $\bar{\mathcal{R}}_{1:\tau}(\mathcal{X}_0)$



**FIGURE 5.** Control limits. The NN control policy (red) is bounded by  $\pi^U, \pi^L$  from CROWN (blue), but none of these respect the control limits (green). While adding a clip could be non-convex (right block), § VI-A borrows the idea from [21] to append a few layers to the NN to handle control limits.

where clip saturates the control if it exceeds the limits. However, clip could be non-convex depending on the domain of  $\mathbf{x}_t$  (and violates DCP rules in cvxpy [33] regardless), as visualized in Fig. 5.

We instead utilize the insights from [21] that clip looks like two ReLUs stitched together (one is inverted). In particular, we make the following changes to the NN:

- 1) Subtract  $\underline{\mathbf{u}}_t$  from final layer's bias,  $\mathbf{b}^{(m)} \leftarrow \mathbf{b}^{(m)} - \underline{\mathbf{u}}_t$
- 2) Append ReLU activation
- 3) Append linear layer,  $\mathbf{W}^{(m+1)} \leftarrow -I_{n_u}, \mathbf{b}^{(m+1)} \leftarrow \bar{\mathbf{u}}_t - \underline{\mathbf{u}}_t$
- 4) Append ReLU activation
- 5) Append linear layer,  $\mathbf{W}^{(m+2)} \leftarrow -I_{n_u}, \mathbf{b}^{(m+2)} \leftarrow \bar{\mathbf{u}}_t$ .

All of the preceding results apply to this modified NN for a system with box control input constraints. Note that adding these ReLUs does not limit which types of activations can be used in other layers of the NN.

While these additional layers exactly mimic clip for any particular  $\mathbf{y}_t$ , they can induce additional conservatism because the additional ReLUs must be relaxed as well. Additionally, while CROWN provides several options for the slope of the lower bound for a ReLU relaxation, note that the two additional ReLUs should be relaxed with a zero-slope lower bound to ensure no values exceed the control limits.

## B. EXTENSION TO POLYNOMIAL DYNAMICS

This section extends the above work to systems with polynomial dynamics. Consider the system with dynamics

$$\begin{aligned}\mathbf{x}_{t+1} &= \mathbf{A}_t \mathbf{x}_t + \bar{f}(\mathbf{x}_t) + \mathbf{B}_t \mathbf{u}_t + \mathbf{c}_t + \boldsymbol{\omega}_t \\ \mathbf{y}_t &= \mathbf{C}_t^T \mathbf{x}_t + \mathbf{v}_t,\end{aligned}\quad (47)$$

where  $\bar{f}(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}^n$  is a polynomial and the other terms are defined in the same way as in (1). As a result, the counterpart of (20) and (21) in Lemma 4.2 are respectively as follows

$$\gamma_{t+1,j}^U = \max_{\mathbf{x}_t \in \mathcal{X}_t} \mathbf{M}_{i,j}^U \mathbf{x}_t + f_j(\mathbf{x}_t) + \mathbf{n}_i^U \quad (48)$$

$$\gamma_{t+1,j}^L = \min_{\mathbf{x}_t \in \mathcal{X}_t} \mathbf{M}_{i,j}^L \mathbf{x}_t + f_j(\mathbf{x}_t) + \mathbf{n}_i^L, \quad (49)$$

where  $f_j(\mathbf{x}_t) = \mathbf{A}_{i,j}^{\text{out}} \bar{f}(\mathbf{x}_t)$ . Because of the polynomial term  $f_j(\mathbf{x}_t)$  in (48) and (49), these objectives can be nonconvex and NP-Hard to solve for global optimality [37]. Thus, we adopt a principled convex relaxation technique [38] to compute a guaranteed upper and lower bound for (48) and (49), respectively.

We first transform the polynomial term  $f(\mathbf{x}_t)$  to a linear or quadratic term by introducing new variables  $\mathbf{s} \in \mathbb{R}^{n_s}$  and new quadratic constraints. For example,  $f(\mathbf{x}_t) = \mathbf{x}_{t;1} \mathbf{x}_{t;2} \mathbf{x}_{t;3} \mathbf{x}_{t;4}$  is equivalent to  $f(\mathbf{s}) = \mathbf{s}_1 \mathbf{s}_2$ , with  $\mathbf{s}_1 = \mathbf{x}_{t;1} \mathbf{x}_{t;2}$  and  $\mathbf{s}_2 = \mathbf{x}_{t;3} \mathbf{x}_{t;4}$ . Then (48) can be transformed as a quadratically constrained quadratic programming (QCQP) in the following

$$\gamma^U = \max_{\mathbf{x} \in \mathcal{X}, \bar{\mathbf{s}} \in \mathcal{Q}} \mathbf{M}^U \mathbf{x} + \bar{\mathbf{s}}^T \mathbf{Q}_0 \bar{\mathbf{s}} + \mathbf{b}_0^T \bar{\mathbf{s}} + \mathbf{n}^U \quad (50)$$

where the subscripts are dropped here for clarity,  $\bar{\mathbf{s}} = [\mathbf{x}^T, \mathbf{s}^T]^T \in \mathbb{R}^{n+n_s}$ ,  $\mathbf{Q}_0 \in \mathcal{S}^{n+n_s}$ ,  $\mathbf{b}_0 \in \mathbb{R}^{n+n_s}$  and  $\mathcal{Q} = \{\bar{\mathbf{s}} \mid \bar{\mathbf{s}}^T \mathbf{Q}_i \bar{\mathbf{s}} + \mathbf{b}_i^T \bar{\mathbf{s}} + c_i = 0, \mathbf{Q}_i \in \mathcal{S}^{n+n_s}, \mathbf{b}_i \in \mathbb{R}^{n+n_s}, c_i \in \mathbb{R}, \forall i = 1, \dots\}$  is formulated by a sequence of quadratic constraints. For QCQP, the semidefinite relaxation can give the tightest bounds among existing convex relaxation techniques [38]. The semidefinite relaxation of (50) can be formulated as

$$\gamma^U = \max_{\mathbf{x} \in \mathcal{X}, \mathbf{S} \in \bar{\mathcal{Q}}, \mathbf{S} \succeq 0} \mathbf{M}^U \mathbf{x} + \text{Tr}(\bar{\mathbf{Q}}_0 \mathbf{S}) + \mathbf{n}^U \quad (51)$$

where  $\mathbf{S} = [1, \bar{\mathbf{s}}^T]^T [1, \bar{\mathbf{s}}^T] \in \mathcal{S}^{1+n+n_s}$ ,  $\bar{\mathbf{Q}}_i = [0, \mathbf{b}_i^T/2; \mathbf{b}_i/2, \mathbf{Q}_i]$ ,  $i = 0, \dots$  and  $\bar{\mathcal{Q}} = \{\mathbf{S} \mid \text{Tr}(\bar{\mathbf{Q}}_i \mathbf{S}) + c_i = 0, i = 1, \dots\}$ . (51) is equivalent to the original QCQP (50) if an extra nonconvex constraint  $\text{rank}(\mathbf{S}) = 1$  is added.

As a result, solving the convex program (51) will yield a guaranteed upper bound for (48). The lower bound for (49) follows similarly.

## VII. BACKWARD REACHABILITY ANALYSIS

So far, this paper has considered the *forward* reachability problem: starting from a set  $\mathcal{X}_0$ , what are all the possible states the system could occupy at time  $t$ ? In this section, we discuss the opposite problem of *backward* reachability: given a target set  $\mathcal{X}_T$ , from which states will the system end up in the target set? Depending on whether the initial or target state sets are known a priori and/or change frequently, one of these two paradigms would be a better fit for a given application. Furthermore, the target set could represent a goal set or avoid/collision set – either way, this section provides a framework for guaranteeing the system will reach/avoid the target set.

Traditionally, it is straightforward to switch between forward and backward reachability analysis through a change of variables [15], [22]. However, including a NN in the analysis adds new challenges. In particular, a key challenge with propagating sets *backward* through NNs is that common NN activations have finite range (e.g.,  $\text{ReLU}(x) = 0$  could correspond to any  $x \leq 0$ ), which causes the sets to quickly explode to  $\infty$ . Moreover, even using an infinite range activation (e.g., Leaky ReLU), the weight matrices may not be invertible (e.g., singular, insufficient rank). Our approach is able to avoid these issues and still leverage forward propagation tools (e.g., CROWN [10]) despite thinking backward in time.

Recent related work has developed specific invertible-by-design NN architectures [39] and modifications to the training procedure to regularize for invertibility [40]. In contrast, our approach can be applied to any NN for which an affine relaxation can be computed, i.e., the same broad class of NN architectures (e.g., feedforward NNs with general activation functions) with arbitrary training processes that CROWN (or recent extensions [41]–[43]) operates on.

### A. BACKREACHABLE & BACKPROJECTION SETS


Fig. 6 illustrates the important concepts for this analysis. Three different sets (left) contain states that lead to the target set  $\mathcal{X}_T$  under different control inputs, as defined below. The nomenclature of these sets is motivated by [44], but with slightly different definitions.

The first of these sets is the *backward reachable set*,

$$\begin{aligned}\mathcal{R}_{t-1}^B(\mathcal{X}_T) &= \{\mathbf{x} \mid \exists \mathbf{u} \in \mathcal{U}_{t-1} \text{ s.t.} \\ &\quad \mathbf{A}_{t-1} \mathbf{x} + \mathbf{B}_{t-1} \mathbf{u} + \mathbf{c}_{t-1} \in \mathcal{R}_t^B(\mathcal{X}_T)\} \forall t \leq 0 \\ \mathcal{R}_0^B(\mathcal{X}_T) &= \mathcal{X}_T,\end{aligned}\quad (52)$$

which consists of states  $\mathbf{x}_t$  for which some admissible control trajectory,  $\mathbf{u}_{t:T}$ , takes the system to  $\mathcal{X}_T$ . To simplify the notation, we assume no noise (i.e.,  $\boldsymbol{\omega}_t = \bar{\boldsymbol{\omega}}_t = \mathbf{v}_t = \bar{\mathbf{v}}_t = \mathbf{0}$ ) and perfect observations,  $\mathbf{C}_t^T = \mathbf{I} \Rightarrow \mathbf{y}_t = \mathbf{x}_t$  in this section. Backward reachable sets are particularly useful when the controller is truly a “black box” or has not yet been defined





**FIGURE 6.** Backreachable, backprojection, and target sets. Given a target set,  $\mathcal{X}_T$ , the backreachable set  $\mathcal{R}_{-1}^B(\mathcal{X}_T)$  contains all states for which some control exists to move the system to  $\mathcal{X}_T$  in one timestep. The backprojection set,  $\mathcal{P}_{-1}(\mathcal{X}_T)$  contains all states for which the NN controller leads the system to  $\mathcal{X}_T$ . Inner-approximation  $\bar{\mathcal{P}}_{-1}(\mathcal{X}_T)$  contains all states for which all controls that the relaxed NN could apply lead the system to  $\mathcal{X}_T$ .

(since the definition only uses the fact that  $\mathbf{u} \in \mathcal{U}$ , rather than  $\pi$ ).

In the case of a trained NN control policy, we can be more precise about the closed loop system’s safety properties. In particular, for the closed-loop system  $f$  from (2), we denote  $\mathcal{P}_{-t}(\mathcal{X}_T)$  the *backprojection set*, which contains all states that will lead to a given target set  $\mathcal{X}_T \subseteq \mathbb{R}^{n_x}$  in exactly  $t$  timesteps, as defined by the recursion

$$\begin{aligned} \mathcal{P}_{t-1}(\mathcal{X}_T) &= \{\mathbf{x} \mid \mathbf{A}_{t-1}\mathbf{x} + \mathbf{B}_{t-1}\pi(\mathbf{x}) + \mathbf{c}_{t-1} \in \mathcal{P}_t(\mathcal{X}_T)\} \\ &\quad \forall t \leq 0 \\ \mathcal{P}_0(\mathcal{X}_T) &= \mathcal{X}_T. \end{aligned} \tag{53}$$

Recall that in § III-C, we computed over-approximations of the reachable sets because the NN makes exact reachability analysis computationally challenging. Analogously, we will compute *under-approximations* of the backprojection sets,  $\bar{\mathcal{P}}_{-t}(\mathcal{X}_T) \subseteq \mathcal{P}_{-t}(\mathcal{X}_T)$ . Under-approximations are still useful because if  $\bar{\mathcal{P}}_{-t}(\mathcal{X}_T) \supseteq \mathcal{X}_0$ , then  $\mathcal{P}_{-t}(\mathcal{X}_T) \supseteq \mathcal{X}_0$  as well, meaning for every state in  $\mathcal{X}_0$ , the control policy  $\pi$  is guaranteed to lead the system to  $\mathcal{X}_T$  in  $T$  timesteps. Thus, it is desirable for  $\bar{\mathcal{P}}_{-t}(\mathcal{X}_T)$  to be as close to  $\mathcal{P}_{-t}(\mathcal{X}_T)$  as possible.

**B. DERIVING BACKPROJECTION SETS**

We first provide a summary of the proposed approach:

- 1) Ignoring the NN, find (hyper-)rectangular bounds on the backreachable set (this is a superset of the backprojection set associated with the particular NN controller)
- 2) Within this region, relax the NN controller to acquire upper/lower affine bounds on control
- 3) Compute the states that will lead to the target set for every control effort within the upper/lower bounds

This last step provides an under-approximation of the backprojection set, which is the set of interest.

The following lemma provides polytope bounds on  $\bar{\mathcal{P}}_t(\mathcal{X}_T)$  for a single timestep, which is the key component of the recursive algorithm introduced in the next section.

*Lemma 7.1:* Given an  $m$ -layer NN control policy  $\pi : \mathbb{R}^{n_y} \rightarrow \mathbb{R}^{n_u}$ , closed-loop dynamics  $f : \mathbb{R}^{n_x} \times \Pi \rightarrow \mathbb{R}^{n_x}$  as in Eqs. (1) and (2), and bounds on the state set,  $\underline{\mathbf{x}}_{t+1} \leq \mathbf{x}_{t+1} \leq \bar{\mathbf{x}}_{t+1}$ , the following polytope bounds on the previous state set hold:

$$\begin{bmatrix} \mathbf{A}_t + \mathbf{Z}^{\mathbf{B}\Psi\Phi} \\ -(\mathbf{A}_t + \mathbf{Z}^{\mathbf{B}\Phi\Psi}) \\ \mathbf{I}_{n_x} \\ -\mathbf{I}_{n_x} \end{bmatrix} \mathbf{x}_t \leq \begin{bmatrix} \bar{\mathbf{x}}_{t+1} - (\mathbf{z}^{\mathbf{B}\alpha\beta} + \mathbf{c}_t) \\ -\underline{\mathbf{x}}_{t+1} + \mathbf{z}^{\mathbf{B}\beta\alpha} + \mathbf{c}_t \\ \bar{\mathbf{x}}_t \\ -\underline{\mathbf{x}}_t \end{bmatrix}, \tag{54}$$

where  $\mathbf{Z}^{\mathbf{B}\Phi\Psi}, \mathbf{Z}^{\mathbf{B}\Psi\Phi} \in \mathbb{R}^{n_x \times n_y}$  and  $\mathbf{z}^{\mathbf{B}\beta\alpha}, \mathbf{z}^{\mathbf{B}\alpha\beta} \in \mathbb{R}^{n_x}$  are defined  $\forall k \in [n_x]$  (by row/element),

$$\left[ \mathbf{Z}^{\mathbf{B}\Phi\Psi} \right]_{k,:} = \mathbf{B}_{t;k,:} s(\mathbf{B}_{t;k,:}, \Phi, \Psi) \tag{55}$$

$$\left[ \mathbf{Z}^{\mathbf{B}\Psi\Phi} \right]_{k,:} = \mathbf{B}_{t;k,:} s(\mathbf{B}_{t;k,:}, \Psi, \Phi) \tag{56}$$

$$\left[ \mathbf{z}^{\mathbf{B}\beta\alpha} \right]_k = \mathbf{B}_{t;k,:} s(\mathbf{B}_{t;k,:}, \beta, \alpha) \tag{57}$$

$$\left[ \mathbf{z}^{\mathbf{B}\alpha\beta} \right]_k = \mathbf{B}_{t;k,:} s(\mathbf{B}_{t;k,:}, \alpha, \beta), \tag{58}$$

and  $\underline{\mathbf{x}}_t, \bar{\mathbf{x}}_t$  are computed elementwise by solving the LPs in Eqs. (59) to (64).

*Proof:* Given dynamics from Eqs. (1) and (2), solve the following optimization problems for each state  $k \in [n_x]$ ,

$$\bar{\mathbf{x}}_{t;k} = \max_{\mathbf{x}_t, \mathbf{u}_t} \mathbf{x}_{t;k} \tag{59}$$

$$\text{s.t. } \mathbf{A}_t \mathbf{x}_t + \mathbf{B}_t \mathbf{u}_t + \mathbf{c}_t \in \mathcal{X}_{t+1} \tag{60}$$

$$\mathbf{u}_t \in \mathcal{U}_t \tag{61}$$

$$\underline{\mathbf{x}}_{t;k} = \min_{\mathbf{x}_t, \mathbf{u}_t} \mathbf{x}_{t;k} \tag{62}$$

$$\text{s.t. } \mathbf{A}_t \mathbf{x}_t + \mathbf{B}_t \mathbf{u}_t + \mathbf{c}_t \in \mathcal{X}_{t+1} \tag{63}$$

$$\mathbf{u}_t \in \mathcal{U}_t, \tag{64}$$

which provides a (hyper-)rectangular outer bound ( $\underline{\mathbf{x}}_t \leq \mathbf{x}_t \leq \bar{\mathbf{x}}_t$ ) on the *backreachable set*. Note that this is a LP for convex  $\mathcal{X}_{t+1}, \mathcal{U}_t$  as we will use here.

Given that bound,  $\underline{\mathbf{x}}_t \leq \mathbf{x}_t \leq \bar{\mathbf{x}}_t$ , Theorem 3.1 provides  $\Psi, \Phi, \alpha, \beta$  which lead to the following inequalities based on Theorem 4.1 with  $\mathbf{A}^{\text{out}} = \mathbf{I}_{n_x}, \forall k \in [n_x]$ ,

$$\begin{aligned} &\left[ \mathbf{A}_t \mathbf{x}_t + \mathbf{B}_t (s(\mathbf{B}_{t;k,:}, \Phi, \Psi) \mathbf{y}_t + s(\mathbf{B}_{t;k,:}, \beta, \alpha)) + \mathbf{c}_t \right]_k \\ &\leq \mathbf{x}_{t+1;k} \\ &\leq \left[ \mathbf{A}_t \mathbf{x}_t + \mathbf{B}_t (s(\mathbf{B}_{t;k,:}, \Psi, \Phi) \mathbf{y}_t + s(\mathbf{B}_{t;k,:}, \alpha, \beta)) + \mathbf{c}_t \right]_k. \end{aligned} \tag{65}$$

After substituting in for  $\mathbf{y}_t$ , grouping  $\mathbf{x}_t$  terms together, and using Eqs. (55) to (58),

$$\begin{aligned} &\left( \mathbf{A}_t + \mathbf{Z}^{\mathbf{B}\Phi\Psi} \right) \mathbf{x}_t + \mathbf{z}^{\mathbf{B}\beta\alpha} + \mathbf{c}_t \\ &\leq \mathbf{x}_{t+1} \\ &\leq \left( \mathbf{A}_t + \mathbf{Z}^{\mathbf{B}\Psi\Phi} \right) \mathbf{x}_t + \mathbf{z}^{\mathbf{B}\alpha\beta} + \mathbf{c}_t. \end{aligned} \tag{66}$$

To ensure no NN control pushes the system into a state beyond  $[\underline{\mathbf{x}}_{t+1}, \bar{\mathbf{x}}_{t+1}]$ , the following inequalities hold,

$$\underline{\mathbf{x}}_{t+1} \leq \left( \mathbf{A}_t + \mathbf{Z}^{\mathbf{B}\Phi\Psi} \right) \mathbf{x}_t + \mathbf{z}^{\mathbf{B}\beta\alpha} + \mathbf{c}_t \tag{67}$$

**Algorithm 4** Backprojection Set Estimation

**Input:** target state set  $\mathcal{X}_T$ , trained NN control policy  $\pi$ , time horizon  $\tau$ , partition parameter  $\mathbf{r}$

**Output:** backprojection set approximations  $\bar{\mathcal{P}}_{-\tau:0}(\mathcal{X}_T)$

```

1:  $\bar{\mathcal{P}}_0(\mathcal{X}_T) \leftarrow \mathcal{X}_T$ 
2: for  $t$  in  $\{-1, -2, \dots, -\tau\}$  do
3:    $\bar{\mathcal{P}}_t(\mathcal{X}_T) \leftarrow \emptyset$ 
4:    $\bar{\mathcal{R}}_t^B(\mathcal{X}_T) = [\underline{\mathbf{x}}_t, \bar{\mathbf{x}}_t] \leftarrow \text{backreach}(\bar{\mathcal{P}}_{t+1}(\mathcal{X}_T), \mathcal{U}_t)$ 
5:    $\mathcal{S} \leftarrow \text{partition}([\underline{\mathbf{x}}_t, \bar{\mathbf{x}}_t], \mathbf{r})$ 
6:   for  $[\underline{\mathbf{x}}_t, \bar{\mathbf{x}}_t]$  in  $\mathcal{S}$  do
7:      $\Psi, \Phi, \alpha, \beta \leftarrow \text{CROWN}(\pi, [\underline{\mathbf{x}}_t, \bar{\mathbf{x}}_t])$ 
8:      $\mathbf{P} \leftarrow \begin{bmatrix} \mathbf{A}_t + \mathbf{Z}^B \Psi \Phi \\ -(\mathbf{A}_t + \mathbf{Z}^B \Phi \Psi) \\ \mathbf{I}_{n_x} \\ -\mathbf{I}_{n_x} \end{bmatrix}$ 
9:      $\mathbf{p} \leftarrow \begin{bmatrix} \bar{\mathbf{x}}_{t+1} - (\mathbf{z}^B \alpha \beta + \mathbf{c}_t) \\ -\underline{\mathbf{x}}_{t+1} + \mathbf{z}^B \beta \alpha + \mathbf{c}_t \\ \bar{\mathbf{x}}_t \\ -\underline{\mathbf{x}}_t \end{bmatrix}$ 
10:     $\mathcal{A} \leftarrow \{\mathbf{x} \mid \mathbf{P}\mathbf{x} \leq \mathbf{p}\}$ 
11:     $\bar{\mathcal{P}}_t(\mathcal{X}_T) \leftarrow \bar{\mathcal{P}}_t(\mathcal{X}_T) \cup \mathcal{A}$ 
12:  end for
13: end for
14: return  $\bar{\mathcal{P}}_{-\tau:0}(\mathcal{X}_T)$ 

```

$$(\mathbf{A}_t + \mathbf{Z}^B \Psi \Phi) \mathbf{x}_t + \mathbf{z}^B \alpha \beta + \mathbf{c}_t \leq \bar{\mathbf{x}}_{t+1}. \quad (68)$$

Written in matrix form,

$$\begin{bmatrix} \mathbf{A}_t + \mathbf{Z}^B \Psi \Phi \\ -(\mathbf{A}_t + \mathbf{Z}^B \Phi \Psi) \end{bmatrix} \mathbf{x}_t \leq \begin{bmatrix} \bar{\mathbf{x}}_{t+1} - (\mathbf{z}^B \alpha \beta + \mathbf{c}_t) \\ -\underline{\mathbf{x}}_{t+1} + \mathbf{z}^B \beta \alpha + \mathbf{c}_t \end{bmatrix}. \quad (69)$$

To arrive at (54), append  $\underline{\mathbf{x}}_t \leq \mathbf{x}_t \leq \bar{\mathbf{x}}_t$  to (69).  $\square$

*Corollary 7.2:* The set  $\bar{\mathcal{P}}_{-1}(\mathcal{X}_T)$  composed of all  $\mathbf{x}_t$  that satisfy (54) is a subset of the one-step backprojection set,  $\mathcal{P}_{-1}(\mathcal{X}_T)$ .

*Proof:* Suppose  $\mathbf{x}_t$  satisfies (54), i.e.,  $\mathbf{x}_t \in \bar{\mathcal{P}}_{-1}(\mathcal{X}_T)$ . Thus,  $\mathbf{x}_t$  also satisfies the inequalities (67) and (68), and thus also (66). Since  $\mathbf{x}_{t+1} = f(\mathbf{x}_t; \pi)$ ,  $\mathbf{x}_t$  also satisfies  $\underline{\mathbf{x}}_{t+1} \leq f(\mathbf{x}_t; \pi) \leq \bar{\mathbf{x}}_{t+1}$ , which means that  $\mathbf{x}_t \in \mathcal{P}_{-1}(\mathcal{X}_T)$  as well. Thus,  $\bar{\mathcal{P}}_{-1}(\mathcal{X}_T) \subseteq \mathcal{P}_{-1}(\mathcal{X}_T)$ .  $\square$

**C. ALGORITHM FOR COMPUTING BACKPROJECTION SETS**

Theorem 7.1 provides a formula for computing  $\bar{\mathcal{P}}_{-1}(\mathcal{X}_T)$ , i.e., an under-approximation of the backprojection set for a single timestep. This section extends that idea to enable computing  $\bar{\mathcal{P}}_{-\tau:0}(\mathcal{X}_T)$  for a time horizon  $\tau$  and also describes how partitioning can help improve the results, particularly when (54) is the empty set. The proposed procedure is summarized in Algorithm 4.

After initializing the zeroth backprojection set as the target set (Line 1), we recursively compute  $\bar{\mathcal{P}}_t(\mathcal{X}_T)$  backward in time (Line 2). For each timestep, we first compute an over-approximation on the backreachable set,  $\bar{\mathcal{R}}_t^B(\mathcal{X}_T)$ ,

by solving the LPs in Eqs. (59) to (60) and Eqs. 62 and (63) (Line 4).

Since  $\bar{\mathcal{R}}_t^B(\mathcal{X}_T) \supseteq \mathcal{P}_t(\mathcal{X}_T)$ , we can relax the NN over  $\bar{\mathcal{R}}_t^B(\mathcal{X}_T)$ , and that relaxation will also hold for all states in  $\mathcal{P}_t(\mathcal{X}_T)$  and  $\bar{\mathcal{P}}_t(\mathcal{X}_T)$ . However,  $\bar{\mathcal{R}}_t^B(\mathcal{X}_T)$  could be large, which could lead to a loose NN relaxation. Thus, instead of analyzing  $\bar{\mathcal{R}}_t^B(\mathcal{X}_T)$  as a single set, we uniformly partition the set  $\bar{\mathcal{R}}_t^B(\mathcal{X}_T)$  into  $\mathbf{r} \in \mathbb{N}^{n_x}$  cells (defined per dimension).

For each cell in the partition, we have a (hyper-)rectangular set of states,  $[\underline{\mathbf{x}}_t, \bar{\mathbf{x}}_t]_i$ , that could be passed as inputs to the NN controller. The NN can be relaxed over this set using CROWN (Theorem 3.1) or a similar type of algorithm. In particular, we extract  $\Psi, \Phi, \alpha, \beta$  from CROWN, which define  $\mathbf{P}$  and  $\mathbf{p}$ , the constraints in (54) (Line 8 and 9). These constraints define a polytope of states that are guaranteed to lead to  $\bar{\mathcal{P}}_{t+1}(\mathcal{X}_T)$  at the next timestep (for any control within CROWN's bounds), so those states should be added to  $\bar{\mathcal{P}}_t(\mathcal{X}_T)$  (Line 11). Note that it is possible, despite partitioning, that  $\mathcal{A} = \emptyset$ . After following this procedure for each timestep, the algorithm returns a sequence of inner-approximations of the backprojection sets (Line 14).

**VIII. EXPERIMENTAL RESULTS**

This section demonstrates our convex reachability analysis tool, Reach-LP, on simulated scenarios. We first show an example verification task and quantify the improvement in runtime vs. bound tightness over the state-of-the-art [21] for a double integrator system. We then apply the algorithm on a 6D quadrotor model subject to multiple sources of noise.

**A. DOUBLE INTEGRATOR**

Consider the LTI double integrator system from [21],

$$\mathbf{x}_{t+1} = \underbrace{\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}}_{\mathbf{A}_t} \mathbf{x}_t + \underbrace{\begin{bmatrix} 0.5 \\ 1 \end{bmatrix}}_{\mathbf{B}_t} \mathbf{u}_t, \quad (70)$$

with  $\mathbf{c}_t = 0$ ,  $\mathbf{C}_t = \mathbf{I}_2$  and no noise, discretized with sampling time  $t_s = 1s$ . As in [21], we implemented a linear MPC with prediction horizon  $N_{MPC} = 10$ , weighting matrices  $\mathbf{Q} = \mathbf{I}_2$ ,  $\mathbf{R} = 1$ , and terminal weighting matrix  $\mathbf{P}_\infty$  synthesized from the discrete-time Riccati equation, subject to state constraints  $\mathcal{A}^C = [-5, 5] \times [-1, 1]$  and input constraint  $\mathbf{u}_t \in [-1, 1] \forall t$ . We used MPC to generate 2420 samples of state and input pairs then trained a NN with Keras [45] for 20 epochs with batch size 32.

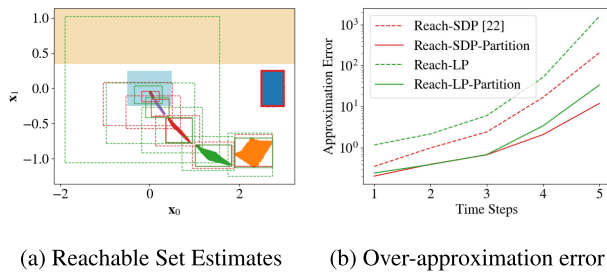
**B. COMPARISON WITH BASELINE**

Table 1 and Fig. 7 compare several algorithms on the double integrator system using a NN with [5,5] neurons and ReLU activations.<sup>1</sup> The key takeaway is that Reach-LP-Partition provides a  $5\times$  improvement in reachable set tightness over the prior state-of-the-art, Reach-SDP [21] (which does not use input set partitioning), while requiring

<sup>1</sup>The proposed methods apply similarly to NNs with other activation functions, including sigmoid and tanh

**TABLE 1.** Reachable Sets for Double Integrator. The first two methods analyze the NN and dynamics separately, leading to high error accumulation, while the next 4 methods analyze the NN and dynamics together. Reach-LP is 4,000× faster to compute but 7× looser than Reach-SDP [21]. Reach-LP-Partition refines the Reach-LP bounds by splitting the input set into 16 subsets, giving 150× faster computation time and 5× tighter bounds than Reach-SDP [21].

Algorithm	Runtime [s]	Error
CL-CROWN	0.023 ± 0.000	654556
CL-SG-IBP [20]	0.888 ± 0.011	70218
Reach-SDP [21]	42.57 ± 0.54	207
Reach-SDP-Partition	670.26 ± 2.91	12
Reach-LP	0.017 ± 0.000	1590
Reach-LP-Partition	0.263 ± 0.001	34



**FIGURE 7.** Reachable sets for double integrator. In (a), all reachable set algorithms bound sampled states across the timesteps, starting from the blue  $\mathcal{X}_0$ , and the tightness of these bounds is quantified per timestep in (b).

150× less computation time. We implemented Reach-SDP in Python with `cvxpy` and `MOSEK` [46]. All computation times (until § VIII-I) are reported from an i7-6700K CPU with 16GB RAM on Ubuntu 20.04.

We also note the large error accumulation that occurs when the NN and dynamics are analyzed separately, as in [20], [32], as that assumes the NN could output its extreme values from any state in  $\mathcal{X}_t$  (top 2 rows of Table 1). CL-CROWN is a simple baseline where, at each timestep, the NN output bounds  $\mathcal{U}_t^{\text{NN}}$  are computed with CROWN [10], then (9) is solved (in closed-form) assuming  $u_t \in \mathcal{U}_t^{\text{NN}}$ . Similarly, CL-SG-IBP is the method from [20], where  $\mathcal{U}_t^{\text{NN}}$  is instead computed with SimGuided-IBP to resolution  $\epsilon = 0.1$ .

Fig. 7a shows sampled trajectories, where each colored cluster of points represents sampled reachable states at a particular timestep (blue→orange→green, etc.). Recall that sampling trajectories could miss possible reachable states, whereas these algorithms are guaranteed to over-approximate the reachable sets. Reachable set bounds are visualized for various algorithms: Reach-SDP [21], Reach-LP, and those two algorithms after partitioning the input set into 16 cells. The key takeaway is that while all approaches provide outer bounds on the sampled trajectories, the algorithms provide various degrees of tightness to the sampled points.

We quantify tightness as the ratio of areas between the smallest axis-aligned bounding box on the sampled points and the provided reachable set (minus 1), shown in Fig. 7b as the system progresses forward in time. Note that as expected, all algorithms get worse as the number of timesteps increase,

**TABLE 2.** Runtime in seconds shows over an order of magnitude speed-up by using closed-form (C.F.) instead of linear program (L.P.) to solve (20) when applicable (i.e.,  $\ell_p$ -ball sets).

Solver	# Partitions		
	1	4	16
L.P.	0.229 ± 0.018	0.856 ± 0.046	3.308 ± 0.091
C.F.	0.017 ± 0.000	0.066 ± 0.000	0.265 ± 0.002

(a) Double Integrator (5 timesteps)

Solver	# Partitions		
	1	4	16
L.P.	2.375 ± 0.039	9.424 ± 0.121	37.932 ± 0.508
C.F.	0.110 ± 0.002	0.432 ± 0.001	1.738 ± 0.003

(b) 6D Quadrotor (2 seconds)

but that Reach-LP-Partition and Reach-SDP-Partition perform the best and similarly. This provides numerical comparisons of the rectangle sizes from Fig. 7a.

Note that both Reach-LP and Reach-SDP methods could be improved by properly choosing the direction of polytope facets. Additionally, while Reach-SDP can provide ellipsoidal bounds given the quadratic nature of the formulation, we implement only the polytope bounds in this comparison.

### C. VERIFICATION

A primary application of reachable sets is to verify reach-avoid properties. In Fig. 7a, we consider a case with an avoid set  $\mathcal{A} = \{x \mid x_1 \geq 0.35\}$  (orange) and a goal set  $\mathcal{G} = [-0.5, 0.5] \times [-0.25, 0.25]$  (cyan). Reach-LP and Reach-SDP-Partition, verify these properties for this 5-step scenario, highlighting the importance of tight reachable sets.

### D. SCALABILITY TO DEEP NNS

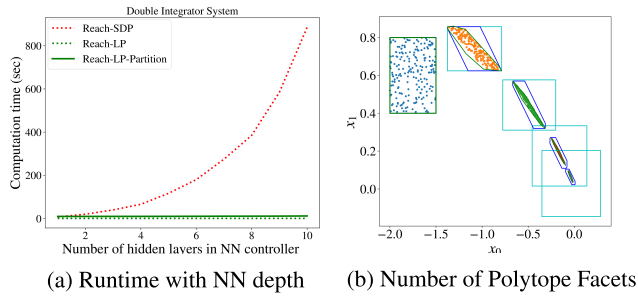
To demonstrate the scalability of the method, we trained NNs with 1-10 hidden layers of 5 neurons and report the average runtime of 5 trials of reachability analysis of the double integrator system. In Fig. 8a, while Reach-SDP appears to grow exponentially (taking > 800s for a 10-layer NN), our proposed Reach-LP methods remain very efficient (< 0.75s for Reach-LP on all NNs). Note that we omit Reach-SDP-Partition (~ 16× more than Reach-SDP) from this plot to maintain reasonable scale. Recall that CROWN itself has  $O(m^2n^3)$  time complexity [10], for an  $m$ -layer network with  $n$  neurons per layer and  $n$  outputs.

### E. ABLATION STUDY: $\ell_\infty$ VS. POLYTOPES

Recall that § IV described reachable sets as either polytopes or  $\ell_\infty$ -balls. Fig. 8b shows the effect of that choice: as the number of sides of the polytope increases, the reachable set size decreases. The tradeoff is that the computation time scales linearly with the number of sides on the polytope. Note that a  $\ell_\infty$ -ball is a 4-polytope, and that  $\mathcal{X}_0$  was chosen to show a different scenario than Fig. 7.

### F. ABLATION STUDY: CLOSED-FORM VS. LP

The computational benefit of the closed-form solution in Theorem 4.3 is shown in Table 2. For both the double



**FIGURE 8.** (a) Our linear relaxation-based methods (Reach-LP, Reach-LP-Partition) scale well for deeper NNs (Reach-LP: 0.6 to 0.74s), whereas SDP-based methods grow to intractable runtimes. Note that input set partitioning multiplies computation time by a scalar. (b) Using Reach-LP, the bounding shapes correspond to  $l_\infty$ -ball, 8-Polytope, and 35-Polytope. Reachable sets become tighter with more facets.

integrator (a) and 6D quadrotor (b), the runtime of the closed-form (C.F.) solution is over an order of magnitude faster than with the LP solver. This speedup is observed across various partition resolutions, with times reported as the mean  $\pm$  one standard deviation of 5 repeats. Note that the C.F. and LP solvers return the same reachable sets, so there is no tightness tradeoff – solely a computational speedup.

### G. GREEDY SIM-GUIDED PARTITIONING

Fig. 9 shows the refinement of reachable sets under the Closed-Loop Greedy Sim-Guided partitioner and CL-CROWN propagator. The blue reachable set estimates become much tighter to the Monte Carlo samples, while still providing a guaranteed over-approximation of the true system’s forward reachable sets. Note that the partition of  $\mathcal{X}_0$  is not uniform, as seen most clearly in Fig. 9e.

### H. 6D QUADROTOR WITH NOISE

Consider the 6D nonlinear quadrotor from [21], [47],

$$\dot{\mathbf{x}} = \underbrace{\begin{bmatrix} \mathbf{0}_{3 \times 3} & \mathbf{I}_3 \\ \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \end{bmatrix}}_{\mathbf{A}_t} \mathbf{x}_t + \underbrace{\begin{bmatrix} g & 0 & 0 \\ \mathbf{0}_{3 \times 3} & 0 & -g & 0 \\ 0 & 0 & 1 \end{bmatrix}^T}_{\mathbf{B}_t} \underbrace{\begin{bmatrix} \tan(\theta) \\ \tan(\phi) \\ \tau \end{bmatrix}}_{\mathbf{u}_t} + \underbrace{\begin{bmatrix} \mathbf{0}_{5 \times 1} \\ -g \end{bmatrix}}_{\mathbf{c}_t} + \boldsymbol{\omega}_t, \quad (71)$$

which differs from [21], [47] in that we add  $\boldsymbol{\omega}_t$  as a uniform process noise, and that the output is measured as in (1) with  $\mathbf{C}_t = \mathbf{I}_6$ , subject to uniform sensor noise. As in [21], the state vector contains 3D positions and velocities,  $[p_x, p_y, p_z, v_x, v_y, v_z]$ , while nonlinearities from [47] are absorbed into the control as functions of  $\theta$  (pitch),  $\phi$  (roll), and  $\tau$  (thrust) (subject to the same actuator constraints as [21]). We implemented a similar nonlinear MPC as [21] in MATLAB to collect  $(\mathbf{x}_t, \mathbf{u}_t)$  training pairs, then trained a [32,32] NN with Keras as above. We use Euler integration to account for (71) in our discrete time formulation.

Fig. 10 shows the reachable sets with and without noise. Note that while these plots only show  $(x, y, z)$  position, the reachable sets are estimated in all 6D. The first key takeaway is that the blue boxes (Reach-LP with  $l_\infty$ -balls) provide meaningful bounds for a long horizon (12 steps, 1.2s shown). Secondly, unlike Reach-SDP, Reach-LP is guaranteed to bound worst-case noise realizations. Computing these bounds took a total of  $0.11 \pm 0.003$  seconds.

### I. HIGH-DIMENSIONAL SYSTEM

To demonstrate that Reach-LP scales to high-dimensional systems, Fig. 11 uses the International Space Station (ISS) Component 1R model,<sup>2</sup> which has  $n_x = 270$  and  $n_u = 3$ . After training a NN in the same manner as § VIII-A, reachable sets for all 270 states were computed, with only  $(\mathbf{x}_0, \mathbf{x}_1)$  shown. The average time for each reachability analysis is about 25.34 sec on a desktop computer with i7-9700K CPU@3.60GHz (8 cores) and 32 GB RAM.

### J. POLYNOMIAL DYNAMICS

To show the SDP-based method from § VI-B on a system with polynomial dynamics, we consider the Duffing oscillator with the following dynamics

$$\dot{\mathbf{x}}_1 = \mathbf{x}_2, \quad (72)$$

$$\dot{\mathbf{x}}_2 = -\mathbf{x}_1 - 2\zeta \mathbf{x}_2 - \mathbf{x}_1^3 + u, \quad (73)$$

where  $\mathbf{x} \in \mathbb{R}^2$  and  $u \in \mathbb{R}$  are the state and control inputs, respectively, and  $\zeta = 0.3$  is the damping coefficient. Furthermore,  $\mathbf{c}_t = 0$ ,  $\mathbf{C}_t = \mathbf{I}_2$ , zero noise is considered, and the dynamics are discretized with sampling time  $t_s = 0.3s$ . We fit a neural network controller described in [48]. Applying the techniques introduced in § VI-B, the reachable sets are illustrated in Fig. 13. The estimated reachable sets contain the sampled states, but due to the relaxation, the bounds of the reachable sets are relatively loose compared to the linear systems that did not require relaxations for the dynamics.

### K. BACKWARD REACHABILITY

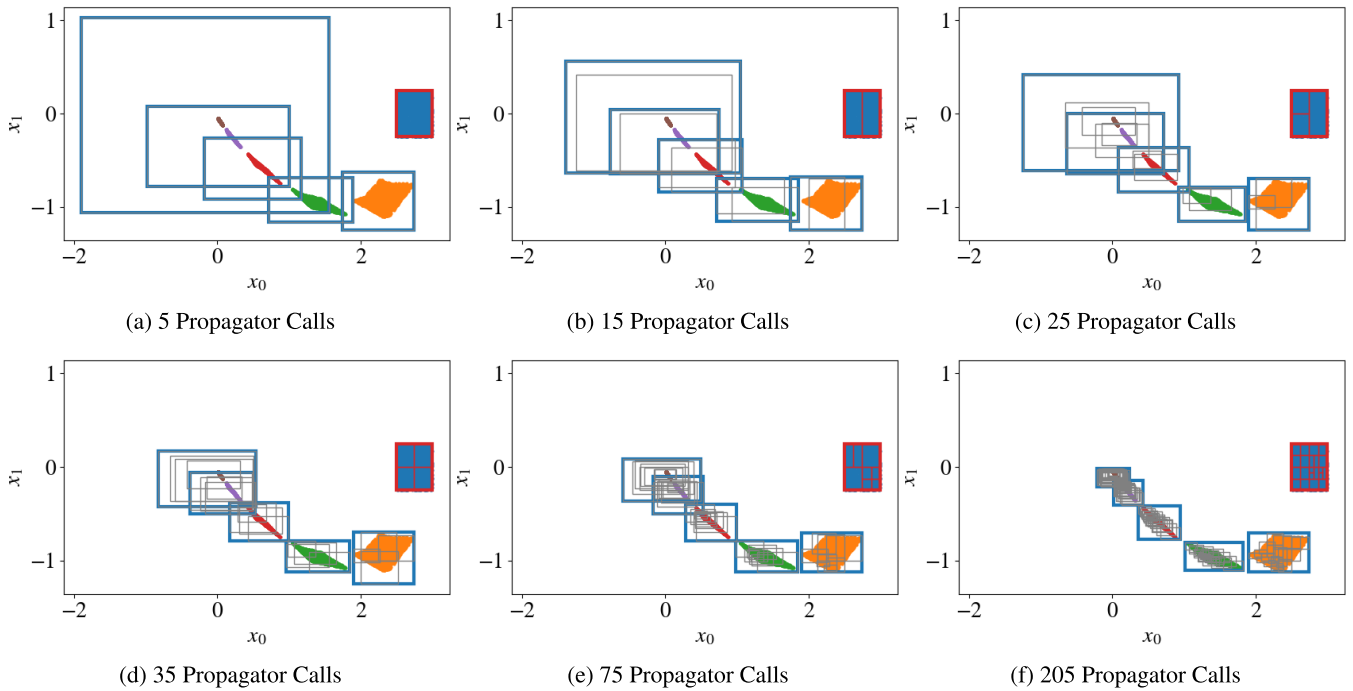
For the double integrator system and NN described in § VIII-A, Fig. 12 shows one-step backprojection sets for a given target set,  $\mathcal{X}_T$ . Recall that this NN is not invertible, both because it uses ReLU activations and some weight matrix dimensions prevent inversion.

The green rectangle represents  $\mathcal{X}_T = [2.5, 3.0] \times [-0.25, 0.25]$ .  $\mathcal{X}_T$  is partitioned into a uniform grid of various resolution in (a)-(d), ranging from  $2 \times 2$  to  $16 \times 16$ . The estimated backprojection set of each cell is computed and visualized as a red polytope. Thus, the estimated backprojection set of  $\mathcal{X}_T$  is the union of all red polytopes. In practice, if  $\mathcal{X}_T$  represents a goal set, one could check that the system’s starting state is inside one of the gray polytopes before “pressing go” on the system.

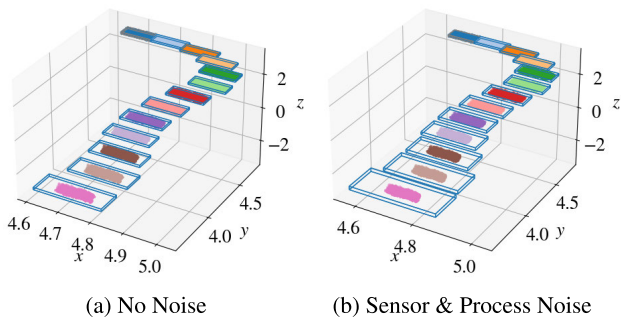
To demonstrate that points inside the backprojection set will move the system to the target set under the NN control

<sup>2</sup><http://slicot.org/20-site/126-benchmark-examples-for-model-reduction>






**FIGURE 9.** Reach-LP with Closed-Loop Greedy Simulation-Guided Partitioning. As  $\mathcal{X}_0$  (red box) is partitioned, the reachable set estimates (blue boxes) tighten to the Monte Carlo samples. Gray boxes show the reachable sets corresponding to cells of the  $\mathcal{X}_0$  partition.



**FIGURE 10.** Reachable Sets for 6D Quadrotor. Only  $(x, y, z)$  states are shown, even though the reachable sets are computed in 6D. Blue boxes (Reach-LP) bound the clusters of sampled points at each discrete timestep, starting from the blue  $\mathcal{X}_0$ . It took 0.11 sec to compute the 12 reachable sets per scenario. In (b),  $v \sim \text{Unif}(\pm 0.001 \cdot \mathbf{1}_6)$ ,  $\omega \sim \text{Unif}(\pm 0.005 \cdot \mathbf{1}_6)$ .

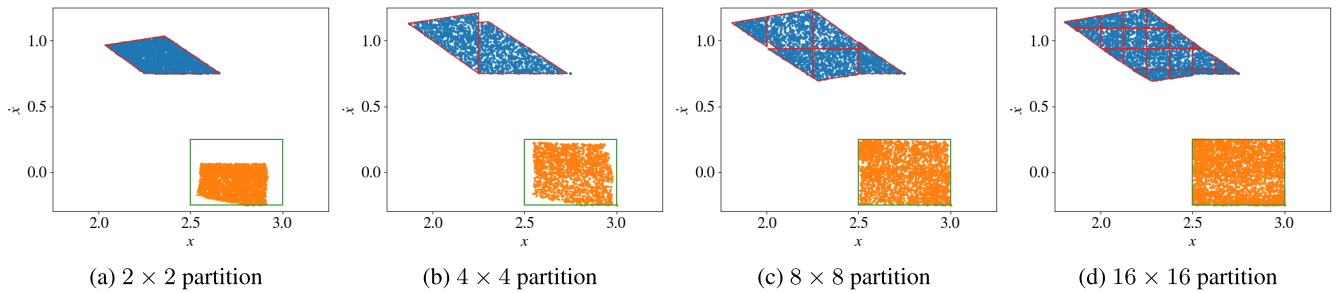
policy, we uniformly sampled points in each polytope of the estimated backprojection set and simulated the system forward one step (using the NN control policy). In Fig. 12, the blue points represent these samples at timestep  $T - 1$ , and the orange points are the resulting system state at timestep  $T$ . All of the orange points lie within the green rectangle of  $\mathcal{X}_T$ , which suggests the backprojection sets were computed properly.

Note that in (a), the under-approximation of the backprojection set has some clear gaps. There are a few reasons for this, but the key issue is that  $\mathcal{A} = \emptyset$  (from Algorithm 4) for some of the cells of the  $\mathcal{X}_T$  partition. What causes  $\mathcal{A} = \emptyset$ ? If  $\mathcal{X}_T$  is large, the backward reachable set could be as well.




**FIGURE 11.** Reachable Sets for the International Space Station (ISS) dynamics model with 270 states and a NN controller. Reach-LP scales to this high-dimensional system.

Thus, even a partitioned  $\mathcal{X}_T$  could lead to a large backward reachable set partition, which is used as the NN input set for CROWN. If the input set to CROWN is large, the NN relaxations can be relatively loose. Loose bounds on the NN restrict the size of  $\bar{\mathcal{P}}$ , since  $\bar{\mathcal{P}}$  only includes states for which that entire range of relaxed control inputs will lead to  $\mathcal{X}_T$  (see bottom arrow of Fig. 6). To get around this issue, (b)-(d) repeat the analysis with finer partitions, which substantially improves the coverage of samples across  $\mathcal{X}_T$ . Future work could investigate better methods for partitioning  $\mathcal{X}_T$  (e.g., using simulation as guidance [20], [36]).



**FIGURE 12.** 1-step Backprojection Sets for Double Integrator. The green rectangle represents  $\mathcal{X}_T$ , which is partitioned into an  $N \times N$  grid in (a)–(d) for various  $N$ . The estimated backprojection set of each cell of  $\mathcal{X}_T$  is computed and visualized as a single red polytope (when non-empty). Thus, the estimated backprojection set of  $\mathcal{X}_T$  is the union of all red polytopes. To confirm the results, blue points are sampled from the estimated backprojection sets, and the system state at the next timestep (under the NN control policy) are shown in orange. The orange points do indeed lie within  $\mathcal{X}_T$ , with better coverage of  $\mathcal{X}_T$  as the number of partitions increases.



**FIGURE 13.** Reachable Sets for the Duffing oscillator. Starting from  $\mathcal{X}_0$  (gray rectangle, blue samples), the Reach-LP extension from § VI-B is used to estimate reachable sets for 6 timesteps into the future, despite the nonlinear (polynomial) terms in the dynamics.

## IX. FUTURE DIRECTIONS

Many open research directions remain in analyzing NFLs. We first describe several natural extensions of this work. One existing challenge is in mitigating the conservatism due to the accumulation of approximation error over many timesteps. This could be addressed, for example, by replacing the recursive calculations with a method for computing  $n$ -step reachable sets at once. In addition, capturing other types of uncertainties and nonlinearities (e.g., uncertainty in  $\mathbf{A}_t$  and  $\mathbf{B}_t$  matrices) will expand the set of neural feedback loops that can be analyzed. Continuous time systems, which likely have hybrid dynamics due to the NN control policy, present additional opportunities for expanding the framework beyond the Euler integration approach described here.

More broadly, an open challenge is in synthesizing provably robust control policies. While this work focused on analyzing NFLs with a pre-trained NN, future work could explore modifications to the training process to enable faster or tighter online analysis. Moreover, bringing similar ideas to stability analysis could provide additional notions of robustness beyond reachability.

Finally, the adoption of these analysis methods on real safety-critical systems will require realistic measurement/

perception models, as many modern systems use high-dimensional sensors (e.g., camera, lidar), which are often fed directly into perception NNs (e.g., for object detection or segmentation). The effects of uncertainty propagating through these modular pipelines presents new challenges before such systems are ready to be deployed in safety-critical settings, such as robots operating alongside humans.


## X. CONCLUSION

This paper proposed a convex relaxation-based algorithm for computing forward reachable sets and backprojection sets of NFLs, which are closed-loop systems with NN controllers. Prior work is limited to shallow NNs and is computationally intensive, which limits applicability to real systems. Furthermore, our method accounts for measurement of sensor and process noise, as well as nonlinearities in the dynamics. The results show that this work advances the state-of-the-art in guaranteeing properties of systems that employ NNs in the feedback loop.

## REFERENCES

- [1] R. Ehlers, “Formal verification of piece-wise linear feed-forward neural networks,” in *Proc. ATVA*, 2017, pp. 269–286.
- [2] G. Katz, C. W. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer, “Reluplex: An efficient SMT solver for verifying deep neural networks,” in *Proc. 29th Int. Conf. CAV*, Heidelberg, Germany, Jul. 2017, pp. 97–117.
- [3] X. Huang, M. Kwiatkowska, S. Wang, and M. Wu, “Safety verification of deep neural networks,” in *Computer Aided Verification*, R. Majumdar and V. Kunčák, Eds. Cham, Switzerland: Springer, 2017, pp. 3–29.
- [4] A. Lomuscio and L. Maganti, “An approach to reachability analysis for feed-forward Relu neural networks,” *CoRR*, vol. abs/1706.07351, pp. 1–10, Jun. 2017.
- [5] V. Tjeng, K. Y. Xiao, and R. Tedrake, “Evaluating robustness of neural networks with mixed integer programming,” in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2019, pp. 1–21.
- [6] T. Gehr, M. Mirman, D. Drachler-Cohen, P. Tsankov, S. Chaudhuri, and M. Vechev, “AI2: Safety and robustness certification of neural networks with abstract interpretation,” in *Proc. IEEE Symp. Secur. Privacy (SP)*, May 2018, pp. 3–18.
- [7] S. Gowal, K. Dvijotham, R. Stanforth, R. Bunel, C. Qin, J. Uesato, R. Arandjelovic, T. Mann, and P. Kohli, “On the effectiveness of interval bound propagation for training verifiably robust models,” 2018, *arXiv:1810.12715*.
- [8] T. Weng, H. Zhang, H. Chen, Z. Song, C. Hsieh, L. Daniel, D. Boning, and I. Dhillon, “Towards fast computation of certified robustness for Relu networks,” in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2018, pp. 5276–5285.
- [9] G. Singh, T. Gehr, M. Mirman, M. Püschel, and M. Vechev, “Fast and effective robustness certification,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 10802–10813.

- [10] H. Zhang, T.-W. Weng, P.-Y. Chen, C.-J. Hsieh, and L. Daniel, "Efficient neural network robustness certification with general activation functions," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 4939–4948.
- [11] E. Wong and J. Z. Kolter, "Provable defenses against adversarial examples via the convex outer adversarial polytope," in *Proc. ICML*, vol. 80, 2018, pp. 5283–5292.
- [12] A. Raghunathan, J. Steinhardt, and P. Liang, "Certified defenses against adversarial examples," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2018, pp. 1–15.
- [13] M. Fazlyab, M. Morari, and G. J. Pappas, "Safety verification and robustness analysis of neural networks via quadratic constraints and semidefinite programming," 2019, *arXiv:1903.01287*.
- [14] C. J. Tomlin, J. Lygeros, and S. S. Sastry, "A game theoretic approach to controller design for hybrid systems," *Proc. IEEE*, vol. 88, no. 7, pp. 949–970, Jul. 2000.
- [15] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin, "Hamilton–Jacobi reachability: A brief overview and recent advances," in *Proc. IEEE 56th Annu. Conf. Decis. Control (CDC)*, Dec. 2017, pp. 2242–2253.
- [16] S. Dutta, X. Chen, and S. Sankaranarayanan, "Reachability analysis for neural feedback systems using regressive polynomial rule inference," in *Proc. 22nd ACM Int. Conf. Hybrid Syst., Comput. Control*, Apr. 2019, pp. 157–168.
- [17] C. Huang, J. Fan, W. Li, X. Chen, and Q. Zhu, "ReachNN: Reachability analysis of neural-network controlled systems," *ACM Trans. Embedded Comput. Syst.*, vol. 18, no. 5, pp. 1–22, Oct. 2019.
- [18] J. Fan, C. Huang, X. Chen, W. Li, and Q. Zhu, "ReachNN: A tool for reachability analysis of neural-network controlled systems," in *Proc. Int. Symp. Autom. Technol. Verification Anal.* New York, NY, USA: ACM, 2020, pp. 537–542.
- [19] R. Ivanov, J. Weimer, R. Alur, G. J. Pappas, and I. Lee, "Verisig: Verifying safety properties of hybrid systems with neural network controllers," in *Proc. 22nd ACM Int. Conf. Hybrid Syst., Comput. Control*, 2019, pp. 169–178.
- [20] W. Xiang, H.-D. Tran, X. Yang, and T. T. Johnson, "Reachable set estimation for neural network control systems: A simulation-guided approach," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 5, pp. 1821–1830, May 2021.
- [21] 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 *Proc. 59th IEEE Conf. Decis. Control (CDC)*, Dec. 2020, pp. 5929–5934.
- [22] L. C. Evans, *Partial Differential Equations* (Graduate Studies in Mathematics), vol. 19, no. 2. Boston, MA, USA: International Press of Boston, 1998.
- [23] M. Everett, G. Habibi, and J. P. How, "Efficient reachability analysis of closed-loop systems with neural network controllers," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2021, pp. 4384–4390.
- [24] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2014, pp. 1–10.
- [25] G. Frehse, C. L. Guernic, A. Donzé, S. Cotton, R. Ray, O. Lebeltel, R. Ripado, A. Girard, T. Dang, and O. Maler, "SpaceEx: Scalable verification of hybrid systems," in *Proc. Int. Conf. Comput. Aided Verification*. Berlin, Germany: Springer, 2011, pp. 379–395.
- [26] X. Chen, E. Ábrahám, and S. Sankaranarayanan, "Flow: An analyzer for non-linear hybrid systems," in *Proc. Int. Conf. Comput. Aided Verification*. Berlin, Germany: Springer, 2013, pp. 258–263.
- [27] M. Althoff, "An introduction to CORA 2015," in *Proc. Workshop Appl. Verification Continuous Hybrid Syst.*, 2015, pp. 1–28.
- [28] P. S. Duggirala, S. Mitra, M. Viswanathan, and M. Potok, "C2E2: A verification tool for stateflow models," in *Proc. Int. Conf. Tools Algorithms Construct. Anal. Syst.* Berlin, Germany: Springer, 2015, pp. 68–82.
- [29] C. Fan, B. Qi, S. Mitra, M. Viswanathan, and P. S. Duggirala, "Automatic reachability analysis for nonlinear hybrid models with C2E2," in *Proc. Int. Conf. Comput. Aided Verification*. Cham, Switzerland: Springer, 2016, pp. 531–538.
- [30] A. Papachristodoulou and S. Prajna, "On the construction of Lyapunov functions using the sum of squares decomposition," in *Proc. 41st IEEE Conf. Decis. Control*, vol. 3, Dec. 2002, pp. 3482–3487.
- [31] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. Autom. Control*, vol. 62, no. 8, pp. 3861–3876, Aug. 2017.
- [32] G. Yang, G. Qian, P. Lv, and H. Li, "Efficient verification of control systems with neural network controllers," in *Proc. 3rd Int. Conf. Vis., Image Signal Process.*, Aug. 2019, pp. 1–7.
- [33] S. Diamond and S. Boyd, "CVXPY: A Python-embedded modeling language for convex optimization," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 2909–2913, Jan. 2016.
- [34] M. Everett, B. Lütjens, and J. P. How, "Certifiable robustness to adversarial state uncertainty in deep reinforcement learning," *IEEE Trans. Neural Netw. Learn. Syst.*, early access, Feb. 15, 2021, doi: 10.1109/TNNLS.2021.3056046.
- [35] W. Xiang, H.-D. Tran, and T. T. Johnson, "Output reachable set estimation and verification for multilayer neural networks," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 11, pp. 5777–5783, Mar. 2018.
- [36] M. Everett, G. Habibi, and J. P. How, "Robustness analysis of neural networks via efficient partitioning with applications in control systems," *IEEE Control Syst. Lett.*, vol. 5, no. 6, pp. 2114–2119, Dec. 2021.
- [37] P. A. Parrilo and B. Sturmfels, "Minimizing polynomial functions," in *Algorithmic and Quantitative Real Algebraic Geometry (DIMACS Series in Discrete Mathematics and Theoretical Computer Science)*, vol. 60. 2003, pp. 83–99.
- [38] Z.-Q. Luo, W.-K. Ma, A. M.-C. So, Y. Ye, and S. Zhang, "Semidefinite relaxation of quadratic optimization problems," *IEEE Signal Process. Mag.*, vol. 27, no. 3, pp. 20–34, May 2010.
- [39] L. Ardizzone, J. Kruse, C. Rother, and U. Köthe, "Analyzing inverse problems with invertible neural networks," in *Proc. Int. Conf. Learn. Represent.*, 2018, pp. 1–20.
- [40] J. Behrmann, W. Grathwohl, R. T. Chen, D. Duvenaud, and J.-H. Jacobsen, "Invertible residual networks," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 573–582.
- [41] 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," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 1–13.
- [42] K. Xu, H. Zhang, S. Wang, Y. Wang, S. Jana, X. Lin, and C.-J. Hsieh, "Fast and complete: Enabling complete neural network verification with rapid and massively parallel incomplete verifiers," in *Proc. Int. Conf. Learn. Represent.*, 2021, pp. 1–15.
- [43] S. Wang, H. Zhang, K. Xu, X. Lin, S. Jana, C.-J. Hsieh, and J. Z. Kolter, "Beta-CROWN: Efficient bound propagation with per-neuron split constraints for complete and incomplete neural network robustness verification," 2021, *arXiv:2103.06624*.
- [44] S. M. LaValle, *Planning Algorithms*. Cambridge, U.K.: Cambridge Univ. Press, 2006.
- [45] F. Chollet. (2015). *Keras*. [Online]. Available: <https://keras.io>
- [46] E. D. Andersen and K. D. Andersen, "The MOSEK interior point optimizer for linear programming: An implementation of the homogeneous algorithm," in *High Performance Optimization*. Springer, 2000, pp. 197–232.
- [47] D. M. Lopez, P. Musau, H.-D. Tran, and T. T. Johnson, "Verification of closed-loop systems with neural network controllers," *EPiC Ser. Comput.*, vol. 61, pp. 201–210, May 2019.
- [48] T. Dang and R. Testylier, "Reachability analysis for polynomial dynamical systems using the Bernstein expansion," *Rel. Comput.*, vol. 17, no. 2, pp. 128–152, 2012.



**MICHAEL EVERETT** (Member, IEEE) received the S.B., S.M., and Ph.D. degrees in mechanical engineering from the Massachusetts Institute of Technology (MIT), in 2015, 2017, and 2020, respectively. He is currently a Research Scientist with the Department of Aeronautics and Astronautics, MIT. His research interests include machine learning, robotics, and control theory, with specific interests in the theory and application of safe and robust neural feedback loops. He was an author of

works that won the Best Paper Award on Cognitive Robotics at IROS 2019, the Best Student Paper Award, and a Finalist for the Best Paper Award on Cognitive Robotics at IROS 2017, and a Finalist for the Best Multi-Robot Systems Paper Award at ICRA 2017. He has been interviewed live on the air by BBC Radio and his team's robots were featured by Today Show and the Boston Globe.



**GOLNAZ HABIBI** (Member, IEEE) received the B.Sc. degree in electrical and control engineering from the K. N. Toosi University of Technology, Iran, in 2005, the M.Sc. degree in control engineering from Tarbiat Modares University, Iran, in 2007, and the Ph.D. degree in computer science from Rice University, in 2015. She is currently a Research Scientist with the Department of Aeronautics and Astronautics, MIT. Her research interests include robotics, control systems, machine learning, multi agent systems, visual navigation, reliable communication, and improving the safety and reliability of autonomous agents. Her paper has been nominated for best student paper award in DARS 2012 and she received the K2I award by Chevron Corporation, in 2013.



**CHUANGCHUANG SUN** (Member, IEEE) received the B.S. degree in aerospace engineering from the Beijing University of Aeronautics and Astronautics, China, in 2013, and the Ph.D. degree in aerospace engineering from Ohio State University, in August 2018. He is currently a Postdoctoral Associate with the Department of Aeronautics and Astronautics, MIT. His research interests include control, optimization, reinforcement learning with applications in robotics, and space systems.



**JONATHAN P. HOW** (Fellow, IEEE) received the B.A.Sc. degree in aerospace from the University of Toronto in 1987, and the S.M. and Ph.D. degrees in aeronautics and astronautics from MIT, in 1990 and 1993, respectively, and then studied for 1.5 years at MIT as a Postdoctoral Associate. Prior to joining MIT in 2000, he was an Assistant Professor with the Department of Aeronautics and Astronautics, Stanford University. He was the Editor-in-Chief of the *IEEE Control Systems Magazine* (2015–2019). He is currently an Associate Editor of the *AIAA Journal of Aerospace Information Systems* and the IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS. He is also the Richard C. Maclaurin Professor in the Department of Aeronautics and Astronautics with the Massachusetts Institute of Technology. He was an Area Chair of the International Joint Conference on Artificial Intelligence, in 2019, and will be the Program Vice-Chair (tutorials) for the Conference on Decision and Control, in 2021. He was elected to the Board of Governors of the IEEE Control System Society (CSS), in 2019, and is a member of the IEEE CSS Technical Committee on Aerospace Control and the Technical Committee on Intelligent Control. He is also the Director of the Ford-MIT Alliance and was a member of the USAF Scientific Advisory Board (SAB), from 2014 to 2017. His research focuses on robust planning and learning under uncertainty with an emphasis on multi-agent systems, and he was the planning and control lead for the MIT DARPA Urban Challenge team. His work has been recognized with multiple awards, including the 2020 AIAA Intelligent Systems Award. He is a fellow of AIAA.

• • •