

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/344447310>

Robustness Analysis of Neural Networks via Efficient Partitioning: Theory and Applications in Control Systems

Preprint · October 2020

CITATIONS

0

READS

139

3 authors:



Michael Everett

Northeastern University

52 PUBLICATIONS 2,041 CITATIONS

[SEE PROFILE](#)



Golnaz Habibi

Rice University

26 PUBLICATIONS 154 CITATIONS

[SEE PROFILE](#)



Jonathan How

Massachusetts Institute of Technology

930 PUBLICATIONS 29,554 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



PI Controller for Non-affine Systems [View project](#)



Path Planning and Controller for DARPA Urban Challenge [View project](#)

Robustness Analysis of Neural Networks via Efficient Partitioning: Theory and Applications in Control Systems

Michael Everett, Golnaz Habibi, Jonathan P. How

Abstract—Neural networks (NNs) are now routinely implemented on systems that must operate in uncertain environments, but the tools for formally analyzing how this uncertainty propagates to NN outputs are not yet commonplace. Computing tight bounds on NN output sets (given an input set) provides a measure of confidence associated with the NN decisions and is essential to deploy NNs on safety-critical systems. Recent works approximate the propagation of sets through nonlinear activations or partition the uncertainty set to provide a guaranteed outer bound on the set of possible NN outputs. However, the bound looseness causes excessive conservatism and/or the computation is too slow for online analysis. This paper unifies propagation and partition approaches to provide a family of robustness analysis algorithms that give tighter bounds than existing works for the same amount of computation time (or reduced computational effort for a desired accuracy level). Moreover, we provide new partitioning techniques that are aware of their current bound estimates and desired boundary shape (e.g., lower bounds, weighted ℓ_∞ -ball, convex hull), leading to further improvements in the computation-tightness tradeoff. The paper demonstrates the tighter bounds and reduced conservatism of the proposed robustness analysis framework with examples from model-free RL and forward kinematics learning.

I. INTRODUCTION

Neural networks (NNs) are ubiquitous across robotics for perception, planning, and control tasks. While empirical performance statistics can indicate that a NN has learned a useful input-output mapping, there are still concerns about how much confidence to associate with decisions resulting from a learned system. One direction toward providing a confidence measure is to consider how the various sources of uncertainty in training/execution processes map to uncertainty in outputs of trained NNs. Many of these uncertainties appear at the NN input (e.g., from noisy/adversarially attacked sensing, unknown initial conditions), thus this work focuses on the problem of propagating input uncertainties through NNs to bound the set of possible NN outputs online.

Analysis of how a set of possible inputs propagates through a NN has an inherent tradeoff between computation time and conservatism. Exact methods [1]–[6] are computationally intractable for online analysis, so we focus on finding guaranteed outer bounds on the network outputs. Most existing methods propagate the entire input set through the NN – we refer to these as *Propagators* [7]–[12].

Although some of these propagators scale to high dimensional NNs, large input sets (e.g., from high state uncertainty) induce massive conservatism, even for small

The authors are with the Aerospace Controls Laboratory at the Massachusetts Institute of Technology. {mfe, ghabibi, jhow}@mit.edu.




Fig. 1: Robustness Analysis Architecture. This work efficiently bounds the set of NN outputs for a given input set.

NNs. *Partitioners* [13]–[16] are a promising direction toward propagating large input sets through NNs, particularly when the number of uncertain NN inputs is relatively small, as in many control systems. Nonetheless, current partitioners spend excessive computational effort when refining cells, have limited theoretical justification, and suffer from simple propagation strategies.

The contributions of this work are: (i) a family of analysis tools that provide tighter guaranteed bounds on possible NN outputs for the same computational effort compared to the state-of-the-art, (ii) two new partitioning algorithms that efficiently refine the input set partition based on desired output set shapes, (iii) theoretical justification for partitioning via analytically quantifying the reduction in output set size, and (iv) applications on control systems to improve a robust RL agent’s navigation efficiency by 22% and reduce reachable set estimation error by 96% for a robot arm.

II. BACKGROUND AND RELATED WORK

Problem Statement: Given a trained NN and a set of possible NN inputs, the objective is to find the tightest guaranteed over-approximation of the set of possible NN outputs. The exact set of possible outputs is called the NN’s *reachable set*. In general, finding the exact reachable set is computationally intractable for reasonably sized NNs; instead, the goal is to compute the over-estimate \mathcal{U}_e such that $\mathcal{U} \subseteq \mathcal{U}_e$ for a given input set.

A. Propagators

Propagators estimate how the full input set moves through the network, and they primarily differ in approximation strategies of the nonlinear activation functions. At one extreme, Interval Bound Propagation (IBP) [7] approximates the output of each layer with a tight ℓ_∞ ball, leading to conservative but fast-to-compute bounds of the final layer. Convex relaxation-based techniques [17] often achieve tighter bounds with more computation by approximating nonlinear activations with linear bounds – some of these




Fig. 2: Partitioner Intuition. (Top) Large input sets cause loose bounds on NN output sets, even for this simple 2-layer NN with linear activations. (Bottom) Tighter bounds can be achieved by partitioning the input set, propagating each cell through the NN, and merging the output sets [13], [14].

can be solved in closed-form [10], [18]. Other propagators provide tighter analysis at the cost of higher computation time, including approaches based on QP/SDP [8], [9], and convex relaxation refinements [19]. While this paper focuses on analysis of trained NNs, several recent works consider the orthogonal problem of how to use these propagation techniques during training [20].

B. Partitioners

Partitioners break the input set into smaller regions, compute the reachable set of each small region, and return the total reachable set as the union of each smaller region’s reachable set. The idea is depicted in Fig. 2 for a simple NN with linear activations. In the top row (without partitioning) IBP operates on the full input set, leading to excessive conservatism in the final output bound (top right: large red dashes vs. red shaded region). The bottom row shows how IBP on two halves of the input set leads to a tighter approximation of the output set. This volume reduction is quantified analytically in Section III-B.

The key difference between partitioning approaches is the strategy for how to split the input set. Some works make one bisection of the input set [15], [16]; [13] splits the input set into a uniform grid. The current state-of-art partitioner, a Simulation-Guided approach (SG) [14], uses a partitioning strategy where Monte Carlo samples of the exact NN output are used as guidance for efficient partitioning of the input set, reducing the amount of computation required for the same level of bound tightness. SG used IBP to compute output sets, and the two ideas of Partitioners and Propagators have been developed separately toward a similar objective. This work addresses key gaps in the partitioning literature: we unify Partitioners with state-of-art Propagators for better performance, propose new partitioners that are flexible in the desired output set shape, and analyze the theoretical reduction in conservatism. We then show how improvements in robustness analysis map directly to reduced conservatism in control tasks.

III. APPROACH

This section introduces the overall architecture, describes our new partitioning algorithms, then analyzes the reduction




Fig. 3: (Left) GSG selects from candidates in M : the input set corresponding to the magenta output set is chosen for refinement, because $d_2 > d_1$ (where $[u_{sim}]$ is the black box). (Right) AGSG decomposition $[\eta] \setminus [\eta_e]$ into four new intervals $[\eta']_{1,2,3,4}$.

in conservatism from partitioning.

Figure 1 shows a schematic of the proposed framework with its three nested modules: Analyzer, Partitioner and Propagator. The Analyzer is aware of the desired output shape (e.g., lower bounds, ℓ_∞ -ball, convex hull) and termination condition (e.g., computation time, number of Propagator calls, improvement per step). The Analyzer specifies a Propagator (e.g., CROWN [10], IBP [7], SDP [9], Fast-Lin [18]) and a Partitioner (e.g., Uniform [13], Sim-Guided [14] or the algorithms proposed in this section). The Partitioner decides how to split the input set into cells, and the Propagator is used by the Partitioner to estimate the output set corresponding to an input set or its cells.

A. Greedy Simulation-Guided Partitioning

The state-of-art partitioning algorithm, SG [14], tightens IBP’s approximated boundary with the following key steps: (1) acquire N Monte Carlo samples of the NN outputs to under-approximate the reachable set as the interval $[u_{sim}]$, (2) using IBP, compute the reachable set of the full input set and add this set to a stack M , and (3) (iteratively) pop an element from M , and either stop refining that cell if its computed reachable set is within $[u_{sim}]$, or bisect the cell, compute each bisection’s reachable set, and add both to the queue. The SG algorithm terminates when one of the cell’s dimensions reaches some threshold, and the returned reachable set estimate is the weighted ℓ_∞ -ball that surrounds the union of all of the cells remaining on the queue and $[u_{sim}]$.

We propose a partitioning algorithm with better bound tightness for the same amount of computation, called Greedy-Sim-Guided (GSG), by modifying the choice of which cell in M to refine at each step. Rather than popping the first element from the stack (LIFO) as in SG, GSG refines the input cell with corresponding output range that is furthest outside the output boundary of the N samples (Fig. 3). This heuristic gives the opportunity to reduce the boundary estimate at each iteration. While the core SG algorithm remains the same, the greedy strategy can greatly improve the algorithm’s performance.

The Adaptive-Greedy-Sim-Guided (AGSG) algorithm extends GSG’s initialization procedure to reduce wasted computation time refining unimportant input regions. After computing $[u_{sim}]$, AGSG initializes $[\eta_e]$ as the input point whose output is at the middle of simulated outputs. This cell is repeatedly expanded with step size e_s as long as its output

set (computed by a Propagator) remains inside $[u_{\text{sim}}]$. The remaining input is decomposed into a set of disjoint intervals $[\eta']_i$. Fig. 3 shows this decomposition in 2D, which creates four new intervals $[\eta']_{1,2,3,4}$ (some of which could be empty) of rectangle shape¹. The new intervals $[\eta']_i$ are passed to GSG as the initial M . The output set boundary estimate returned by AGSG merges that GSG output and the initial expanded cell's output boundary.

Boundary Specification: While SG only computes a ℓ_∞ -ball over-approximation, GSG/AGSG optimize for the desired output set shape. For example, if the objective is to find a tight convex hull over-approximation, GSG/AGSG modify the idea from Fig. 3 to select the input set that is furthest from the convex hull boundary (instead of the ℓ_∞ -ball, as in SG [14]).

B. Analysis of Partitioning

This section quantifies the reduction in conservatism from partitioning, as measured by output set volume. While [14] proves that smaller input intervals lead to a smaller output intervals, that analysis did not consider unions of output sets.

Overview: For a given input set, we first compute the volume of the output set when the full input set is propagated through a layer (using IBP [7]), then we compute the volume of the union of two output sets corresponding to two pieces of the input set. We then quantify the volume of the full output set and the union of the two pieces. The estimate of the output set's volume (as the union of output sets) is shown to improve from repeated splitting of the input set.

Consider a single NN layer with weight $\mathbf{W} \in \mathbb{R}^{n_{\text{out}} \times n_{\text{in}}}$.

Definition III.1. Let the input set be described by the n_{in} -paralleloptope, $\mathcal{P}(\mathbf{U}) \in \mathbb{R}^{n_{\text{in}}}$, with generating vectors,

$$\mathbf{U} = [\mathbf{u}_1 \mathbf{u}_2 \dots \mathbf{u}_{n_{\text{in}}}] \in \mathbb{R}^{n_{\text{in}} \times n_{\text{in}}}. \quad (1)$$

Definition III.2. The pre-activation set of the layer is a n_{in} -paralleloptope, $\mathcal{P}(\mathbf{V}) \in \mathbb{R}^{n_{\text{out}}}$, with

$$\mathbf{V} = \mathbf{W}\mathbf{U} = [\mathbf{v}_1 \mathbf{v}_2 \dots \mathbf{v}_{n_{\text{in}}}] \in \mathbb{R}^{n_{\text{out}} \times n_{\text{in}}}. \quad (2)$$

Definition III.3. Let $\mathcal{B}(\mathbf{V})$ denote the smallest axis-aligned hyperrectangle² containing $\mathcal{P}(\mathbf{V})$, i.e., the Cartesian product of n_{out} intervals, each of which is defined by the minimal and maximal value of the corresponding coordinate for the vertices of $\mathcal{P}(\mathbf{V})$.

Thus, $\mathcal{B}(\mathbf{V})$ has n_{out} -volume,

$$\text{Vol}(\mathcal{B}(\mathbf{V})) = \prod_{i=1}^{n_{\text{out}}} \sum_{j=1}^{n_{\text{in}}} |\mathbf{v}_{j,i}| = \prod_{i=1}^{n_{\text{out}}} l_i, \quad (3)$$

where $\mathbf{v}_{j,i}$ is the value of vector \mathbf{v}_j in i -th dimension.

Consider splitting the input set into two parts, where we choose to split \mathbf{u}_1 , w.l.o.g., by a scalar ratio $r \in [0, 1]$.

Definition III.4. The “left” and “right” input paralleloptopes, $\mathcal{P}(\mathbf{U}^L), \mathcal{P}(\mathbf{U}^R)$, are described by the generating vectors (\mathbf{U}^R 's corner is offset from \mathbf{U}^L 's corner by $r\mathbf{u}_1$),

$$\mathbf{U}^L = [r\mathbf{u}_1 \mathbf{u}_2 \dots \mathbf{u}_{n_{\text{in}}}] ; \mathbf{U}^R = [(1-r)\mathbf{u}_1 \mathbf{u}_2 \dots \mathbf{u}_{n_{\text{in}}}] . \quad (4)$$

Lemma III.5. Given a single layer of a NN, with weights \mathbf{W} and input set $\mathcal{P}(\mathbf{U})$ as in Definition III.1, partitioning $\mathcal{P}(\mathbf{U})$ at \mathbf{u}_1 by ratio r as in Definition III.4 reduces the volume of the pre-activation set by (11).

Proof. Using Definitions III.2 and III.4, the “left” and “right” pre-activation sets, $\mathcal{P}(\mathbf{V}^L), \mathcal{P}(\mathbf{V}^R)$ have generating vectors,

$$\mathbf{V}^L = \mathbf{W}\mathbf{U}^L ; \quad \mathbf{V}^R = \mathbf{W}\mathbf{U}^R \quad (5)$$

Using Definition III.3, $\mathcal{P}(\mathbf{V}^L), \mathcal{P}(\mathbf{V}^R)$ can each be bounded by a hyperrectangle, denoted as $\mathcal{B}(\mathbf{V}^L), \mathcal{B}(\mathbf{V}^R)$, respectively, with volumes,

$$\text{Vol}(\mathcal{B}(\mathbf{V}^L)) = \prod_{i=1}^{n_{\text{out}}} (r|\mathbf{v}_{1,i}| + z_i) \quad (6)$$

$$\text{Vol}(\mathcal{B}(\mathbf{V}^R)) = \prod_{i=1}^{n_{\text{out}}} ((1-r)|\mathbf{v}_{1,i}| + z_i) \quad (7)$$

using the notation $z_i = \sum_{j=2}^{n_{\text{out}}} |\mathbf{v}_{j,i}|$.

Because \mathbf{U}^L and \mathbf{U}^R have $n_{\text{in}} - 1$ vectors in common, $\mathcal{B}(\mathbf{V}^L)$ and $\mathcal{B}(\mathbf{V}^R)$ both contain the hyperrectangle with diagonal $\mathbf{d} = \mathbf{v}_2 + \dots + \mathbf{v}_{n_{\text{in}}}$. Their intersection's volume is

$$\text{Vol}(\mathcal{B}(\mathbf{V}^L) \cap \mathcal{B}(\mathbf{V}^R)) = \prod_{i=1}^{n_{\text{out}}} z_i = \prod_{i=1}^{n_{\text{out}}} \sum_{j=2}^{n_{\text{out}}} |\mathbf{v}_{j,i}|, \quad (8)$$

and their union's volume is

$$\begin{aligned} \text{Vol}(\mathcal{B}(\mathbf{V}^L) \cup \mathcal{B}(\mathbf{V}^R)) &= \prod_{i=1}^{n_{\text{out}}} (r|\mathbf{v}_{1,i}| + z_i) \\ &+ \prod_{i=1}^{n_{\text{out}}} ((1-r)|\mathbf{v}_{1,i}| + z_i) - \prod_{i=1}^{n_{\text{out}}} z_i. \end{aligned} \quad (9)$$

The difference in volume between $\mathcal{B}(\mathbf{V})$ and the union of the two pre-activation sets, $\mathcal{B}(\mathbf{V}^L) \cup \mathcal{B}(\mathbf{V}^R)$, is,

$$\begin{aligned} V_{\text{red}} &= \text{Vol}(\mathcal{B}(\mathbf{V})) - \text{Vol}(\mathcal{B}(\mathbf{V}^L) \cup \mathcal{B}(\mathbf{V}^R)) \\ &= \sum_{k=2}^{n_{\text{out}}} (1 - r^k - (1-r)^k) \\ &\cdot \left[\sum_{i=1}^{(n_{\text{out}})} \left[\prod_{j \in C_i} |\mathbf{v}_{1,j}| \right] \prod_{t \in C \setminus C_i} z_t \right], \end{aligned} \quad (10)$$

where $C = \{1, \dots, n_{\text{out}}\}$, $P_k(C)$ is the set containing elements of the power set of C with exactly k elements, and each element of $P_k(C)$ is denoted C_i . \square

The volume of pre-activation space in consideration reduces with each input partition when $V_{\text{red}} > 0$. Note that $V_{\text{red}} \geq 0$ (partitioning cannot hurt), since no vertex of $\mathcal{P}(\mathbf{V}^L)$ or $\mathcal{P}(\mathbf{V}^R)$ can exceed a vertex of $\mathcal{P}(\mathbf{V})$.

Corollary III.6. For $n_{\text{out}} = 2$, (11) reduces to $V_{\text{red}} = 2r(1-r)(|\mathbf{v}_{1,1}| \cdot |\mathbf{v}_{1,2}|)$.

¹Extension of the decomposition to higher dims. is left as future work.

²See https://en.wikipedia.org/wiki/Minimum_bounding_box.

Corollary III.7. For $n_{out} = 2$ or 3, the optimal r is $\frac{1}{2}$.

Lemma III.8. For $r = \frac{1}{2}$ and $n_{out} = 2$, the total volume reduction from repeatedly splitting \mathbf{u}_1 is $|\mathbf{W}_{:,1} \cdot \mathbf{u}_1| |\mathbf{W}_{:,2} \cdot \mathbf{u}_1|$.

Proof. Corollary III.6 is written in terms of \mathbf{u}_1 as

$$V_{\text{red}}(\mathbf{u}_1) = \frac{1}{2} (|\mathbf{W}_{:,1} \cdot \mathbf{u}_1| |\mathbf{W}_{:,2} \cdot \mathbf{u}_1|). \quad (12)$$

Say we bisect each of the two cells created in the first split; each will provide a further volume reduction of

$$V_{\text{red}}\left(\frac{1}{2}\mathbf{u}_1\right) = \left(\frac{1}{2}\right)^2 V_{\text{red}}(\mathbf{u}_1). \quad (13)$$

On the i -th round of bisecting every cell in the input set,

$$V_{\text{red}}\left(\left(\frac{1}{2}\right)^{i-1}\mathbf{u}_1\right) = \left(\left(\frac{1}{2}\right)^{i-1}\right)^2 V_{\text{red}}(\mathbf{u}_1), \quad (14)$$

per cell, with 2^i cells providing this volume reduction.

Thus, adding up all the volume reductions up to and including the i -th round, the total volume reduction is

$$V_{\text{red},i}(\mathbf{u}_1) = \sum_{j=1}^i 2^j \left(\left(\frac{1}{2}\right)^{j-1}\right)^2 V_{\text{red}}(\mathbf{u}_1) \quad (15)$$

$$= \sum_{j=1}^i \left(\frac{1}{2}\right)^{j-1} V_{\text{red}}(\mathbf{u}_1) \quad (16)$$

In the limit of partitioning one dimension over and over,

$$V_{\text{red},\infty}(\mathbf{u}_1) = \lim_{i \rightarrow \infty} V_{\text{red},i}(\mathbf{u}_1) = 2 \cdot V_{\text{red}}(\mathbf{u}_1). \quad (17)$$

□

This analysis justifies partitioning as a strategy to reduce output set volume/conservatism; extensions for nonlinear activations and N -layer NNs are left for future work.

IV. EXPERIMENTAL RESULTS

This section shows example partitions, applies the ideas to a robotic arm task, demonstrates better closed-loop behavior in collision avoidance, shows an ability to scale to various network sizes/architectures, and measures improvement along the time vs. tightness tradeoff.

A. Partitions for Different Output Shapes

The ability to partition efficiently for different output shapes is shown in Fig. 4 for a randomly initialized NN with 2 inputs, 2 outputs, and 50 nodes in hidden layer, *i.e.*, $(2, 50, 2)$, with ReLU activations, and input set $[0, 1] \times [0, 1]$. Each of (a-c) uses GSG with CROWN for 2 seconds. Recall that SG [13] would only return one output set for (a-c).

B. Comparisons to Baselines & Ablation Study

Four partitioning algorithms are compared in Fig. 5 for the same $(2, 50, 2)$ NN and input set. Each analyzer runs for 2 seconds to compute an estimated output set. The true output set is obtained by exhaustively sampling from the input space, and error is reported as percent extra area, $\frac{A_{\text{estimate}} - A_{\text{true}}}{A_{\text{true}}}$. The proposed partitioning algorithms GSG (c) and AGSG (d) use only 280 and 275 partitions respectively, and their approximation error is 0.018, which indicates more




Fig. 4: Input & Output Sets for Different Output Set Shapes. The estimated bounds (black) are “tight” when they are close to the bounds from exhaustive sampling (dashed red). The GSG partitioner with CROWN [10] propagator ran for 2 sec.

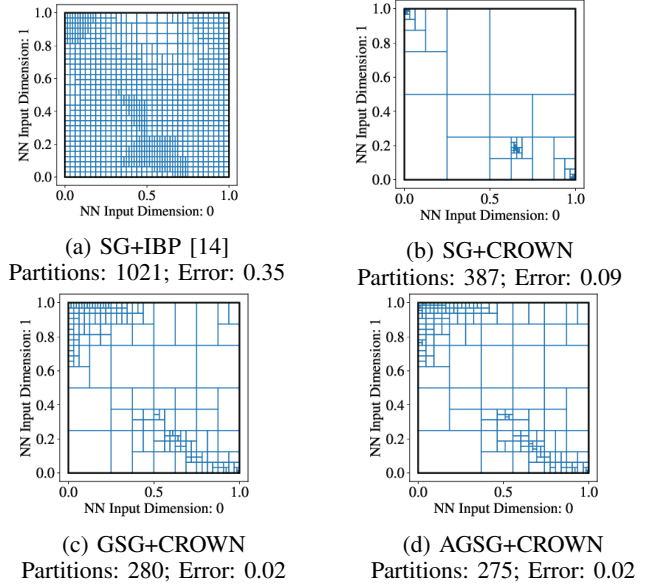


Fig. 5: Input partitions for a random $(2, 50, 2)$ ReLU NN analyzed for 2 seconds. New methods (b, c, d) reduce the number of input partitions and output set (convex hull) error.

than 79% improvement over SG -CROWN (b) and 95% over the state of the art SG-IBP [14]. In addition to quantitative improvement, Fig. 5 illustrates the input set partitions of each algorithm, which highlights how GSG refines different/fewer cells as SG, and that AGSG does not strictly make bisections.

C. Applications in Robotics and Control

1) *Reachable Set Analysis of Robotic Arm:* Borrowing the robotic arm model from [13], [14], we compare our algorithm to [14] for reachable set estimation of a forward kinematics model. Fig. 6a shows the 2 DOF robot arm, with input (θ_1, θ_2) as joint angles and output (x, y) as end effector position. The nonlinear dynamics are modeled by a small $(2, 5, 2)$ NN with tanh activations. We assume a time limit of 2 sec to approximate the convex hull of end effector positions from the set of joint angles $(\theta_1, \theta_2) \in [\frac{\pi}{3}, \frac{2\pi}{3}] \times [\frac{\pi}{3}, \frac{2\pi}{3}]$. As shown in Fig. 6a, AGSG-CROWN reduces the error from [14] by 96%. Only switching the partitioner (SG-IBP vs. (A)GSG-IBP) still achieves 80% error reduction. The estimated boundary is shown in (c, d).

2) *Multiagent Collision Avoidance:* Deep RL methods are popular in multiagent collision avoidance literature [21],

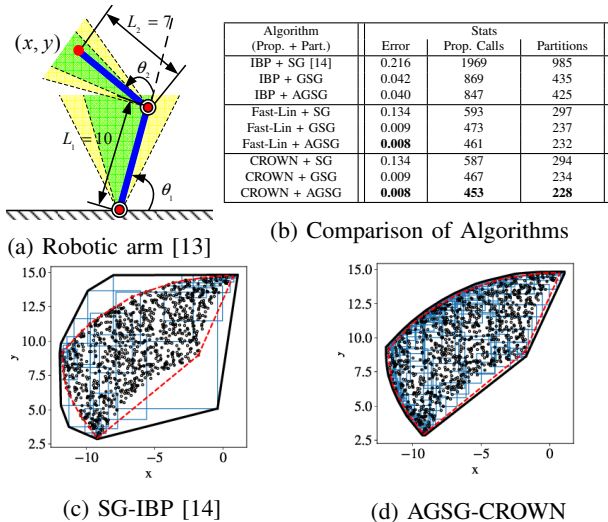


Fig. 6: Reachable set estimate of a robotic arm. GSG/AGSG-CROWN achieves 96% lower error than [14] (2 sec. computation).

[22], but they rarely account for measurement uncertainty. [23] proposed a certifiably robust deep RL algorithm, but the performance degraded with large input uncertainty (ϵ_{rob} in that work), partially due to overly conservative lower bounds from Fast-Lin [18].

The robust-but-conservative behavior caused by CROWN/Fast-Lin is shown in Fig. 7a ($\pm 0.5m$ uncertainty on the blue agent’s position at each timestep). By instead using GSG-CROWN to estimate worst-case Q-values (Fig. 7b), the orange agent reaches the goal much faster while still avoiding the blue agent. This improved behavior is a result of tighter estimates of worst-case Q-values, shown at a single timestep in Fig. 7c. For this experiment, a (11, 64, 64, 11) DQN (11 states & 11 discrete actions) was trained with perfect measurements in the gym-collision-avoidance environment [24]. Furthermore, this application is a case where only *lower bounds* on the NN outputs are needed, which motivates the use of our proposed algorithms that can focus computation toward this objective.

D. Scalability to Larger NNs

The proposed algorithms scale to bigger NNs as shown in Table I and improve on existing methods by orders of magnitude. For a given number of propagator calls (100), the average error (close to 0 is best) of 10 random NNs is reported for 4 different architectures (small, deep, higher dimension, LSTM), and for different boundary types. The (4, 100, 10) NN has uncertainty on the full 4D input set and the LSTM uses 2D uncertainty on the last timestep. While this work’s approaches scale well to deep NNs and various architectures, future work should consider the challenges from settings with even higher dimensional input uncertainties.

E. Computation-Performance Tradeoff

Throughout this paper, we have leveraged the idea of partitioning the input set to tighten the approximated bound-

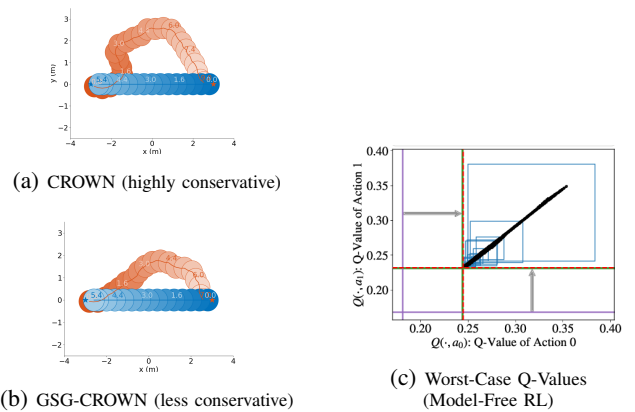


Fig. 7: Multiagent collision avoidance under uncertainty. In (a), a robust but conservative trajectory from a robust RL formulation [23] that used CROWN/Fast-Lin [10] to estimate worst-case Q-values under uncertainty on obstacle positions. In (b), the proposed GSG-CROWN algorithm enables the orange agent to reach its goal faster (7 vs. 9 sec) while still avoiding the blue agent. This improved behavior is a result of tighter estimates of worst-case Q-values, shown at a single timestep in (c) (purple \rightarrow green).

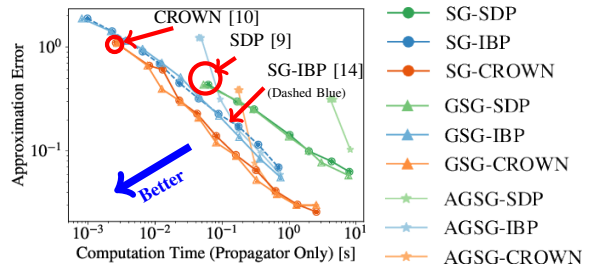


Fig. 8: Improvement in Computation Time vs. Accuracy Tradeoff. Colors indicate Propagator; markers indicate Partitioner. GSG-CROWN (orange triangles) requires $\sim 5\times$ less computation for the same accuracy as SG-IBP [14]. This work unified partitioning and propagation ideas to give many methods that exceed the prior state-of-art.

ary. To empirically show that bounds tighten with additional computational effort, we plot several combinations of partitioners and propagators in Fig. 8 over time. Each color corresponds to a propagator (IBP, CROWN, SDP) and each marker corresponds to a partitioner (SG, GSG, AGSG). This result uses the robotic arm model from before, but with ReLU activations.

A first key takeaway is that additional computation time leads to reduced error (increased tightness)³. Another key takeaway is that our framework provides many algorithms that exceed the performance of previous state-of-art algorithms [9], [10], [14]. Except the blue dashed line [14] and leftmost green/orange points [9], [10], all of the options are new algorithms proposed by this work. The analysis provided in the plots informs the choice of propagator and partitioner for a particular application with, say, a desired level of accuracy or budgeted resources (memory/computation). Overall,

³Comp. Time is reported only for propagator calls, which should dominate computation for larger NNs.

NN	Algorithm (Prop. + Part.)	Lower Bounds	Boundary Type ℓ_{∞} -ball	Convex Hull
(2, 100, 2)	IBP [7]	1.50	7.77×10^1	9.06
	IBP + SG [14]	2.47×10^{-1}	4.06	1.49
	IBP + GSG	1.70×10^{-1}	3.44	1.44
	Fast-Lin [11]	2.78×10^{-1}	4.62	1.90
	Fast-Lin + SG	1.70×10^{-3}	1.70×10^{-2}	1.12×10^{-1}
	Fast-Lin + GSG	3.94×10^{-3}	5.48×10^{-2}	7.23×10^{-2}
	CROWN [10]	2.15×10^{-1}	3.29	1.55
	CROWN + SG	1.34×10^{-3}	1.23×10^{-2}	1.09×10^{-1}
	CROWN + GSG	3.49×10^{-3}	5.32×10^{-2}	6.65×10^{-2}
	SDP [9]	1.20×10^{-1}	1.90	1.06
(2, 100, 100, 100, 100, 2)	IBP [7]	1.69×10^2	8.17×10^9	1.07×10^5
	IBP + SG [14]	3.32×10^1	3.16×10^8	2.10×10^4
	IBP + GSG	3.07×10^1	2.67×10^8	1.93×10^4
	Fast-Lin [11]	2.32	1.57×10^6	1.48×10^3
	Fast-Lin + SG	2.65×10^{-4}	3.83×10^{-1}	2.91×10^{-1}
	Fast-Lin + GSG	9.32×10^{-5}	2.12×10^{-1}	2.30×10^{-1}
	CROWN [10]	8.96×10^{-1}	2.42×10^5	5.74×10^2
	CROWN + SG	1.61×10^{-4}	2.03×10^{-1}	1.91×10^{-1}
	CROWN + GSG	5.40×10^{-5}	1.18×10^{-1}	1.65×10^{-1}
	(4, 100, 10)	IBP [7]	3.11×10^1	1.41×10^{17}
IBP + SG [14]		1.40×10^1	6.89×10^{13}	-
IBP + GSG		1.33×10^1	4.37×10^{13}	-
Fast-Lin [11]		6.18	5.95×10^{10}	-
Fast-Lin + SG		8.24×10^{-1}	7.81×10^3	-
Fast-Lin + GSG		7.44×10^{-1}	4.59×10^3	-
CROWN [10]		4.51	4.01×10^9	-
CROWN + SG		5.60×10^{-1}	8.17×10^2	-
CROWN + GSG		5.00×10^{-1}	5.52×10^2	-
LSTM (8, 8, 64, 2)		IBP [7]	1.56×10^{-2}	1.16×10^2
	IBP + SG [14]	2.44×10^{-3}	5.55	1.83
	IBP + GSG	1.80×10^{-3}	4.90	1.80

TABLE I: Approximation Error (closer to 0 is better) for four different types of NNs: small, deep, higher dimension, and LSTM. Reported values are average error across 10 randomly initialized NNs after up to 100 propagator calls.

for this task GSG-CROWN almost always provides the best accuracy vs. computation time tradeoff, requiring $\sim 5 \times$ less computation for the same accuracy as SG-IBP [14].

V. CONCLUSION

This work proposed a suite of algorithms for online robustness analysis of NNs that can provide confidence in NN decisions under uncertainty. We build on recent work for handling large uncertainties by proposing new, flexible partitioning algorithms and give theoretical rationale for partitioning as a strategy for reducing conservatism. Furthermore, we show how recent methods that efficiently relax NN nonlinearities can be unified with partitioning in a single framework, which provides many new state-of-art algorithmic choices for robotics applications. Along with showing improved aggregate performance on random NNs with various sizes/architectures, we show how these ideas can be applied to other learning tasks for control systems, showing a 22% improvement in robust RL for multiagent collision avoidance and a 96% reduction in conservatism for a learned robotic arm kinematic model.

ACKNOWLEDGMENT

This work was supported by Ford Motor Company.

REFERENCES

- [1] R. Ehlers, “Formal verification of piece-wise linear feed-forward neural networks,” in *ATVA*, 2017.
- [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 *Computer Aided Verification - 29th International Conference, CAV 2017, Heidelberg, Germany, July 24-28, 2017, Proceedings, Part I*, pp. 97–117, 2017.
- [3] X. Huang, M. Kwiatkowska, S. Wang, and M. Wu, “Safety verification of deep neural networks,” in *Computer Aided Verification* (R. Majumdar and V. Kunčak, eds.), (Cham), pp. 3–29, Springer International Publishing, 2017.

- [4] A. Lomuscio and L. Maganti, “An approach to reachability analysis for feed-forward relu neural networks,” *CoRR*, vol. abs/1706.07351, 2017.
- [5] V. Tjeng, K. Y. Xiao, and R. Tedrake, “Evaluating robustness of neural networks with mixed integer programming,” in *International Conference on Learning Representations (ICLR)*, 2019.
- [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 *2018 IEEE Symposium on Security and Privacy (SP)*, pp. 3–18, May 2018.
- [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,” *arXiv preprint arXiv:1810.12715*, 2018.
- [8] A. Raghunathan, J. Steinhardt, and P. Liang, “Certified defenses against adversarial examples,” in *International Conference on Learning Representations (ICLR)*, 2018.
- [9] M. Fazlyab, M. Morari, and G. J. Pappas, “Safety verification and robustness analysis of neural networks via quadratic constraints and semidefinite programming,” *arXiv preprint arXiv:1903.01287*, 2019.
- [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 *Advances in neural information processing systems*, pp. 4939–4948, 2018.
- [11] 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 *International Conference on Machine Learning (ICML)*, 2018.
- [12] G. Singh, T. Gehr, M. Mirman, M. Püschel, and M. Vechev, “Fast and effective robustness certification,” in *Advances in Neural Information Processing Systems*, pp. 10802–10813, 2018.
- [13] W. Xiang, H.-D. Tran, and T. T. Johnson, “Output reachable set estimation and verification for multilayer neural networks,” *IEEE transactions on neural networks and learning systems*, vol. 29, no. 11, pp. 5777–5783, 2018.
- [14] W. Xiang, H.-D. Tran, X. Yang, and T. T. Johnson, “Reachable set estimation for neural network control systems: A simulation-guided approach,” *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [15] B. G. Anderson, Z. Ma, J. Li, and S. Sojoudi, “Tightened convex relaxations for neural network robustness certification,” *arXiv preprint arXiv:2004.00570*, 2020.
- [16] V. Rubies-Royo, R. Calandra, D. M. Stipanovic, and C. Tomlin, “Fast neural network verification via shadow prices,” *arXiv preprint arXiv:1902.07247*, 2019.
- [17] H. Salman, G. Yang, H. Zhang, C.-J. Hsieh, and P. Zhang, “A convex relaxation barrier to tight robustness verification of neural networks,” in *Advances in Neural Information Processing Systems*, pp. 9835–9846, 2019.
- [18] T.-W. Weng, H. Zhang, H. Chen, Z. Song, C.-J. Hsieh, D. Boning, I. S. Dhillon, and L. Daniel, “Towards fast computation of certified robustness for relu networks,” *arXiv preprint arXiv:1804.09699*, 2018.
- [19] G. Singh, R. Ganvir, M. Püschel, and M. Vechev, “Beyond the single neuron convex barrier for neural network certification,” in *Advances in Neural Information Processing Systems*, pp. 15098–15109, 2019.
- [20] H. Zhang, H. Chen, C. Xiao, S. Gowal, R. Stanforth, B. Li, D. Boning, and C.-J. Hsieh, “Towards stable and efficient training of verifiably robust neural networks,” *arXiv preprint arXiv:1906.06316*, 2019.
- [21] P. Long, T. Fan, X. Liao, W. Liu, H. Zhang, and J. Pan, “Towards optimally decentralized multi-robot collision avoidance via deep reinforcement learning,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6252–6259, IEEE, 2018.
- [22] Y. F. Chen, M. Liu, M. Everett, and J. P. How, “Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning,” in *2017 IEEE international conference on robotics and automation (ICRA)*, pp. 285–292, IEEE, 2017.
- [23] M. Everett, B. Lutjens, and J. P. How, “Certified adversarial robustness for deep reinforcement learning,” *arXiv preprint arXiv:2004.06496*, 2020.
- [24] M. Everett, Y. F. Chen, and J. P. How, “Motion planning among dynamic, decision-making agents with deep reinforcement learning,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, (Madrid, Spain), Sept. 2018.