Tight Neural Network Verification via Semidefinite Relaxations and Linear Reformulations

Jianglin Lan[1], **Yang Zheng**[2], **Alessio Lomuscio**[1]
[1] Department of Computing, Imperial College London, UK
[2] Department of Electrical and Computer Engineering, University of California San Diego, USA
j.lan@imperial.ac.uk, zhengy@eng.ucsd.edu, a.lomuscio@imperial.ac.uk

Abstract

We present a novel semidefinite programming (SDP) relaxation that enables tight and efficient verification of neural networks. The tightness is achieved by combining SDP relaxations with valid linear cuts, constructed by using the reformulation-linearisation technique (RLT). The computational efficiency results from a layerwise SDP formulation and an iterative algorithm for incrementally adding RLT-generated linear cuts to the verification formulation. The layer RLT-SDP relaxation here presented is shown to produce the tightest SDP relaxation for ReLU neural networks available in the literature. We report experimental results based on MNIST neural networks showing that the method outperforms the state-of-the-art methods while maintaining acceptable computational overheads. For networks of approximately 10k nodes (1k, respectively), the proposed method achieved an improvement in the ratio of certified robustness cases from 0% to 82% (from 35% to 70%, respectively).

Introduction

While progress in training methods for neural networks (NNs) continues, it is well-known that NNs are susceptible to adversarial attacks (Goodfellow, Shlens, and Szegedy 2014). This is highly problematic for uses of NNs in safety-critical systems such as the aircraft domain (Kouvaros et al. 2021; Akintunde et al. 2020b,a; Julian and Kochenderfer 2021; Manzanas Lopez et al. 2021) or in any application where miss-classifications need to be minimised. The area of verification of NNs (Liu et al. 2020; Bak, Liu, and Johnson 2021) aims to develop methods to guarantee that NNs are robust with respect to small perturbations, with particular emphasis to noise perturbations.

Existing NN verification methods can be divided into two categories (Liu et al. 2020): *complete* and *incomplete* approaches. Complete approaches guarantee to resolve any verification query, but may incur high computational cost. Incomplete approaches normally leverage forms of convex over-approximations of NNs to enable faster verification. While incomplete approaches tend to scale to larger networks, the looser their approximation is, the more likely it is that the approach may be unable to verify the problem instance. Thus, one central objective in incomplete approaches is to develop tighter convex approximations while retaining computation efficiency (Salman et al. 2019).

In this paper we provide a novel relaxation method which combines semidefinite programming (SDP) (Vandenberghe and Boyd 1996) with the reformulation-linearisation technique (RLT) (Anstreicher 2009) to verify NNs. The new SDP relaxation is provably tighter than existing SDP approaches, whilst enjoying a competitive efficiency.

Related work. Complete approaches to the verification of NNs are typically based on mixed-integer linear programming (MILP) (Bastani et al. 2016; Lomuscio and Maganti 2017; Tjeng, Xiao, and Tedrake 2019; Anderson et al. 2020; Botoeva et al. 2020), satisfiability modulo theories (Katz et al. 2019; Ehlers 2017) or bound propagations combined with input refinement (Henriksen and Lomuscio 2020, 2021; Wang et al. 2021; Hashemi, Kouvaros, and Lomuscio 2021). While these approaches can provide theoretical termination guarantees, their scalability to large NNs is often problematic. Incomplete approaches for NN verification are normally based on bound propagations (Weng et al. 2018; Singh et al. 2019a; Tjandraatmadja et al. 2020; Müller et al. 2021), combinations between linear programming (LP) and relaxations (Ehlers 2017), or duality relaxation (Dvijotham et al. 2018; Wong and Kolter 2018; Dathathri et al. 2020). The triangle relaxation (Ehlers 2017) gives the tightest convex approximation of a single ReLU node and has inspired several other approaches (Salman et al. 2019; Li et al. 2020). While these methods often achieve state-of-the-art (SoA) performance, they have limited efficacy: even optimally tuned LP-based convex relaxations may fail to obtain tight bounds on the certified robustness ratio (Salman et al. 2019).

Two lines of research have attempted to alleviate this problem. The first aims to provide tighter LP relaxations by exploring interdependencies among multiple neurons and inputs, *e.g.*, `DeepPoly` (Singh et al. 2019a), `kPoly` (Singh et al. 2019b), `OptC2V` (Tjandraatmadja et al. 2020), and `PRIMA` (Müller et al. 2021). The second seeks alternative, stronger relaxations beyond LPs. The most promising relaxation combining tightness with efficiency is presently based on SDPs (Raghunathan, Steinhardt, and Liang 2018; Fazlyab, Morari, and Pappas 2022; Batten et al. 2021).

It has been empirically observed that the relaxations generated with the SDP method in (Raghunathan, Steinhardt, and Liang 2018) are considerably tighter than standard LP

relaxations. Using geometric techniques, it has been shown that the SDP relaxation for a variant of the NN verification problem is exact over a single hidden layer under mild assumptions, but becomes loose for several hidden layers (Zhang 2020). To obtain tighter SDP relaxations, effective linear cuts were identified in (Batten et al. 2021); non-convex cuts were investigated in (Ma and Sojoudi 2020).

SDPs are harder to solve than LPs. To overcome this, a dual SDP relaxation was formulated and subsequently solved using a subgradient algorithm in (Dathathri et al. 2020). A layer SDP relaxation has been recently proposed in (Batten et al. 2021) by exploiting cascading network structures based on graph decomposition (Zheng, Fantuzzi, and Papachristodoulou 2021). To the best of our knowledge, layer SDP provides the tightest relaxations that have so far been achieved by combining SDP relaxations with triangle relaxations (Ehlers 2017). Even so, the relaxation gap is still considerable in large NNs, as observed in (Batten et al. 2021). This leads to an increasing number of verification queries that cannot be resolved as the model size increases, thereby limiting the applicability of the approach.

Contributions. In this paper, we advance the SoA SDP relaxations for NN verification by using the RLT (Anstreicher 2009) in this context. Specifically, we propose a new layer RLT-SDP relaxation with valid linear cuts obtained from RLT that offers provably tighter relaxations. The linear cuts capture both *inter-layer dependencies* and *intra-layer interactions* of the network, which are presently not exploited in the existing relaxation methods. Due to the computational cost of using large numbers of linear cuts, we refine this method by introducing an iterative algorithm to integrate the RLT-generated linear cuts and the SDP relaxation. At each iteration only a portion of linear cuts are added, with their priorities being determined by the network weights. Our theoretical analysis shows that the relaxations here obtained are provably tighter than any other approach previously considered. In the experiments we report, the method obtained considerably tighter relaxations than the present SoA leading to several more queries being answerable.

Problem Statement and Preliminaries

Notation. We use the symbol $\mathbb{R}^n$ to denote the n-dimensional Euclidean space. We use $\text{diag}(X)$ to stack the main diagonals of matrix X as a column, and $|X|$ to get its absolute value element-wise. We use $\odot$ to denote the element-wise product, and $\mathbb{I}_b$ with $b > 0$ to denote a sequence of nonzero integers from 0 to b. We use $\|\cdot\|_\infty$ to refer to the standard ℓ_∞ norm.

NN verification problem. We focus on feed-forward fully-connected NNs with ReLU activations. A NN $f(x_0) : \mathbb{R}^{n_0} \to \mathbb{R}^{n_{L+1}}$ with L hidden layers, n_0 inputs and n_{L+1} outputs is defined as follows: x_0 is the input, $f(x_0)$ is the output, and $\hat{x}_i, x_i \in \mathbb{R}^{n_i}$, $i = 1, 2, \ldots, L$, are the pre-activation and activation vectors of the i-th layer, respectively. The NN output is $f(x_0) = W_L x_L + b_L$ with $x_{i+1} = \text{ReLU}(\hat{x}_{i+1})$ and $\hat{x}_{i+1} = W_i x_i + b_i$, $i \in \mathbb{I}_{L-1}$, where $W_i \in \mathbb{R}^{n_{i+1} \times n_i}$ and $b_i \in \mathbb{R}^{n_{i+1}}$ are the weights and biases, respectively. For a vector $z \in \mathbb{R}^n$, the ReLU function is defined as $\text{ReLU}(z) = [\max(z_1, 0), \ \cdots, \ \max(z_n, 0)]^\mathsf{T}$. We focus on classification networks whereby an input x_0 is assigned to the class associated with the NN output with the highest value: $j^\star = \arg\max_{j=1,2,\cdots,n_{L+1}} f(x_0)_j$.

We now present the local robustness verification problem.

Definition 1. *Given a NN $f : \mathbb{R}^{n_0} \to \mathbb{R}^{n_{L+1}}$, an input $\bar{x}$ and a perturbation radius $\epsilon \in \mathbb{R}$, f is robust on $\bar{x}$ if $f(\bar{x})_{j^\star} - f(x_0)_j > 0$ when $j \neq j^\star$, for all x_0 s.t. $\|x_0 - \bar{x}\|_\infty \leq \epsilon$.*

Given a NN f, an input $\bar{x}$ and a perturbation ϵ, the local robustness problem concerns determining whether it is the case that f meets Definition 1 for $\bar{x}$ and ϵ. This problem can be formulated and solved as an optimisation problem:

$$
\begin{aligned}
\gamma^* := \min_{\{x_i\}_{i=0}^L} \ & c^\mathsf{T} x_L + c_0 \\
\text{s.t.} \ & x_{i+1} = \text{ReLU}(W_i x_i + b_i),\ i \in \mathbb{I}_{L-1}, && \text{(1a)} \\
& \|x_0 - \bar{x}\|_\infty \leq \epsilon, && \text{(1b)} \\
& l_{i+1} \leq x_{i+1} \leq u_{i+1},\ i \in \mathbb{I}_{L-1}, && \text{(1c)}
\end{aligned}
$$

where $c^\mathsf{T} = W_L(j^\star, :) - W_L(j, :)$ and $c_0 = b_L(j^\star) - b_L(j)$. l_{i+1} and u_{i+1} are the lower and upper bounds of the activation vectors, which can be computed using bound propagation methods, see *e.g.*, (Wong and Kolter 2018; Wang et al. 2018). The optimisation problem (1) is solved for every potential adversarial target $j \neq j^\star$. If $\gamma^* > 0$ in all cases, then the network is robust on $\bar{x}$ against the adversarial input x_0.

Due to the nonlinear ReLU constraint (1a), the problem (1) is non-convex and thus hard to solve generally. In the literature, two SDP relaxations have been proposed: a global SDP relaxation (Raghunathan, Steinhardt, and Liang 2018) and a layer SDP relaxation (Batten et al. 2021).

Global SDP relaxation. The ReLU constraint (1a) is equivalent to a set of linear and quadratic constraints:

$$
\begin{aligned}
& x_{i+1} \geq 0,\ x_{i+1} \geq W_i x_i + b_i,\ i \in \mathbb{I}_{L-1}, \\
& x_{i+1} \odot (x_{i+1} - W_i x_i - b_i) = 0,\ i \in \mathbb{I}_{L-1}.
\end{aligned} \tag{2}
$$

The input constraints in (1b) and (1c) can be equivalently represented as the quadratic constraints:

$$
x_i \odot x_i - (l_i + u_i) \odot x_i + l_i \odot u_i \leq 0,\ i \in \mathbb{I}_L, \tag{3}
$$

where $l_0 = \bar{x} - \epsilon\mathbf{1}$, $u_0 = \bar{x} + \epsilon\mathbf{1}$ and $\mathbf{1}$ denotes column of ones. Replacing (1a)-(1c) with (2)-(3) yields an equivalent quadratically constrained quadratic programming (QCQP):

$$
\gamma^* := \left\{ \min_{\{x_i\}_{i=0}^L} \ c^\mathsf{T} x_L + c_0 \mid (2), (3) \right\}. \tag{4}
$$

This QCQP problem is still non-convex due to the quadratic constraints. The techniques of polynomial lifting (Parrilo 2000; Lasserre 2009) can be used to reformulate them as linear constraints. The lifting matrix P is defined as

$$
P = \mathbf{x}\mathbf{x}^\mathsf{T}, \text{ with } \mathbf{x} = [1,\ x_0^\mathsf{T},\ x_1^\mathsf{T},\ \cdots,\ x_L^\mathsf{T}]^\mathsf{T} \in \mathbb{R}^{\bar{n}}, \tag{5}
$$

where $\bar{n} = 1 + \sum_{i=0}^L n_i$. The above can be reformulated as:

$$
P \succeq 0,\ P[1] = 1,\ \text{rank}(P) = 1. \tag{6}
$$

By using (5) and (6) and dropping $\text{rank}(P) = 1$, the QCQP (4) is relaxed into a global SDP (Raghunathan, Steinhardt, and Liang 2018):

$$\gamma^*_{\text{GlobalSDP}} := \min_{P} \; c^\mathsf{T} P[x_L] + c_0$$

$$\text{s.t. } P[x_{i+1}] \geq 0,\; P[x_{i+1}] \geq W_i P[x_i] + b_i, i \in \mathbb{I}_{L-1}, \quad (7a)$$

$$\text{diag}(P[x_{i+1}x_{i+1}^\mathsf{T}] - W_i P[x_i x_{i+1}^\mathsf{T}]) - b_i \odot P[x_{i+1}] = 0,\; i \in \mathbb{I}_{L-1}, \quad (7b)$$

$$\text{diag}(P[x_i x_i^\mathsf{T}]) - (l_i + u_i) \odot P[x_i] + l_i \odot u_i \leq 0,\; i \in \mathbb{I}_L, \quad (7c)$$

$$P[1] = 1, P \succeq 0, \quad (7d)$$

where the symbolic indexing $P[\cdot]$ is used to index the elements of P. Since $\text{rank}(P) = 1$ is dropped, problem (7) gives a relaxed solution to the QCQP, *i.e.*, $\gamma^*_{\text{GlobalSDP}} \leq \gamma^*$.

Layer SDP relaxation. The layer SDP relaxation exploits the deep structure of the layers of a NN, where the activation vector of a layer depends only on its immediate preceding layer. This structure is exploited in (Batten et al. 2021) to develop a layer-based SDP formulation of (7), where the dimension of the matrix constraint is reduced, thus improving the computational efficiency. In this work, instead of using a single large lifting matrix P for the entire network, each hidden layer i is assigned a smaller matrix P_i defined as

$$P_i = \mathbf{x}_i \mathbf{x}_i^\mathsf{T}, \text{ with } \mathbf{x}_i = [1,\; x_i^\mathsf{T},\; x_{i+1}^\mathsf{T}]^\mathsf{T} \in \mathbb{R}^{\bar{n}_i}, \quad (8)$$

where $\bar{n}_i = 1 + n_i + n_{i+1}$. Similar to (6), the constraint $P_i = \mathbf{x}_i \mathbf{x}_i^\mathsf{T}$ is equivalent to $P_i \succeq 0,\; P_i[1] = 1,\; \text{rank}(P_i) = 1$. By using (8), the layer SDP relaxation is formulated as

$$\gamma^*_{\text{LayerSDP}} := \min_{\{P_i\}_{i=0}^{L-1}} \; c^\mathsf{T} P_{L-1}[x_L] + c_0$$

$$\text{s.t. } P_i[x_{i+1}] \geq 0,\; i \in \mathbb{I}_{L-1}, \quad (9a)$$

$$P_i[x_{i+1}] \geq W_i P_i[x_i] + b_i,\; i \in \mathbb{I}_{L-1}, \quad (9b)$$

$$\text{diag}(P_i[x_{i+1}x_{i+1}^\mathsf{T}] - W_i P_i[x_i x_{i+1}^\mathsf{T}]) - b_i \odot P_i[x_{i+1}] = 0,\; i \in \mathbb{I}_{L-1}, \quad (9c)$$

$$\text{diag}(P_i[x_i x_i^\mathsf{T}]) - (l_i + u_i) \odot P_i[x_i] + l_i \odot u_i \leq 0,\; i \in \mathbb{I}_{L-1}, \quad (9d)$$

$$\text{diag}(P_{L-1}[x_L x_L^\mathsf{T}]) - (l_L + u_L) \odot P_{L-1}[x_L] + l_L \odot u_L \leq 0, \quad (9e)$$

$$P_i[1] = 1,\; P_i \succeq 0,\; i \in \mathbb{I}_{L-1}, \quad (9f)$$

$$P_i[\bar{x}_{i+1}\bar{x}_{i+1}^\mathsf{T}] = P_{i+1}[\bar{x}_{i+1}\bar{x}_{i+1}^\mathsf{T}],\; i \in \mathbb{I}_{L-2}, \quad (9g)$$

$$P_i[x_{i+1}] \leq A_i P_i[x_i] + B_i,\; i \in \mathbb{I}_{L-1}, \quad (9h)$$

where $\bar{x}_{i+1} = [1,\; x_{i+1}^\mathsf{T}]^\mathsf{T}$, $A_i = k_i \odot W_i$, $B_i = k_i \odot (b_i - \hat{l}_{i+1}) + \text{ReLU}(\hat{l}_{i+1})$, $k_i = (\text{ReLU}(\hat{u}_{i+1}) - \text{ReLU}(\hat{l}_{i+1}))/(\hat{u}_{i+1} - \hat{l}_{i+1})$, with $\hat{u}_{i+1}, \hat{l}_{i+1}$ being upper and lower bounds of the pre-activation vector $\hat{x}_{i+1}$ (Wong and Kolter 2018; Wang et al. 2018). Note that the constraint (9g) appears to ensure input-output consistency between layers. Compared to (7), the new constraint (9h) is obtained from the triangle relaxation. It is shown in (Batten et al. 2021) that without (9h), the layer SDP relaxation (9) is equivalent to the global SDP relaxation (7) based on graph decomposition (Vandenberghe and Andersen 2015; Zheng 2019). Also, this layer SDP relaxation achieves faster verification than the global SDP relaxation by dealing with lower dimensional constraints, and obtains a tighter relaxation by introducing (9h), *i.e.*, $\gamma^*_{\text{GlobalSDP}} \leq \gamma^*_{\text{LayerSDP}} \leq \gamma^*$.

Source of SDP relaxation gap. Let $\tilde{x} = [x_0^\mathsf{T},\; \cdots,\; x_L^\mathsf{T}]^\mathsf{T}$ and $P[\tilde{x}\tilde{x}^\mathsf{T}] = \tilde{x}\tilde{x}^\mathsf{T}$. It follows from (6) that $P = \mathbf{x}\mathbf{x}^\mathsf{T}$ is equivalent to $P[1] = 1,\; P[\tilde{x}\tilde{x}^\mathsf{T}] = \tilde{x}\tilde{x}^\mathsf{T}$. Observe that the global SDP relaxation (7) only includes (Eq. (7d)) the relaxed constraints $P[1] = 1$ and $P[\tilde{x}\tilde{x}^\mathsf{T}] \succeq \tilde{x}\tilde{x}^\mathsf{T}$, reformulated as $P \succeq 0$ via Schur complement. A consequence of this is that P may not be sufficiently bounded, thereby resulting in loose SDP solutions. The same argument applies for the layer SDP relaxation (9).

To bridge the gap of the global SDP relaxation, the nonconvex constraint $P[\tilde{x}\tilde{x}^\mathsf{T}] \preceq \tilde{x}\tilde{x}^\mathsf{T}$ is imposed in (Ma and Sojoudi 2020) via secant approximation. Valid linear cuts are generated by an iterative algorithm, where the global SDP relaxation together with an LP need to be solved at each iteration. Unfortunately, it is known that global SDP relaxation itself is already computationally expensive. Therefore, the tightening in (Ma and Sojoudi 2020) is unlikely to lead to scalable NN verification.

Tightening Layer SDP Relaxation via RLT

Having highlighted the existing relaxation gap in the SoA, we now present an approach for tightening the SDP relaxation for NN verification while retaining an acceptable computational overhead. Our method combines layer SDP relaxations (9) with the RLT (Anstreicher 2009).

Motivation. We first denote a few terms for each layer of a NN: $\tilde{x}_{i+1} = [x_i^\mathsf{T},\; x_{i+1}^\mathsf{T}]^\mathsf{T}$, $\tilde{l}_{i+1} = [l_i^\mathsf{T},\; l_{i+1}^\mathsf{T}]^\mathsf{T}$ and $\tilde{u}_{i+1} = [u_i^\mathsf{T},\; u_{i+1}^\mathsf{T}]^\mathsf{T}$. Since $0 \leq \tilde{l}_i \leq \tilde{x}_{i+1} \leq \tilde{u}_i$, we have $\tilde{x}_{i+1}\tilde{l}_{i+1}^\mathsf{T} \leq \tilde{x}_{i+1}\tilde{x}_{i+1}^\mathsf{T} \leq \tilde{x}_{i+1}\tilde{u}_{i+1}^\mathsf{T}$. These nonlinear constraints can be reformulated as linear constraints on the elements of P_i:

$$P_i[\tilde{x}_{i+1}]\tilde{l}_{i+1}^\mathsf{T} \leq P_i[\tilde{x}_{i+1}\tilde{x}_{i+1}^\mathsf{T}] \leq P_i[\tilde{x}_{i+1}]\tilde{u}_{i+1}^\mathsf{T}. \quad (10)$$

The method aims to bound $P_i[\tilde{x}_{i+1}\tilde{x}_{i+1}^\mathsf{T}]$ within the region given in (10). The constraints in (10) are linear and could be directly added to (9). However, they introduce $2(n_i + n_{i+1})^2$ new inequalities, thereby increasing the computational effort required to solve the verification problem. Therefore, it is desirable to develop efficient strategies for imposing the constraints in (10). In the following we: (i) use RLT to construct valid linear cuts that are provably stronger than (10), and (ii) provide a computationally-efficient strategy for integrating the linear cuts with the layer SDP relaxation (9).

Construction of valid linear cuts using RLT. RLT involves the construction of valid linear cuts on the lifting matrices $\{P_i\}_{i=0}^{L-1}$ by using products of the existing linear constraints in (9) on the original variables $\{x_i\}_{i=0}^{L}$. Under the constraints (9a) and (9d), the variables x_i and x_{i+1} satisfy: $x_i \geq 0$, $x_i - l_i \geq 0$, $x_i - u_i \leq 0$, $x_{i+1} - l_{i+1} \geq 0$, $x_{i+1} - u_{i+1} \leq 0$. These can be used to construct the constraints: $x_i l_i^\mathsf{T} \leq x_i x_i^\mathsf{T} \leq x_i u_i^\mathsf{T}$, $(x_{i+1} - l_{i+1})(x_i - l_i)^\mathsf{T} \geq 0$,

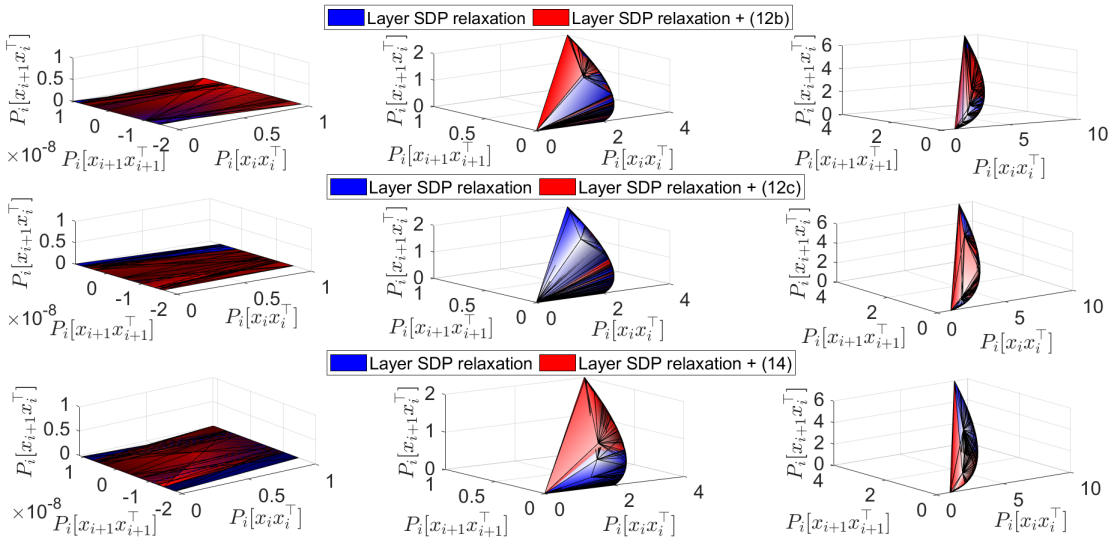

Figure 1: Feasible region of the triple $(P_i[x_i x_i^\mathsf{T}], P_i[x_{i+1} x_{i+1}^\mathsf{T}], P_i[x_{i+1} x_i^\mathsf{T}])$ by adding linear cuts (12b), (12c) or (14), with $x_{i+1} = \text{ReLU}(x_i - 1)$ and $\hat{l}_{i+1} \le x_i - 1 \le \hat{u}_{i+1}$. Left to right columns: 1) *inactive* neuron $\hat{l}_{i+1} = -1$, $\hat{u}_{i+1} = 0$; 2) *unstable* neuron $\hat{l}_{i+1} = -1$, $\hat{u}_{i+1} = 1$; 3) *strictly active* neuron $\hat{l}_{i+1} = 0$, $\hat{u}_{i+1} = 2$.

$(x_{i+1} - l_{i+1})(x_i - u_i)^\mathsf{T} \le 0$, $(x_i - l_i)(x_{i+1} - u_{i+1})^\mathsf{T} \le 0$. By using (8), these nonlinear constraints are linearised as

$$P_i[x_i] l_i^\mathsf{T} \le P_i[x_i x_i^\mathsf{T}] \le P_i[x_i] u_i^\mathsf{T}, \tag{11a}$$
$$P_i[x_{i+1} x_i^\mathsf{T}] \ge P_i[x_{i+1}] l_i^\mathsf{T} + l_{i+1}(P_i[x_i^\mathsf{T}] - l_i^\mathsf{T}), \tag{11b}$$
$$P_i[x_{i+1} x_i^\mathsf{T}] \le P_i[x_{i+1}] u_i^\mathsf{T} + l_{i+1}(P_i[x_i^\mathsf{T}] - u_i^\mathsf{T}), \tag{11c}$$
$$P_i[x_i x_{i+1}^\mathsf{T}] \le P_i[x_i] u_{i+1}^\mathsf{T} + l_i(P_i[x_{i+1}^\mathsf{T}] - u_{i+1}^\mathsf{T}). \tag{11d}$$

We now make the following remarks.

Observation 1. *The linear cuts (11a) - (11d) imply (10).*

Observation 2. *The existing constraints (9d) and (9f) are stronger than the first part of (11a); while (9d) is stronger than the diagonal components of the second part of (11a).*

These observations show that the targeted bounding in (10) can be realised by adding to the layer SDP relaxation (9) the following linear cuts for each P_i, $i \in \mathbb{I}_{L-1}$:

$$P_i[x_i x_i^\mathsf{T}] \le P_i[x_i] u_i^\mathsf{T}, \tag{12a}$$
$$P_i[x_{i+1} x_i^\mathsf{T}] \ge P_i[x_{i+1}] l_i^\mathsf{T} + l_{i+1}(P_i[x_i^\mathsf{T}] - l_i^\mathsf{T}), \tag{12b}$$
$$P_i[x_{i+1} x_i^\mathsf{T}] \le \min\{P_i[x_{i+1}] u_i^\mathsf{T} + l_{i+1}(P_i[x_i^\mathsf{T}] - u_i^\mathsf{T}), \\ u_{i+1} P_i[x_i^\mathsf{T}] + (P_i[x_{i+1}] - u_{i+1}) l_i^\mathsf{T}\}, \tag{12c}$$

where the diagonal components of (12a) are redundant.

It has been shown above that adding the linear cuts in (12) to the layer SDP relaxation (9) is efficient to bound $P_i[\tilde{x}_{i+1} \tilde{x}_{i+1}^\mathsf{T}]$ and subsequently the matrix P_i. Problem (9) also has other existing linear constraints (9a), (9b) and (9h) that can be used to construct the new constraints:

$$(x_{i+1} - W_i x_i - b_i)(x_{i+1} - W_i x_i - b_i)^\mathsf{T} \ge 0, \tag{13a}$$
$$(x_{i+1} - A_i x_i - B_i) x_{i+1}^\mathsf{T} \le 0, \tag{13b}$$
$$(x_{i+1} - A_i x_i - B_i)(x_{i+1} - A_i x_i - B_i)^\mathsf{T} \ge 0. \tag{13c}$$

Observation 3. *Linear cut (13a) is weaker than (9f); while (13c) is weaker than the conjunction of (9a) - (9c) and (9h).*

Observation 4. *Adding the linear cut (13b) can tighten the layer SDP relaxation, but only its off-diagonals cut the feasible region, while the diagonals are implied by (9c).*

These observations reveal that the constraint (13b) has not been included in the layer SDP relaxation (9) and it can narrow the relaxation gap. By defining $P_i[x_{i+1} x_{i+1}^\mathsf{T}] = x_{i+1} x_{i+1}^\mathsf{T}$ and $P_i[x_i x_{i+1}^\mathsf{T}] = x_i x_{i+1}^\mathsf{T}$ and recalling that $P_i[x_{i+1}] = P_{i+1}[x_{i+1}]$ under (9g), the constraints (12a) and (13b) are merged as a linear cut for each P_i, $i \in \mathbb{I}_{L-1}$:

$$P_i[x_{i+1} x_{i+1}^\mathsf{T}] \le \\ \min\{P_i[x_{i+1}] u_{i+1}^\mathsf{T},\ A_i P_i[x_i x_{i+1}^\mathsf{T}] + B_i P_i[x_{i+1}^\mathsf{T}]\}. \tag{14}$$

When $i = 0$, $P_1[x_0 x_0^\mathsf{T}] \le P_1[x_0] u_0^\mathsf{T}$ is also needed.

Integration of linear cuts with the layer SDP relaxation. The above analysis identifies the valid linear cuts (12b), (12c) and (14) for each matrix P_i, $i \in \mathbb{I}_{L-1}$. Adding them to (9) yields the layer RLT-SDP relaxation:

$$\begin{aligned} \gamma^*_{\text{RLT-SDP}} := \min_{\{P_i\}_{i=0}^{L-1}} \ & c^\mathsf{T} P_{L-1}[x_L] + c_0 \\ \text{s.t.} \ & (9a) - (9h), (12b), (12c), (14). \end{aligned} \tag{15}$$

Simple numerical examples in Figure 1 show that adding each of linear cuts (12b), (12c) and (14) shrinks the relaxation region of $(P_i[x_i x_i^\mathsf{T}], P_i[x_{i+1} x_{i+1}^\mathsf{T}], P_i[x_{i+1} x_i^\mathsf{T}])$ and thus tightens the layer SDP relaxation. It follows that the layer RLT-SDP relaxation (15) offers a tighter bound than the layer SDP relaxation (9), or formally:

Theorem 1. $\gamma^*_{\textit{GlobalSDP}} \le \gamma^*_{\textit{LayerSDP}} \le \gamma^*_{\textit{RLT-SDP}} \le \gamma^*$.

Algorithm 1: Implementation of layer RLT-SDP relaxation

Input: NN parameters, $\{p_s\}_{s=1}^r$, $k_{\max}$.
Initialise: Set the order to add the linear cuts using matrices $\mathcal{O}_i$, $i \in \mathbb{I}_{L-1}$. Set $k = 1$.
while $\gamma^*_{\text{RLT-SDP}} < 0$ and $k \leq k_{\max}$ **do**
Set $\bar{p}_i$ as the integer part of the product $p_k n_i$, $i \in \mathbb{I}_{L-1}$.
Solve (15), where for each P_i, $i \in \mathbb{I}_{L-1}$, adding only $\bar{p}_i$ elements (with corresponding indexes in $\mathcal{O}_i$) of each row in (12b) and (12c).
Set $k = k + 1$.
end while
Output: $\gamma^*_{\text{RLT-SDP}}$

Proof. Following the principle of RLT, the added linear cuts (12b), (12c) and (14) can always be deduced from the original linear constraints in the layer SDP relaxation (9). Hence, the optimal objective value of the layer RLT-SDP relaxation (15) still serves as a lower bound to that of the QCQP and the original verification problem (1). Moreover, adding these valid linear cuts can shrink the feasible region, as shown in Observations 1 and 4. This means that every solution to (15) is feasible to (9). Therefore, the layer RLT-SDP relaxation is at least as tight as the original layer SDP relaxation. □

Efficient implementation. The number of linear inequalities introduced by (12b), (12c) and (14) for each P_i, $i \in \mathbb{I}_{L-1}$, are $n_i n_{i+1}$, $2n_i n_{i+1}$ and $2n_{i+1}(n_{i+1}-1)$ (by removing diagonals), respectively. For P_0, $n_0(n_0-1)$ extra linear inequalities are needed. The total number of inequalities for each P_i, $i = 1, 2, \ldots, L-1$, is $(2n_{i+1}^2 + 3n_i n_{i+1} - 2n_{i+1})$, and for P_0 is $(2n_1^2 + 3n_0 n_1 - 2n_1 + n_0^2 - n_0)$. Compared to directly imposing the constraints in (10) (which introduces $2(n_i + n_{i+1})^2$ inequalities), adding (12b), (12c) and (14) has a lower computational burden, especially for large NNs. However, adding all the inequalities in (12b), (12c) and (14) is still computationally expensive. An efficient strategy for integrating them with the layer SDP relaxation is thus necessary. The strategy we deploy is based on two observations:

- The linear cuts (12b) and (12c) capture *inter-layer dependencies* (*i.e.*, terms $x_{i+1}x_i^\mathsf{T}$). Since $x_{i+1} = \text{ReLU}(W_i x_i + b_i)$, the dependencies are also reflected in the weighting matrix W_i. Hence, the structure of W_i can be exploited to efficiently adding (12b) and (12c).
- The linear cut (14) captures the *intra-layer interactions* (*i.e.*, terms $x_i x_i^\mathsf{T}$), which cannot be clearly indicated by NN parameters (weights or biases).

Given these observations, in the following we use the linear cuts (12b) and (12c). Moreover, it is straightforward to show that Theorem 1 holds even by adding a portion of (12b) and (12c). Hence, the computational cost of the problem can be reduced by adding only a portion of them.

Algorithm 1 describes an efficient implementation of the layer RLT-SDP relaxation. The fraction of linear cuts added at each iteration are set by choosing the sequence $\{p_s\}_{s=1}^r$, where $0 \leq p_1 < \cdots < p_r \leq 1$. In practice, the sequence $\{p_s\}_{s=1}^r$ and the maximum iteration $k_{\max}$, which satisfies $1 \leq k_{\max} \leq r$, can be adapted to the computational power available. Note that, in principle, a different sequence $\{p_s\}_{s=1}^r$ can be chosen for each individual layer; for simplicity we consider these to be constant. The matrix $\mathcal{O}_i$ stores the ordering (in descending order) of the elements in each row of $|W_i|$. The purpose of creating the ordering is to ensure that the part of linear cuts with larger influences on shrinking the feasible region are added first. This is based on the consideration as follows: For the neuron m at layer $i+1$, its pre-activation is $\hat{x}_{i+1,m} = W_i(m,:)x_i + b_i(m)$, where $W_i(m,:)$ is a row vector. Let w_1 and w_2 be any two elements of $W_i(m,:)$ and their corresponding inputs are $x_{i,1}$ and $x_{i,2}$, respectively. If $|w_1| > |w_2|$, then compared to those linear cuts about $x_{i+1}x_{i,2}$, the linear cuts about $x_{i+1}x_{i,1}$ has bigger influence on the feasible region of $x_{i+1,m}$. Figure 2 provides an example for this, where it is seen that the linear cuts about $x_{i+1}x_{i,2}$ contributes more than $x_{i+1}x_{i,1}$ in shrinking the feasible region of x_{i+1}.

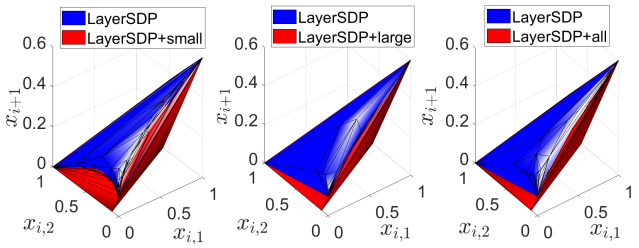

Figure 2: Feasible region of $x_{i+1} = \text{ReLU}(W_i x_i + 0.1)$, with $W_i = [0.5\ -1]$, $x_i = [x_{i,1},\ x_{i,2}]^\mathsf{T}$ and $0 \leq x_{i,1}, x_{i,2} \leq 1$, by adding a part or all of linear cuts (12b) and (12c). Adding only the linear cuts about $x_{i+1}x_{i,2}$ (*i.e.*, the larger element of $|W_i|$) yields a feasible region close to the one with full constraints on $x_{i+1}x_i^\mathsf{T}$.

We can show the following property of Algorithm 1:

Theorem 2. *The relation $\gamma^*_{\text{LayerSDP}} \leq \gamma^*_{\text{RLT-SDP}} \leq \gamma^*$ holds under any choice of $\{p_s\}_{s=1}^r$. At any given iteration k of Algorithm 1, we have that $\gamma^*_{\text{RLT-SDP}_k} \leq \gamma^*_{\text{RLT-SDP}_{k+1}} \leq \gamma^*$.*

Proof. It is straightforward to prove the first part following Theorem 1. At each iteration k, the proportion of linear cuts added is $\bar{p}_i$, the integer part of $p_k n_i$. Since $p_{k+1} > p_k$, the proportion added at iteration $k+1$ is larger than that at iteration k and contains it as a subset. Hence, for any $k \geq 1$, every feasible solution to the optimisation problem solved at iteration $k+1$ is also a solution to the problem solved at iteration k, *i.e.*, $\gamma^*_{\text{RLT-SDP}_k} \leq \gamma^*_{\text{RLT-SDP}_{k+1}}$. □

At each iteration, the layer RLT-SDP relaxation (15) is solved with a total number of $\sum_{i=1}^{L-1} 3\bar{p}_i n_{i+1}$ linear cuts. This is computationally lighter than the problem obtained by adding all the inequalities in (12b) and (12c). Furthermore, before running the algorithm, we can also remove the inactive neurons and simplify the constraints of stable neurons to reduce the sizes of the constraints $P_i \succeq 0$, $i \in \mathbb{I}_{L-1}$. This can be realised by examining the activation pattern of the NN under a given verification query and will not relax the solution. This is a strategy used in (Batten et al. 2021).

Experimental Evaluation

Two sets of experiments were carried out to evaluate the precision and scalability of relaxation proposed as well as Algorithm 1. The experiments were run on a Linux machine with an Intel i9-10920X 3.5 GHz 12-core CPU with 128 GB RAM. The optimisation problems were modelled by using YALMIP (Lofberg 2004) and solved using MOSEK (Andersen and Andersen 2000). We compared the results obtained against presently available SoA methods and tools.

Networks. In Experiment 1, we considered two groups of two-input, two-output, fully-connected random ReLU NNs generated by using the method in (Fazlyab, Morari, and Pappas 2022). *Group 1* had four models with $L = 4, 6, 8, 10$ hidden layers, respectively, and 15 neurons for each hidden layer. *Group 2* had four three-layer models, with $n_i = 10, 15, 50, 100$ neurons per hidden layer, respectively.

In Experiment 2, we considered three groups of fully-connected ReLU NNs trained on the MNIST dataset. These are widely used in all the SDP benchmarks (By "$m \times n$" we mean a NN with $m-1$ hidden layers each having n neurons, which is consistent with (Batten et al. 2021).):

- *Small NNs*: MLP-Adv, MLP-LP and MLP-SDP from (Raghunathan, Steinhardt, and Liang 2018) and tested under the same perturbation $\epsilon = 0.1$ as in (Raghunathan, Steinhardt, and Liang 2018; Batten et al. 2021).
- *Medium NNs*: Models 6×100 and 9×100 from (Singh et al. 2019a) and evaluated under the same $\epsilon = 0.026$ and $\epsilon = 0.015$ as in (Singh et al. 2019a; Tjandraatmadja et al. 2020; Müller et al. 2021; Batten et al. 2021)
- *Large NNs*: Models 8×1024-0.1 and 8×1024-0.3 from (Li et al. 2020), which were trained using CROWN-IBP (Zhang et al. 2019) with adversarial attack $\epsilon = 0.1, 0.3$, respectively. As in (Li et al. 2020), they were tested under the perturbations $\epsilon = 0.1, 0.3$, respectively.

Baseline methods. In Experiment 2, we compared the proposed layer RLT-SDP relaxation (referred to as `RLT-SDP`) against the SoA methods for verification below:

- *Complete methods*: `MILP` (Tjeng, Xiao, and Tedrake 2019), `AI`^{2} (Gehr et al. 2018), and β-`CROWN` (Wang et al. 2021).
- *Linear relaxations*: the standard linear programming relaxation `LP` (Ehlers 2017) and its variants including `IBP` (Gowal et al. 2019), `OptC2V` (Tjandraatmadja et al. 2020), and `PRIMA` (Müller et al. 2021). We did not consider `kPoly` (Singh et al. 2019b) and `DeepPoly` (Singh et al. 2019a) , as they were shown in (Müller et al. 2021) to be weaker than `PRIMA`.
- *SDP relaxations*: `LayerSDP` (Batten et al. 2021), `SDP-IP` (*i.e.*, the global SDP relaxation (7)) (Raghunathan, Steinhardt, and Liang 2018), and `SDP-FO` (Dathathri et al. 2020).

Experiment 1: Efficacy of the proposed strategy. We investigated both the network depth and width by using `RLT-SDP` to obtain an over-approximation of the feasible output region of the NN for a given input set. The test inputs were random values within $[0, 1]$ and the heuristic method

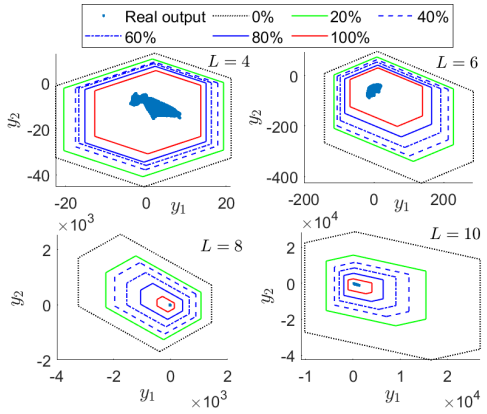

Figure 3: Over-approximated output region by `RLT-SDP` with different percentages of linear cuts for networks of different hidden layers L. The 0% case is `LayerSDP`.

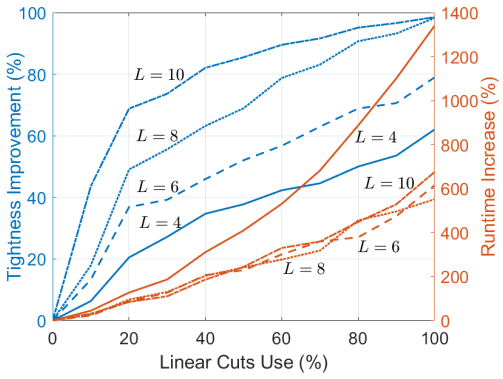

Figure 4: Tightness improvement and runtime increase obtained by solving `RLT-SDP` with different percentages of linear cuts for networks of different hidden layers L. The 0% case is `LayerSDP`.

in (Fazlyab, Morari, and Pappas 2022) was adopted to compute the over-approximations. Algorithm 1 was run with $\{p_s\}_{s=1}^{11} = \{0, 0.1, \cdots, 1\}$ and $k_{\max} = 11$. Without linear cuts ($p_1 = 0$), `RLT-SDP` is equivalent to `LayerSDP`.

We first studied the impact of network depth on the verification method here proposed by using the models in *Group 1*. Figure 3 shows that for all the four models considered, adding a larger percentage of linear cuts yields a tighter over-approximation. As the number of hidden layers L increases, `LayerSDP` becomes looser and the effects of adding linear cuts becomes more significant. The figures show that across all models, even using just 20% of the linear cuts considerably reduces the over-approximation. To further analyse the gain in the approximation versus the corresponding increase in computational complexity, we considered two metrics: the improvement in approximation (or tightness) and the runtime increase. The former is the relative reduction in the feasible output regions obtained by `RLT-SDP` and `LayerSDP`; the latter is the relative increase in their runtime. As expected, it is shown in Figure 4 that adding a larger

		RLT−SDP		LayerSDP		SDP−IP		SDP−FO	LP	OptC2V	PRIMA	β−CROWN	MILP	AI^2	IBP
Models	PGD	Certified	Time*	Cert.†	Time*	Cert.†	Time*	Cert.†	Cert.†	Cert.†	Cert.†	Cert.†	Cert.†	Cert.†	Cert.†
MLP-Adv	94%	**88** | **94**	3622	80 | 91	2164	82	12079	84	65	–	–	–	–	–	–
MLP-LP	80%	**80** | **80**	159	80 | 80	145	80	50733	78	79	–	–	–	–	–	–
MLP-SDP	84%	**84** | **84**	11141	80 | 84	4373	80	43156	64	35	–	–	–	–	–	–
6×100	91%	**82** | **90**	3297	60 | 75	1900	–	6760	◇	21	42.9	51.0	69.9	–	–	–
9×100	86%	**56** | **70**	10800	22 | 35	7119	–	11899	◇	18	38.4	42.8	62.0	–	–	–
8×1024-0.1	89%	**82** |	5883	–	2932	◇	◇	◇	0	–	–	–	67	52	80
8×1024-0.3	26%	**26** |	761	–	469	◇	◇	◇	0	–	–	–	7	16	22

Table 1: Certified robustness (in percentage) and runtime per image (in seconds) for a set of benchmarks with various sizes. Dagger (†): these results are directly taken from the literature: LayerSDP and SDP−FO from (Batten et al. 2021), SDP−IP from (Raghunathan, Steinhardt, and Liang 2018), OptC2V from (Tjandraatmadja et al. 2020), PRIMA from (Müller et al. 2021), β−CROWN from (Wang et al. 2021), while MILP, AI^2 and IBP from (Li et al. 2020); The first four numbers of LP are from (Batten et al. 2021), and the last two are obtained by implementation with interval arithmetic bounds. Dash (–): previously reported numbers are unavailable. Diamond (◇): the methods fail to verify any instance. Star (∗): the runtime is estimated by running over five images using the same interval arithmetic bounds. Vertical line (|): the certified robustness on the left and right are obtained using interval arithmetic bounds and symbolic interval propagation, respectively.

proportion of linear cuts yields a tighter over-approximation, along with an increase in runtime. Adding the same percentage of linear cuts leads to a more significant tightness improvement on larger networks (with larger L) than on smaller ones. For each network, as the percentage of linear cuts increases, the tightness improvement becomes less significant, but the runtime increase becomes more significant. Particularly, experimentally we found that the first 20% of linear cuts contributes most significantly to the improvement in overall tightness of the method. We evaluated the impact of network width by using the models in *Group 2* and observed very similar behaviour of the method.

These results clearly confirm Theorem 2 and demonstrate the efficiency of Algorithm 1. They also indicate that a trade-off needs to be balanced between the tightness improvement and runtime increase. Specifically, the addition of 20% of linear cuts could be sufficient to improve considerably the precision of the SDP approach without incurring the higher computational costs associated with larger problems.

Experiment 2: Comparison with the SoA methods. We benchmarked the technique on the NNs built on the MNIST dataset described above. All experiments were run on the first 100 images of the dataset. The results obtained are reported in Table 1, where the runtime is the solver time. The PGD upper bounds of MLP-Adv, MLP-LP, MLP-SDP, 6×100 and 9×100 are taken from (Batten et al. 2021), while those of 8×1024-0.1 and 8×1024-0.3 are from (Li et al. 2020). We ran Algorithm 1 with the sequence $\{0.1, 0.2\}$ and $k_{\max} = 2$. As in LayerSDP, we further optimised RLT−SDP by removing inactive neurons in the first step.

Our results show that RLT−SDP based on the interval arithmetic bounds is more precise than LayerSDP under the same bounds and all other baseline methods for all the networks. One exception is the 9×100 network, for which β−CROWN achieves the highest precision. By using the tighter symbolic bound propagation (Botoeva et al. 2020), RLT−SDP significantly outperformed all the incomplete/complete baseline methods.

As expected we found RLT−SDP to be significantly more computationally demanding than LayerSDP across all the networks. However, it was still faster than SDP−IP for the small and median networks. Neither SDP−IP nor SDP−FO could verify the two large networks. Also, it is shown in (Batten et al. 2021) that compared to LayerSDP, SDP−FO has a runtime that is much larger for MLP-Adv and MLP-LP, but smaller for MLP-SDP. SDP−FO fails to verify 6×100 and 9×100. These results confirm that RLT−SDP remains competitive in terms of computational efficiency. We note that the runtime of LayerSDP in Table 1 is larger than that reported in (Batten et al. 2021). This is because we directly solved the layer SDP relaxation (9), without implementing SparseColO (Fujisawa and et al. 2009) or the automatic model transformation as in (Batten et al. 2021). Their work shows that using subroutine from SparseColO can balance the size of semidefinite constraints and equality constraints, and using automatic model transformation can reduce YALMIP overhead time, both of which significantly improve the efficiency of LayerSDP. Note, however, that these techniques are also directly applicable to RLT−SDP. Hence, the results presented here provides a like-for-like comparison between LayerSDP and RLT−SDP.

Conclusions

While SDP-based algorithms have shown their promise as a next generation method to verify NN, their resulting over-approximations are still too coarse to verify deep and large NNs. In this paper we put forward a novel SDP relaxation for achieving tight and efficient neural network robustness verification. We did so by combining the layerwise SDP relaxation with RLT. We showed that the method always yields tighter bounds than the present SoA. We also illustrated how a careful choice of linear cuts can mitigate the additional computational cost, thereby resulting in an overall tight and computationally balanced technique. The experiments reported demonstrated that the method achieves SoA on all benchmarks commonly used in the area.

Acknowledgements

This work is partly funded by DARPA under the Assured Autonomy programme (FA8750-18-C-0095). Alessio Lomuscio is supported by a Royal Academy of Engineering Chair in Emerging Technologies.

References

Akintunde, M. E.; Botoeva, E.; Kouvaros, P.; and Lomuscio, A. 2020a. Formal Verification of Neural Agents in Non-deterministic Environments. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS20)*, 25–33.

Akintunde, M. E.; Botoeva, E.; Kouvaros, P.; and Lomuscio, A. 2020b. Verifying Strategic Abilities of Neural-symbolic Multi-agent Systems. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR20)*, volume 17, 22–32.

Andersen, E. D.; and Andersen, K. D. 2000. The MOSEK interior point optimizer for linear programming: an implementation of the homogeneous algorithm. In *High performance optimization*, 197–232. Springer.

Anderson, R.; Huchette, J.; Ma, W.; Tjandraatmadja, C.; and Vielma, J. 2020. Strong mixed-integer programming formulations for trained neural networks. *Mathematical Programming*, 1–37.

Anstreicher, K. M. 2009. Semidefinite programming versus the reformulation-linearization technique for nonconvex quadratically constrained quadratic programming. *Journal of Global Optimization*, 43(2-3): 471–484.

Bak, S.; Liu, C.; and Johnson, T. 2021. The Second International Verification of Neural Networks Competition (VNN-COMP 2021): Summary and Results. *arXiv preprint arXiv:2103.06624*.

Bastani, O.; Ioannou, Y.; Lampropoulos, L.; Vytiniotis, D.; Nori, A.; and Criminisi, A. 2016. Measuring neural net robustness with constraints. In *Advances in Neural Information Processing Systems (NeurIPS16)*, 2613–2621.

Batten, B.; Kouvaros, P.; Lomuscio, A.; and Zheng, Y. 2021. Efficient Neural Network Verification via Layer-based Semidefinite Relaxations and Linear Cuts. In *International Joint Conference on Artificial Intelligence (IJCAI21)*, 2184–2190.

Botoeva, E.; Kouvaros, P.; Kronqvist, J.; Lomuscio, A.; and Misener, R. 2020. Efficient verification of neural networks via dependency analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI20)*, 3291–3299.

Dathathri, S.; Dvijotham, K.; Kurakin, A.; Raghunathan, A.; Uesato, J.; Bunel, R.; Shankar, S.; Steinhardt, J.; Goodfellow, I.; Liang, P.; and Pushmeet, K. 2020. Enabling certification of verification-agnostic networks via memory-efficient semidefinite programming. In *Advances in Neural Information Processing Systems (NeurIPS20)*, 1–14.

Dvijotham, K.; Stanforth, R.; Gowal, S.; Mann, T.; and Kohli, P. 2018. A dual approach to scalable verification of deep networks. *arXiv preprint arXiv:1803.06567*.

Ehlers, R. 2017. Formal verification of piece-wise linear feed-forward neural networks. In *International Symposium on Automated Technology for Verification and Analysis (ATVA17)*, 269–286.

Fazlyab, M.; Morari, M.; and Pappas, G. J. 2022. Safety verification and robustness analysis of neural networks via quadratic constraints and semidefinite programming. *IEEE Transactions on Automatatic Control*, 67(1): 1–15.

Fujisawa, K.; and et al. 2009. User's manual for SparseCoLO: Conversion methods for sparse conic-form linear optimization problems. *Dept. of Math. and Comp. Sci. Japan, Tech. Rep.*, 152–8552.

Gehr, T.; Mirman, M.; Drachsler-Cohen, D.; Tsankov, P.; Chaudhuri, S.; and Vechev, M. 2018. AI^2: Safety and robustness certification of neural networks with abstract interpretation. In *IEEE Symposium on Security and Privacy (SP18)*, 3–18. IEEE.

Goodfellow, I.; Shlens, J.; and Szegedy, C. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*.

Gowal, S.; Dvijotham, K. D.; Stanforth, R.; Bunel, R.; Qin, C.; Uesato, J.; Arandjelovic, R.; Mann, T.; and Kohli, P. 2019. Scalable verified training for provably robust image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (IEEE/CVF19)*, 4842–4851.

Hashemi, V.; Kouvaros, P.; and Lomuscio, A. 2021. OSIP: Tightened Bound Propagation for the Verification of ReLU Neural Networks. In *International Conference on Software Engineering and Formal Methods (SEFM21)*, 463–480. Springer.

Henriksen, P.; and Lomuscio, A. 2020. Efficient neural network verification via adaptive refinement and adversarial search. In *Proceedings of the Twenty-fourth European Conference on Artificial Intelligence (ECAI20)*, 2513–2520. IOS Press.

Henriksen, P.; and Lomuscio, A. 2021. DEEPSPLIT: An Efficient Splitting Method for Neural Network Verification via Indirect Effect Analysis. In *International Joint Conference on Artificial Intelligence (IJCAI21)*, 2549–2555.

Julian, K. D.; and Kochenderfer, M. J. 2021. Reachability Analysis for Neural Network Aircraft Collision Avoidance Systems. *Journal of Guidance, Control, and Dynamics*, 44(6): 1132–1142.

Katz, G.; Huang, D.; Ibeling, D.; Julian, K.; Lazarus, C.; Lim, R.; Shah, P.; Thakoor, S.; Wu, H.; Zeljic, A.; Dill, D.; Kochenderfer, M.; and Barrett, C. 2019. The Marabou framework for verification and analysis of deep neural networks. In *International Conference on Computer Aided Verification (CAV19)*, 443–452.

Kouvaros, P.; Kyono, T.; Leofante, F.; Lomuscio, A.; Margineantu, D.; Osipychev, D.; and Zheng, Y. 2021. Formal Analysis of Neural Network-Based Systems in the Aircraft Domain. In *International Symposium on Formal Methods (FM21)*, 730–740. Springer.

Lasserre, J. B. 2009. *Moments, positive polynomials and their applications*, volume 1. World Scientific.

Li, L.; Qi, X.; Xie, T.; and Li, B. 2020. SoK: Certified robustness for deep neural networks. *arXiv preprint arXiv:2009.04131*.

Liu, C.; Arnon, T.; Lazarus, C.; Strong, C.; Barrett, C.; Kochenderfer, M. J.; et al. 2020. Algorithms for Verifying Deep Neural Networks. *Foundations and Trends® in Optimization*, 3-4: 244–404.

Lofberg, J. 2004. YALMIP: A toolbox for modeling and optimization in MATLAB. In *IEEE International Conference on Robotics and Automation (ICRA04)*, 284–289. IEEE.

Lomuscio, A.; and Maganti, L. 2017. An approach to reachability analysis for feed-forward ReLU neural networks. *CoRR*, abs/1706.07351.

Ma, Z.; and Sojoudi, S. 2020. Strengthened SDP verification of neural network robustness via non-convex cuts. *arXiv preprint arXiv:2010.08603*.

Manzanas Lopez, D.; Johnson, T.; Tran, H.-D.; Bak, S.; Chen, X.; and Hobbs, K. L. 2021. Verification of Neural Network Compression of ACAS Xu Lookup Tables with Star Set Reachability. In *AIAA Scitech 2021 Forum*, 0995.

Müller, M. N.; Makarchuk, G.; Singh, G.; Püschel, M.; and Vechev, M. 2021. PRIMA: Precise and general neural network certification via multi-neuron convex relaxations. *arXiv preprint arXiv:2103.03638*.

Parrilo, P. A. 2000. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. California Institute of Technology.

Raghunathan, A.; Steinhardt, J.; and Liang, P. 2018. Semidefinite relaxations for certifying robustness to adversarial examples. In *Advances in Neural Information Processing Systems (NeurIPS18)*, 10877–10887.

Salman, H.; Yang, G.; Zhang, H.; Hsieh, C.; and Zhang, P. 2019. A convex relaxation barrier to tight robustness verification of neural networks. In *Advances in Neural Information Processing Systems (NeurIPS19)*, 9835–9846.

Singh, G.; Gehr, T.; Püschel, M.; and Vechev, M. 2019a. An abstract domain for certifying neural networks. *Proceedings of the ACM on Programming Languages*, 3(POPL): 41:1–41:30.

Singh, G.; R. Ganvir, R.; Püschel, M.; and Vechev, M. 2019b. Beyond the single neuron convex barrier for neural network certification. In *Advances in Neural Information Processing Systems (NeurIPS19)*, 15098–15109.

Tjandraatmadja, C.; Anderson, R.; Huchette, J.; Ma, W.; Patel, K.; and Vielma, J. P. 2020. The convex relaxation barrier, revisited: Tightened single-neuron relaxations for neural network verification. In *Advances in Neural Information Processing Systems (NeurIPS20)*, 1–12.

Tjeng, V.; Xiao, K.; and Tedrake, R. 2019. Evaluating robustness of neural networks with mixed integer programming. In *International Conference on Learning Representations (ICLR19)*, 1–21.

Vandenberghe, L.; and Andersen, M. S. 2015. Chordal graphs and semidefinite optimization. *Foundations and Trends in Optimization*, 1(4): 241–433.

Vandenberghe, L.; and Boyd, S. 1996. Semidefinite programming. *SIAM Review*, 38(1): 49–95.

Wang, S.; Pei, K.; Whitehouse, J.; Yang, J.; and Jana, S. 2018. Efficient formal safety analysis of neural networks. In *Advances in Neural Information Processing Systems (NeurIPS18)*, 6367–6377.

Wang, S.; Zhang, H.; Xu, K.; Lin, X.; Jana, S.; Hsieh, C.-J.; and Kolter, J. Z. 2021. Beta-crown: Efficient bound propagation with per-neuron split constraints for complete and incomplete neural network verification. *arXiv preprint arXiv:2103.06624*.

Weng, T.; Zhang, H.; Chen, H.; Song, Z.; Hsieh, C.; Boning, D.; Dhillon, I.; and Daniel, L. 2018. Towards fast computation of certified robustness for ReLU networks. In *International Conference on Machine Learning (ICML18)*, 5276–5285.

Wong, E.; and Kolter, Z. 2018. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *International Conference on Machine Learning (ICML18)*, 5286–5295.

Zhang, H.; Chen, H.; Xiao, C.; Gowal, S.; Stanforth, R.; Li, B.; Boning, D.; and Hsieh, C.-J. 2019. Towards stable and efficient training of verifiably robust neural networks. *arXiv preprint arXiv:1906.06316*.

Zhang, R. Y. 2020. On the tightness of semidefinite relaxations for certifying robustness to adversarial examples. *arXiv preprint arXiv:2006.06759*.

Zheng, Y. 2019. *Chordal sparsity in control and optimization of large-scale systems*. Ph.D. thesis, University of Oxford.

Zheng, Y.; Fantuzzi, G.; and Papachristodoulou, A. 2021. Chordal and factor-width decompositions for scalable semidefinite and polynomial optimization. *Annual Reviews in Control*, 52: 243–279.