

---

# A Convex Relaxation Barrier to Tight Robustness Verification of Neural Networks

---

**Hadi Salman\***  
Microsoft Research AI  
hadi.salman@microsoft.com

**Greg Yang**  
Microsoft Research AI  
gregyang@microsoft.com

**Huan Zhang**  
UCLA  
huan@huan-zhang.com

**Cho-Jui Hsieh**  
UCLA  
chohsieh@cs.ucla.edu

**Pengchuan Zhang**  
Microsoft Research AI  
penzhan@microsoft.com

## Abstract

Verification of neural networks enables us to gauge their robustness against adversarial attacks. Verification algorithms fall into two categories: *exact* verifiers that run in exponential time and *relaxed* verifiers that are efficient but incomplete. In this paper, we unify all existing LP-relaxed verifiers, to the best of our knowledge, under a general convex relaxation framework. This framework works for neural networks with diverse architectures and nonlinearities and covers both primal and dual views of neural network verification. We further prove strong duality between the convex-relaxed primal and associated dual problems under very mild conditions. Next, we perform large-scale experiments, amounting to more than 22 CPU-years, to obtain exact solution to the convex-relaxed problem that is optimal within our framework for ReLU networks. We find the exact solution does not significantly improve upon the gap between PGD and existing relaxed verifiers for various networks trained normally or robustly on MNIST and CIFAR datasets. Our results suggest there is an inherent *barrier* to tight verification for the large class of methods captured by our framework. We discuss possible causes of this barrier and potential future directions for bypassing it. Our code can be found at <https://github.com/Hadisalman/robust-verify-benchmark>.

## 1 Introduction

Neural networks (NNs) have had a lot of successes in tasks such as image classification (Krizhevsky et al., 2012; He et al., 2016) and speech recognition (Hinton et al., 2012). However, it has been shown that NNs are vulnerable to *adversarial examples*, for instance perturbed images which are visually indistinguishable from the original to a human, but can change the NN’s classification dramatically (Szegedy et al., 2013; Goodfellow et al., 2015). This becomes a serious problem when NNs are applied to safety-critical applications, or those where security is a requirement. Such concern has prompted very active research in adversarial *attack* (designing adversarial examples) (Goodfellow et al., 2015; Carlini & Wagner, 2017; Chen et al., 2017), *defense* (train- or inference-time procedures to nullify the effectiveness of adversarial examples) (Madry et al., 2017; Wong & Kolter, 2018; Raghunathan et al., 2018a), and *verification* (certifying that no small perturbations of a given input can cause the NN to change its prediction) (Katz et al., 2017; Dvijotham et al., 2018b; Weng et al., 2018). This paper is concerned with the last, *robustness verification*.

Many recent works have proposed robustness verification methods. Typically, such algorithms try to exactly find or upper bound an adversarial loss (or *robust error*), defined to be the worst case

---

\*Work done as part of the Microsoft AI Residency Program.

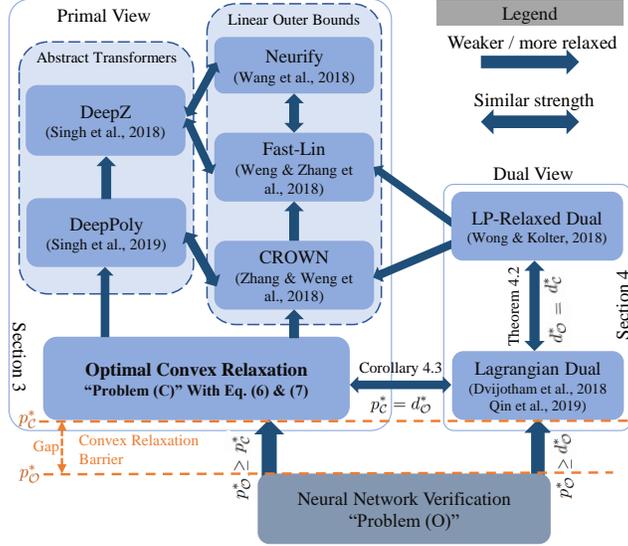


Figure 1: Relationship between existing relaxed algorithms and our framework. See main text.

(maximum) loss (or error) over all possible norm-bounded perturbations of the input. Alternatively, they exactly find or lower bound the minimum adversarial distortion  $\epsilon$ , defined to be the radius of the largest ball (in a specified norm), around the input, that contains no adversarial examples. The algorithms devised so far can be roughly categorized as exact (Ehlers, 2017; Katz et al., 2017; Tjeng et al., 2019) or relaxed methods (Wong & Kolter, 2018; Dvijotham et al., 2018d; Zhang et al., 2018; Gehr et al., 2018; Singh et al., 2018; Weng et al., 2018; Wang et al., 2018b), and among the latter, the most dominant category is *convex relaxation*, which tries to relax the exact adversarial optimization problem into a convex problem. Our **first main contribution** is to give a unifying treatment of convex relaxations used in a large number of recent works (Fig. 1) and reveal the relationships between them. We further prove that these relaxed problems satisfy strong duality under very mild conditions, connecting relaxed verifiers in the primal space (Zhang et al., 2018; Weng et al., 2018) to those in the dual space (Dvijotham et al., 2018d; Wong & Kolter, 2018).

For any NN, our framework prescribes a unique tightest layer-wise relaxation, whose optimal value  $p_C^*$  presents a *barrier* for any method described by the framework. How close is  $p_C^*$  to the true adversarial loss? In our **second main contribution**, we perform extensive experiments with deep ReLU networks to compute  $p_C^*$  and compare with the LP-relaxed dual formulation of Wong & Kolter (2018) (which, as our framework will show, approximates  $p_C^*$ ), the PGD attack of Madry et al. (2017), and the mixed integer linear programming (MILP) exact verifier of Tjeng et al. (2019). Over different models, sizes, training methods, and datasets (MNIST and CIFAR-10), we find that our tightest layer-wise relaxation only very slightly improves the lower bound of the minimum  $l_\infty$  adversarial distortion found by Wong & Kolter (2018), especially when compared with the upper bound provided by the PGD attack, which is consistently 1.5 to 5 times larger. For networks trained using the LP-relaxed dual formulation of Wong & Kolter (2018), when  $\epsilon$  is large and the robust error is also large, our tightest relaxation improves the certified robust error significantly in relative terms, but at most 7% in absolute terms. Furthermore, on both normally and adversarially-trained MNIST classifiers, our tightest relaxation does not significantly close the gap in the *robust error* between the PGD lower bound or MILP exact answer on one hand, and the upper bound of Wong & Kolter (2018) on the other. Thus the convex relaxation barrier we discover theoretically indeed presents a significant challenge to existing algorithms empirically.

## 2 Preliminaries and Related Work

**Robustness to adversarial attacks.** Consider a multi-class classification network  $f(x) \in \mathbb{R}^K$ , where  $f_i(x) \in \mathbb{R}$  is the logit for the  $i$ -th class. Given a nominal input  $x^{\text{nom}}$ , its nominal prediction  $i^{\text{nom}} = \arg \max_i f_i(x^{\text{nom}})$ , and a neighborhood  $\mathcal{S}_{i_n}(x^{\text{nom}})$ , the following optimization problem

determines if the network is robust to adversarial attacks within this neighborhood:

$$\min_{x \in \mathcal{S}_{in}(x^{\text{nom}})} f_{i^{\text{nom}}}(x) - f_i(x). \quad (1)$$

If the minimal value is positive for all  $i \in \{1 \dots K\} \setminus \{i^{\text{nom}}\}$ , the network is robust at this point within  $\mathcal{S}_{in}(x^{\text{nom}})$ . We refer the reader to [Dvijotham et al. \(2018d\)](#), [Dvijotham et al. \(2018a\)](#) and [Qin et al. \(2019\)](#) for other robustness verification problems which can be solved with similar methods.

**Exact verifiers and NP-completeness.** For ReLU networks (piece-wise linear networks in general), exact verifiers solve the robustness verification problem (1) by typically employing MILP solvers ([Cheng et al., 2017](#); [Lomuscio & Maganti, 2017](#); [Dutta et al., 2018](#); [Fischetti & Jo, 2017](#); [Tjeng et al., 2019](#); [Xiao et al., 2019](#)) or Satisfiability Modulo Theories (SMT) solvers ([Scheibler et al., 2015](#); [Katz et al., 2017](#); [Carlini et al., 2017](#); [Ehlers, 2017](#)). However, due to the NP-completeness for solving such a problem ([Katz et al., 2017](#); [Weng et al., 2018](#)), it can be really challenging to scale these to large networks. It can take Reluplex ([Katz et al., 2017](#)) several hours to find the minimum distortion of an example for a ReLU network with 5 inputs, 5 outputs, and 300 neurons. A recent work by [Tjeng et al. \(2019\)](#) uses MILP to exactly verify medium-size networks, but the verification time is very sensitive to how a network is trained; for example, it is fast for networks trained using the LP-relaxed dual formulation of [Wong & Kolter \(2018\)](#), but much slower for normally trained networks. A concurrent work by [Xiao et al. \(2019\)](#) trains networks with the objective of speeding up the MILP verification problem, but this compromises on the performance of the network.

**Relaxed and efficient verifiers.** These verifiers solve a relaxed, but more computationally efficient, version of (1), and have been proposed from different perspectives. From the primal view, one can relax the nonlinearity in (1) into linear inequality constraints. This perspective has been previously explored as in the framework of “abstract transformers” ([Singh et al., 2018, 2019b,a](#); [Gehr et al., 2018](#); [Mirman et al., 2018](#)), via linear outer bounds of activation functions ([Zhang et al., 2018](#); [Weng et al., 2018](#); [Wang et al., 2018a,b](#)), or via interval bound propagation ([Gowal et al., 2018](#); [Mirman et al., 2018](#)). From the dual view, one can study the dual of the relaxed problem ([Wong & Kolter, 2018](#); [Wong et al., 2018](#)) or study the dual of the original nonconvex verification problem ([Dvijotham et al., 2018d,c](#); [Qin et al., 2019](#)). In this paper, we unify both views in a common convex relaxation framework for NN verification, clarifying their relationships (as summarized in Fig. 1).

[Raghunathan et al. \(2018b\)](#) formulates the verification of ReLU networks as a quadratic programming problem and then relaxes and solves this problem with a semidefinite programming (SDP) solver. While our framework does not cover this SDP relaxation, it is not clear to us how to extend the SDP relaxed verifier to general nonlinearities, for example max-pooling, which can be done in our framework on the other hand. Other verifiers have been proposed to certify via an intermediary step of bounding the local Lipschitz constant ([Hein & Andriushchenko, 2017](#); [Weng et al., 2018](#); [Raghunathan et al., 2018a](#); [Zhang et al., 2019](#)). These are outside the scope of our framework.

Combining exact and relaxed verifiers, hybrid methods have shown some effectiveness ([Bunel et al., 2018](#); [Singh et al., 2019a](#)). In fact, many exact verifiers also use relaxation as a subroutine to speed things up, and hence can be viewed as hybrid methods as well. In this paper, we are not concerned with such techniques but only focus on *relaxed verifiers*.

### 3 Convex Relaxation from the Primal View

**Problem setting.** In this paper, we assume that the neighborhood  $\mathcal{S}_{in}(x^{\text{nom}})$  is a convex set. An example of this is  $\mathcal{S}_{in}(x^{\text{nom}}) = \{x : \|x - x^{\text{nom}}\|_{\infty} \leq \epsilon\}$ , which is the constraint on  $x$  in the  $\ell_{\infty}$  adversarial attack model. We also assume that  $f(x)$  is an  $L$ -layer feedforward NN. For notational simplicity, we denote  $\{0, 1, \dots, L - 1\}$  by  $[L]$  and  $\{x^{(0)}, x^{(1)}, \dots, x^{(L-1)}\}$  by  $x^{[L]}$ . We define  $f(x)$  as,

$$\begin{aligned} x^{(l+1)} &= \sigma^{(l)}(\mathbf{W}^{(l)}x^{(l)} + b^{(l)}) \quad l \in [L], \\ f(x) &:= z^{(L)} = \mathbf{W}^{(L)}x^{(L)} + b^{(L)}, \end{aligned} \quad (2)$$

where  $x^{(l)} \in \mathbb{R}^{n^{(l)}}$ ,  $z^{(l)} \in \mathbb{R}^{n_z^{(l)}}$ ,  $x^{(0)} := x \in \mathbb{R}^{n^{(0)}}$  is the input,  $\mathbf{W}^{(l)} \in \mathbb{R}^{n_z^{(l)} \times n^{(l)}}$  and  $b^{(l)} \in \mathbb{R}^{n_z^{(l)}}$  are the weight matrix and bias vector of the  $l^{\text{th}}$  linear layer, and  $\sigma^{(l)} : \mathbb{R}^{n_z^{(l)}} \rightarrow \mathbb{R}^{n^{(l+1)}}$  is a (nonlinear) activation function like (leaky-)ReLU, the sigmoid family (including sigmoid, arctan, hyperbolic

tangent, etc), and the pooling family (MaxPool, AvgPool, etc). Our results can be easily extended to networks with convolutional layers and skip connections as well, similar to what is done in [Wong et al. \(2018\)](#), as these can be seen as special forms of (2).

Consider the following optimization problem  $\mathcal{O}(c, c_0, L, \underline{z}^{[L]}, \bar{z}^{[L]})$ :

$$\begin{aligned} \min_{(x^{[L+1]}, z^{[L]}) \in \mathcal{D}} \quad & c^\top x^{(L)} + c_0 \\ \text{s.t.} \quad & z^{(l)} = \mathbf{W}^{(l)} x^{(l)} + b^{(l)}, l \in [L], \\ & x^{(l+1)} = \sigma^{(l)}(z^{(l)}), l \in [L], \end{aligned} \quad (\mathcal{O})$$

where the optimization domain  $\mathcal{D}$  is

$$\mathcal{D} = \{(x^{[L+1]}, z^{[L]}) : x^{(0)} \in \mathcal{S}_{in}(x^{\text{nom}}), \underline{z}^{(l)} \leq z^{(l)} \leq \bar{z}^{(l)}, l \in [L]\}. \quad (3)$$

If  $c^\top = \mathbf{W}_{i^{\text{nom}}}^{(L)} \cdot -\mathbf{W}_{i^{\text{nom}}}^{(L)}$ ,  $c_0 = b_{i^{\text{nom}}}^{(L)} - b_{i^{\text{nom}}}^{(L)}$ ,  $\underline{z}^{[L]} = -\infty$ , and  $\bar{z}^{[L]} = \infty$ , then  $(\mathcal{O})$  is equivalent to problem (1). However, when we have better information about valid bounds  $\underline{z}^{[l]}$  and  $\bar{z}^{[l]}$  of  $z^{[l]}$ , we can significantly narrow down the optimization domain and, as will be detailed shortly, achieve tighter solutions when we relax the nonlinearities. We denote the minimal value of  $\mathcal{O}(c, c_0, L, \underline{z}^{[L]}, \bar{z}^{[L]})$  by  $p^*(c, c_0, L, \underline{z}^{[L]}, \bar{z}^{[L]})$ , or just  $p_{\mathcal{O}}^*$  when no confusion arises.

**Obtaining lower and upper bounds  $(\underline{z}^{[L]}, \bar{z}^{[L]})$  by solving sub-problems.** This can be done by *recursively* solving  $(\mathcal{O})$  with specific choices of  $c$  and  $c_0$ , which is a common technique used in many works ([Wong & Kolter, 2018](#); [Dvijotham et al., 2018d](#)). For example, one can obtain  $\underline{z}_j^{(\ell)}$ , a lower bound of  $z_j^{(\ell)}$ , by solving  $\mathcal{O}(\mathbf{W}_{j^{\text{nom}}}^{(\ell)\top}, b_{j^{\text{nom}}}^{(\ell)}, \ell, \underline{z}^{[\ell]}, \bar{z}^{[\ell]})$ ; this shows that one can estimate  $\underline{z}^{(l)}$  and  $\bar{z}^{(l)}$  inductively in  $l$ . However, we may have millions of sub-problems to solve because practical networks can have millions of neurons. Therefore, it is crucial to have *efficient* algorithms to solve problem  $(\mathcal{O})$ .

**General nonlinear specifications  $F(x^{(0)}, x^{(L)})$**  are also of interest beyond the linear specification  $c^\top x^{(L)} + c_0$  considered in  $(\mathcal{O})$ . These can be formulated as a special case of  $(\mathcal{O})$  by adding a nonlinear layer  $x^{(L+1)} = F(x^{(0)}, x^{(L)})$  on top of the original network, and setting  $c = 1$  and  $c_0 = 0$ . Because our framework can deal with general nonlinear layers, all the theoretical results in this paper can be applied to verification problems with nonlinear specifications as well.

**Convex relaxation in the primal space.** Due to the nonlinear activation functions  $\sigma^{(l)}$ , the feasible set of  $(\mathcal{O})$  is nonconvex, which leads to the NP-completeness of the neural network verification problem ([Katz et al., 2017](#); [Weng et al., 2018](#)). One natural idea is to do convex relaxation of its feasible set. Specifically, one can relax the nonconvex equality constraint  $x^{(l+1)} = \sigma^{(l)}(z^{(l)})$  to convex inequality constraints, i.e.,

$$\begin{aligned} \min_{(x^{[L+1]}, z^{[L]}) \in \mathcal{D}} \quad & c^\top x^{(L)} + c_0 \\ \text{s.t.} \quad & z^{(l)} = \mathbf{W}^{(l)} x^{(l)} + b^{(l)}, \quad l \in [L], \\ & \underline{\sigma}^{(l)}(z^{(l)}) \leq x^{(l+1)} \leq \bar{\sigma}^{(l)}(z^{(l)}), \quad l \in [L], \end{aligned} \quad (\mathcal{C})$$

where  $\underline{\sigma}^{(l)}(z)$  ( $\bar{\sigma}^{(l)}(z)$ ) is convex (concave) and satisfies  $\underline{\sigma}^{(l)}(z) \leq \sigma^{(l)}(z) \leq \bar{\sigma}^{(l)}(z)$  for  $\underline{z}^{(l)} \leq z \leq \bar{z}^{(l)}$ . We denote the feasible set of  $(\mathcal{C})$  by  $\mathcal{S}_{\mathcal{C}}$  and its minimum by  $p_{\mathcal{C}}^*$ . Naturally, we have that  $\mathcal{S}_{\mathcal{C}}$  is convex and  $p_{\mathcal{C}}^* \leq p_{\mathcal{O}}^*$ .

For example, [Ehlers \(2017\)](#) proposed the following relaxation for the ReLU function  $\sigma(z) = \max(0, z)$ :

$$\underline{\sigma}(z) = \max(0, z), \quad \bar{\sigma}(z) = \frac{\bar{z}}{\bar{z} - \underline{z}} (z - \underline{z}), \quad (4)$$

and the following for MaxPool  $\sigma(z) = \max_k z_k$ :

$$\underline{\sigma}(z) = \max_k z_k \geq \sum_k (z_k - \bar{z}_k) + \max_k \bar{z}_k, \quad \bar{\sigma}(z) = \sum_k (z_k + \underline{z}_k) - \max_k \underline{z}_k. \quad (5)$$

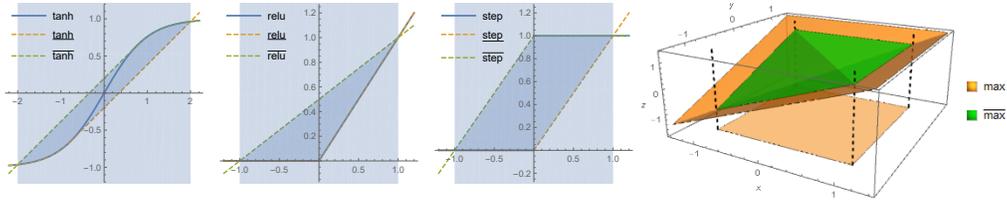


Figure 2: Optimal convex relaxations for common nonlinearities. For tanh, the relaxation contains two linear segments and parts of the tanh function. For ReLU and the step function, the optimal relaxations are written as 3 and 4 linear constraints, respectively. For  $z = \max(x, y)$ , the light orange shadow indicates the pre-activation bounds for  $x$  and  $y$ , and the optimal convex relaxation is lower bounded by the max function itself.

**The optimal layer-wise convex relaxation.** As a special case, we consider the optimal layer-wise convex relaxation, where  $\underline{\sigma}(z)$  is the greatest convex function majored by  $\sigma$ , and  $\bar{\sigma}(z)$  is the smallest concave function majoring  $\sigma$ . They can be equivalently defined as

$$\underline{\sigma}(z) := \sup_{\{(\alpha, \gamma): \alpha^\top z' + \gamma \leq \sigma(z'), \forall z' \in [z, \bar{z}]\}} \{\alpha^\top z + \gamma\}, \quad (6)$$

$$\bar{\sigma}(z) := \inf_{\{(\alpha, \gamma): \alpha^\top z' + \gamma \geq \sigma(z'), \forall z' \in [z, \bar{z}]\}} \{\alpha^\top z + \gamma\}. \quad (7)$$

In Fig. 2, we show the optimal convex relaxation for several common activation functions. It is easy to see that (4) is the optimal convex relaxation for ReLU, but (5) is not optimal for the MaxPool function.

When  $x^{(l+1)} = \sigma^{(l)}(z^{(l)})$  is a nonlinear layer that has a vector output  $x^{(l+1)} \in \mathbb{R}^{n^{(l+1)}}$ , the optimal convex relaxation may not have a simple analytic form as  $\underline{\sigma}^{(l)}(z^{(l)}) \leq x^{(l+1)} \leq \bar{\sigma}^{(l)}(z^{(l)})$ . Fortunately, if there is no interaction (as defined below) among the output neurons, the optimal convex relaxation can be given as a simple analytic form.

**Definition 3.1** (non-interactive layer). *Let  $\sigma : \mathbb{R}^m \rightarrow \mathbb{R}^n$  and  $x = \sigma(z)$  be a nonlinear layer with input  $z \in [z, \bar{z}] \subset \mathbb{R}^m$  and output  $x \in \mathbb{R}^n$ . For each output  $x_j$ , let  $I_j \subset [m]$  be the minimal set of  $z$ 's entries that affect  $x_j$ , where  $x_j = \sigma(z_{I_j})$ . We call the layer  $x = \sigma(z)$  non-interactive if the sets  $I_j$  ( $j \in [n]$ ) are mutually disjoint.*

Commonly used nonlinear activation layers are all non-interactive. It is obvious that all entry-wise nonlinear layers, such as (leaky-)ReLU and sigmoid, are non-interactive. A MaxPool layer with non-overlapping regions (stride no smaller than kernel size) is also non-interactive. Finally, any layer with scalar-valued output is non-interactive. When we treat a general nonlinear specification (as proposed in Qin et al. (2019)) as an additional nonlinear layer  $x^{(L+1)} = F(x^{(0)}, x^{(L)})$ , this layer is automatically non-interactive. This nice property ensures that our framework can deal with very general specifications.

The optimal convex relaxation of a non-interactive layer has a simple analytic form as below.

**Proposition 3.2.** *If the layer  $\sigma^{(l)} : [z^{(l)}, \bar{z}^{(l)}] \rightarrow \mathbb{R}^{n^{(l+1)}}$  is non-interactive, we have*

$$\{(z^{(l)}, x^{(l+1)}) : \underline{\sigma}^{(l)}(z^{(l)}) \leq x^{(l+1)} \leq \bar{\sigma}^{(l)}(z^{(l)})\} = \overline{\text{conv}}(\{(z^{(l)}, x^{(l+1)}) : x^{(l+1)} = \sigma^{(l)}(z^{(l)}), z^{(l)} \leq z^{(l)} \leq \bar{z}^{(l)}\}),$$

where  $\overline{\text{conv}}$  denotes the closed convex hull, and scalar-valued functions  $\underline{\sigma}^{(l)}(z)$  and  $\bar{\sigma}^{(l)}(z)$  are defined in (6) and (7), respectively.

We defer the proof of Proposition 3.2 to Appendix A.1.

**Convex relaxations not included in Problem (C).** We emphasize that by *optimal*, we mean the optimal convex relaxation of the *single* nonlinear constraint  $x^{(l+1)} = \sigma^{(l)}(z^{(l)})$  (see Proposition (3.2)) instead of the optimal convex relaxation of the nonconvex feasible set of the original problem (C). In fact, for neural networks with more than two hidden layers ( $L \geq 2$ ), the optimal convex relaxation of

the nonconvex feasible set of problem (O) is a strict subset of the feasible set of problem (C), even with the tightest bounds  $(\underline{z}^{[L]}, \bar{z}^{[L]})$  and the optimal choice of  $\underline{\sigma}^{(l)}(z)$  and  $\bar{\sigma}^{(l)}(z)$  in (6) and (7). It is possible to obtain other (maybe tighter) convex relaxations (Anderson et al., 2018), but it comes with more assumptions on the nonlinear layers and more complex convex constraints.

For example, Raghunathan et al. (2018b) rewrites the ReLU nonlinearity as a quadratic constraint, and then proposes a semidefinite programming (SDP) relaxation for the resulting quadratic optimization problem. Problem (C) does not cover this SDP-relaxation. Sometimes Problem (C) provides tighter relaxation than the SDP-relaxation, e.g., the case when there is only one neuron in a layer, while sometimes the SDP-relaxation provides tighter relaxation than Problem (C), e.g., the examples provided in Raghunathan et al. (2018b). The SDP-relaxation currently only works for ReLU nonlinearity. It is not clear to us how to extend the SDP-relaxed verifier to general nonlinearities. On the other hand, Problem (C) can handle any non-interactive nonlinear layer and any nonlinear specification.

**Greedily solving the primal with linear bounds.** Before resorting to duality for large problems, it is possible to solve (C) by over-relaxing the problem to efficiently give a lower bound for the objective. This can be done by defining *exactly one* linear upper bound and one linear lower bound for each activation function  $\sigma^{(l)}(z^{(l)})$  in (C) as follows,

$$\bar{\sigma}^{(l)}(z^{(l)}) := \bar{a}^{(l)} z^{(l)} + \bar{b}^{(l)}, \quad \underline{\sigma}^{(l)}(z^{(l)}) := \underline{a}^{(l)} z^{(l)} + \underline{b}^{(l)}.$$

Given these bounds for each activation, the primal variables  $z_i^{(l)}$  can be *greedily* bounded in a layer-by-layer manner: we can derive one linear upper and one linear lower bound of  $z_i^{(l)}$  with respect to  $z^{(l-1)}$  using the fact that  $z_i^{(l)} = \mathbf{W}_{i,:}^{(l)} \sigma^{(l-1)}(z^{(l-1)}) + b^{(l)}$ , and that  $\sigma^{(l-1)}(z^{(l-1)})$  is linearly upper and lower bounded by  $\bar{\sigma}^{(l-1)}(z^{(l-1)})$  and  $\underline{\sigma}^{(l-1)}(z^{(l-1)})$ . Because a linear combination of linear bounds (coefficients are related to the entries in  $\mathbf{W}$ ) is a single linear bound, we can apply this technique again and replace  $z^{(l-1)}$  with its upper and lower bounds with respect to  $z^{(l-2)}$ , obtaining the bound for  $z_i^{(l)}$  with respect to  $z^{(l-2)}$ . Applying this repeatedly eventually leads to linear lower and upper bounds of  $z_i^{(l)}$  with respect to the input  $x^{(0)} \in \mathcal{S}_{in}(x^{\text{nom}})$ .

This perspective covers Fast-Lin (Weng et al., 2018), DeepZ (Singh et al., 2018) and Neurify (Wang et al., 2018b), where the proposed linear lower bound has the same slope as the upper bound, i.e.,  $\underline{a}_k^{(l)} = \bar{a}_k^{(l)}$ . The resulting shape is referred to as a *zonotope* in Gehr et al. (2018) and Singh et al. (2018). In CROWN (Zhang et al., 2018) and DeepPoly (Singh et al., 2019b), this restriction is lifted and they can achieve better verification results than Fast-Lin and DeepZ. Fig. 1 summarizes the relationships between these algorithms. Importantly, each of these works has its own merits on solving the verification problem; our focus here is to give a unified view on how they over-relax C in our framework. See Appendix A.2 for more discussions and other related algorithms.

## 4 Convex Relaxation from the Dual View

We now tackle the verification problem from the dual view and show its connection to the primal view.

**Strong duality for the convex relaxed problem.** As in Wong & Kolter (2018), we introduce the dual variables for (C) and write its Lagrangian dual as

$$\begin{aligned} g_{\mathcal{C}}(\mu^{[L]}, \underline{\lambda}^{[L]}, \bar{\lambda}^{[L]}) := & \min_{(x^{[L+1]}, z^{[L]}) \in \mathcal{D}} c^\top x^{(L)} + c_0 + \sum_{l=0}^{L-1} \mu^{(l)\top} (z^{(l)} - \mathbf{W}^{(l)} x^{(l)} - b^{(l)}) \\ & - \sum_{l=0}^{L-1} \underline{\lambda}^{(l)\top} (x^{(l+1)} - \underline{\sigma}^{(l)}(z^{(l)})) + \sum_{l=0}^{L-1} \bar{\lambda}^{(l)\top} (x^{(l+1)} - \bar{\sigma}^{(l)}(z^{(l)})). \end{aligned} \quad (8)$$

Note that for any choice of  $\mu^{[L]}, \underline{\lambda}^{[L]} \geq 0$  and  $\bar{\lambda}^{[L]} \geq 0$ , the value of the dual function  $g_{\mathcal{C}}$  lower bounds the optimal value of (C), and thus we have

$$d_{\mathcal{C}}^* := \max_{\mu^{[L]}, \underline{\lambda}^{[L]} \geq 0, \bar{\lambda}^{[L]} \geq 0} g_{\mathcal{C}}(\mu^{[L]}, \underline{\lambda}^{[L]}, \bar{\lambda}^{[L]}) \leq p_{\mathcal{C}}^*. \quad (9)$$

This property is known as weak duality (Boyd & Vandenberghe, 2004). In the following, we prove that strong duality between (C) and (9) holds under some mild conditions.

**Theorem 4.1.** *If, for each  $l \in [L]$ , both  $\underline{\sigma}^{(l)}$  and  $\bar{\sigma}^{(l)}$  have a finite Lipschitz constant in the domain  $[\underline{z}^{(l)}, \bar{z}^{(l)}]$ , then strong duality holds between (C) and (9), i.e.,*

$$p_{\mathcal{C}}^* = d_{\mathcal{C}}^*.$$

The assumption of the finite Lipschitz constant is so weak that one can always construct such  $\underline{\sigma}^{(l)}$  and  $\bar{\sigma}^{(l)}$  as long as  $\sigma^{(l)}$  is bounded in the domain  $[\underline{z}^{(l)}, \bar{z}^{(l)}]$ . Notice that we do not make additional geometric assumptions (like Slater's condition) in Theorem (4.1). Instead, we make use of geometric properties of NNs, i.e., their linear stability against perturbations in the layer-wise computation, to obtain our strong duality results. We present the linear stability analysis in Lemma A.1 and provide the strong duality proof in Appendix A.3.

**The optimal dual relaxation.** If we restrict ourselves to the dual of the layer-wise convex relaxed problem (C), the optimal (tightest) bound the dual problem can provide is exactly the bound the optimal layer-wise convex relaxation ((6) and (7)) provides, thanks to the strong duality.

One approach to possibly obtain a tighter dual problem is to directly study the Lagrangian dual of the original problem (O),

$$\begin{aligned} g_{\mathcal{O}}(\mu^{[L]}, \lambda^{[L]}) := & \min_{(x^{[L+1]}, z^{[L]}) \in \mathcal{D}} c^\top x^{(L)} + c_0 \\ & + \sum_{l=0}^{L-1} \mu^{(l)\top} (z^{(l)} - \mathbf{W}^{(l)} x^{(l)} - b^{(l)}) \\ & + \sum_{l=0}^{L-1} \lambda^{(l)\top} (x^{(l+1)} - \sigma^{(l)}(z^{(l)})), \end{aligned} \quad (10)$$

which was first proposed in Dvijotham et al. (2018d). Note that for any choice of  $\mu^{[L]}$  and  $\lambda^{[L]}$ , the value of the dual function  $g_{\mathcal{O}}$  lower bounds the optimal value of (O). Thus,

$$d_{\mathcal{O}}^* := \max_{\mu^{[L]}, \lambda^{[L]}} g_{\mathcal{O}}(\mu^{[L]}, \lambda^{[L]}) \leq p_{\mathcal{O}}^*, \quad (11)$$

It would seem to be strictly better than (9). Unfortunately, the following theorem shows that the dual problem (11) can do no better than the dual of the convex relaxation (9).

**Theorem 4.2.** *Assume that the nonlinear layer  $\sigma^{(l)}$  is non-interactive and  $\underline{\sigma}^{(l)}$  and  $\bar{\sigma}^{(l)}$  are defined in (6) and (7), respectively. Then the lower bound  $d_{\mathcal{C}}^*$  provided by the dual of the convex-relaxed problem (9) and  $d_{\mathcal{O}}^*$  provided by the dual of the original problem (11) are the same, i.e.,*

$$d_{\mathcal{C}}^* = d_{\mathcal{O}}^*.$$

*Proof sketch.* First, we prove that for any  $\mu^{[L]}, \underline{\lambda}^{[L]} \geq 0$  and  $\bar{\lambda}^{[L]} \geq 0$ , with  $\lambda^{[L]} = \bar{\lambda}^{[L]} - \underline{\lambda}^{[L]}$ , we have

$$g_{\mathcal{C}}(\mu^{[L]}, -\lambda_-^{[L]}, \lambda_+^{[L]}) \leq g_{\mathcal{C}}(\mu^{[L]}, \underline{\lambda}^{[L]}, \bar{\lambda}^{[L]}), \quad (12)$$

where  $\lambda_+ = \max(\lambda, 0)$  and  $\lambda_- = \min(\lambda, 0)$ . Therefore, the dual problem (9) can be rewritten as an unconstrained optimization problem as

$$d_{\mathcal{C}}^* = \max_{\mu^{[L]}, \lambda^{[L]}} g_{\mathcal{C}}(\mu^{[L]}, -\lambda_-^{[L]}, \lambda_+^{[L]}). \quad (13)$$

Second, we prove that for any  $\mu^{[L]}$  and  $\lambda^{[L]}$ , we have

$$g_{\mathcal{C}}(\mu^{[L]}, -\lambda_-^{[L]}, \lambda_+^{[L]}) = g_{\mathcal{O}}(\mu^{[L]}, \lambda^{[L]}) \quad (14)$$

Combining (11), (13) and (14), we obtain  $d_{\mathcal{C}}^* = d_{\mathcal{O}}^*$ .  $\square$

The complete proof is in Appendix A.4. From (13) and (14), we can see that these two dual problems not only share the optimal value, but also share the same objective function! We would like to point

out that Theorem 2 in Dvijotham et al. (2018d) is a special case of our Theorem 4.2, when applied to ReLU networks. Our proof makes use of the Fenchel-Moreau theorem (Theorem 12.2 in Rockafellar (2015)) to deal with general nonlinearities, which is different from that in Dvijotham et al. (2018d).

Theorem 4.2 combined with the strong duality result of Theorem 4.1 implies that the primal relaxation (C) and the two kinds of dual relaxations, (9) and (11), are all blocked by the same *barrier* — the bound that the optimal layer-wise convex relaxation ((6) and (7)) provides.

**Corollary 4.3.** *Suppose that the activation functions  $\sigma^{(l)}$  ( $l \in [L]$ ) are non-interactive and are among the following functions: piece-wise linear functions (e.g., ReLU and stair-case functions<sup>2</sup>), ELU, sigmoid, arctan, hyperbolic tangent, polynomials and common pooling functions (e.g., MaxPool and AvgPool). Assume that  $\underline{\sigma}^{(l)}$  and  $\bar{\sigma}^{(l)}$  are defined in (6) and (7), respectively. Then we have that the lower bound  $p_C^*$  provided by the primal relaxation (C) and  $d_O^*$  provided by the dual relaxation (11) are the same, i.e.,*

$$p_C^* = d_O^*.$$

We point out that the constraint on the activation function is to ensure the strong duality between (C) and (9) when  $\underline{\sigma}^{(l)}$  and  $\bar{\sigma}^{(l)}$  are defined in (6) and (7), respectively. This constraint is so weak that it is satisfied by nearly all commonly used activation functions, far more than the functions listed in Corollary 4.3.

**Which problem to solve in practice?** Thanks to the strong duality, the same lower bound can be achieved from both the primal and the dual problems, and thus we have the freedom to choose the problem to solve. When the relaxed upper and lower bounds, i.e.,  $\underline{\sigma}^{(l)}$  and  $\bar{\sigma}^{(l)}$ , are piece-wise linear (e.g. (4) for ReLU networks), both the primal and dual problems are linear programs and can be efficiently solved by existing LP solvers (which is what we use in the coming sections). In other cases, we recommend to solve the dual problem (11) for two reasons. First, the primal relaxed problem (C) is a constrained optimization problem, and its constraints may not have a simple analytic form when  $\underline{\sigma}^{(l)}$  and  $\bar{\sigma}^{(l)}$  are not piecewise linear; see examples in Fig. 2. On the contrary, the dual problem (11) can be framed as an unconstrained optimization problem and its objective function has a simple analytic form for some common activation functions (Dvijotham et al., 2018d). Second, the optimization process of (11) can be stopped anytime to give a lower bound of  $p_C^*$ , thanks to weak duality, but this is not true of the primal view. Of course,  $\underline{\sigma}^{(l)}$  and  $\bar{\sigma}^{(l)}$  must be in the form of (6) and (7) to achieve the optimal value.

**Greedily solving the dual with linear bounds.** When the relaxed bounds  $\underline{\sigma}$  and  $\bar{\sigma}$  are linear, i.e.,

$$\underline{\sigma}^{(l)}(z^{(l)}) := \underline{a}^{(l)} z^{(l)} + \underline{b}^{(l)}, \quad \bar{\sigma}^{(l)}(z^{(l)}) := \bar{a}^{(l)} z^{(l)} + \bar{b}^{(l)},$$

the dual objective (9) can be approximately solved (lower bounded) by a greedy algorithm. For  $l \in \{L-1, L-2, \dots, 1, 0\}$ , this greedy algorithm fixes the already obtained sub-optimal solution of  $(\mu^{[l+1:L]}, \underline{\lambda}^{[l+1:L]}, \bar{\lambda}^{[l+1:L]})$ , then computes a sub-optimal solution for  $(\mu^{(l)}, \underline{\lambda}^{(l)}, \bar{\lambda}^{(l)})$  by maximizing the part of the objective function that only involves  $(\mu^{[l:L]}, \underline{\lambda}^{[l:L]}, \bar{\lambda}^{[l:L]})$ . We provide the algorithm details in Appendix A.5.

It turns out that this greedy algorithm exactly recovers the algorithm proposed in Wong & Kolter (2018) when

$$\underline{\sigma}^{(l)}(z^{(l)}) := \alpha^{(l)} z^{(l)}, \quad \bar{\sigma}^{(l)}(z^{(l)}) := \frac{\bar{z}^{(l)}}{\bar{z}^{(l)} - \underline{z}^{(l)}} (z^{(l)} - \underline{z}^{(l)}),$$

where  $0 \leq \alpha^{(l)} \leq 1$  represents the slope of the lower bound. When  $\alpha^{(l)} = \frac{\bar{z}^{(l)}}{\bar{z}^{(l)} - \underline{z}^{(l)}}$ , the greedy algorithm also recovers Fast-Lin (Weng et al., 2018), which explains the arrow from Wong & Kolter (2018) to Weng et al. (2018) in Fig. 1. When  $\alpha^{(l)}$  is chosen adaptively as in CROWN (Zhang et al., 2018), the greedy algorithm then recovers CROWN, which explains the arrow from Wong & Kolter (2018) to Zhang et al. (2018) in Fig. 1. See Appendix A.2 for more discussions on the relationship between the primal and dual greedy solvers.

<sup>2</sup>Note that we only require  $\underline{\sigma}^{(l)}$  and  $\bar{\sigma}^{(l)}$  defined in (6) and (7) be Lipschitz continuous, while  $\sigma$  itself may not be Lipschitz.

## 5 Optimal LP-relaxed Verification

In the previous sections, we presented a framework that subsumes all existing layer-wise convex-relaxed verification algorithms except that of [Raghunathan et al. \(2018b\)](#). For ReLU networks, being piece-wise linear, these correspond exactly to the set of all existing LP-relaxed algorithms, as discussed above. We showed the existence of a barrier,  $p_C^*$ , that limits all such algorithms. Is this just theoretical babbling or is this barrier actually problematic in practice?

In the next section, we perform extensive experiments on deep ReLU networks, evaluating the tightest convex relaxation afforded by our framework (denoted **LP-ALL**) against a greedy dual algorithm (Algorithm 1 of [Wong & Kolter \(2018\)](#), denoted **LP-GREEDY**) as well as another algorithm **LP-LAST**, intermediate in speed and accuracy between them. Both LP-GREEDY and LP-LAST solve the bounds  $\underline{z}^{[L]}, \bar{z}^{[L]}$  by setting the dual variables heuristically (see previous section), but LP-GREEDY solves the adversarial loss in the same manner while LP-LAST solves this final LP exactly. We also compare them with the opposite bounds provided by PGD attack ([Madry et al., 2017](#)), as well as exact results from MILP ([Tjeng et al., 2019](#)).

For the rest of the main text, we are only concerned with ReLU networks, so (C) subject to (4) is in fact an LP.

### 5.1 LP-ALL Implementation Details

In order to exactly solve the tightest LP-relaxed verification problem of a ReLU network, two steps are required: (A) obtaining the tightest pre-activation upper and lower bounds of all the neurons in the NN, excluding those in the last layer, then (B) solving the LP-relaxed verification problem exactly for the last layer of the NN.

**Step A: Obtaining Pre-activation Bounds.** This can be done by solving sub-problems of the original relaxed problem (C) subject to (4). Given a NN with  $L_0$  layers, for each layer  $l_0 \in [L_0]$ , we obtain a lower (resp. upper) bound  $\underline{z}_j^{(l_0)}$  (resp.  $\bar{z}_j^{(l_0)}$ ) of  $z_j^{(l_0)}$ , for all neurons  $j \in [n^{(l_0)}]$ . We do this by setting

$$\begin{aligned} L &\leftarrow l_0, \\ c^\top &\leftarrow \mathbf{W}_{j,:}^{(l_0)} \text{ (resp. } c^\top \leftarrow -\mathbf{W}_{j,:}^{(l_0)}), \\ c_0 &\leftarrow b_j^{(l_0)} \text{ (resp. } c_0 \leftarrow -b_j^{(l_0)}) \end{aligned}$$

in (C) and computing the exact optimum.

However, we need to solve an LP for each neuron, and practical networks can have millions of them. We utilize the fact that in each layer  $l_0$ , computing the bounds  $\bar{z}_j^{(l_0)}$  and  $\underline{z}_j^{(l_0)}$  for each  $j \in [n^{(l_0)}]$  can proceed independently in parallel. Indeed, we design a scheduler to do so on a cluster with 1000 CPU-nodes. See Appendix C for details.

**Step B: Solving the LP-relaxed Problem for the Last Layer.** After obtaining the pre-activation bounds on all neurons in the network using step (A), we solve the LP in (C) subject to (4) for all  $j \in [n^{(L_0)}] \setminus \{j^{\text{nom}}\}$  obtained by setting

$$\begin{aligned} L &\leftarrow L_0, \\ c^\top &\leftarrow \mathbf{W}_{j^{\text{nom}},:}^{(L_0)} - \mathbf{W}_{j,:}^{(L_0)}, \\ c_0 &\leftarrow b_{j^{\text{nom}}}^{(L_0)} - b_j^{(L_0)} \end{aligned}$$

again in (C) and computing the exact minimum.

Here,  $j^{\text{nom}}$  is the true label of the data point  $x^{\text{nom}}$  at which we are verifying the network. *We can certify the network is robust around  $x^{\text{nom}}$  iff the solutions of all such LPs are positive, i.e. we cannot make the true class logit lower than any other logits.* Again, note that these LPs are also independent of each other, so we can solve them in parallel.

Given any  $x^{\text{nom}}$ , LP-ALL follows steps (A) then (B) to produce a certificate whether the network is robust around a given datapoint or not. LP-LAST on the other hand solves only step (B), and instead of doing (A), it finds the preactivation bounds greedily as in Algorithm 1 of [Wong & Kolter \(2018\)](#).

Table 1: Certified bounds on the robust error on the test set of MNIST for normally and robustly trained networks. The prefix of each network corresponds to the training method used: **ADV** for PGD training (Madry et al., 2017), **NOR** for normal CE loss training, and **LPD** when the LP-relaxed dual formulation of Wong & Kolter (2018) is used for robust training.

NETWORK	$\epsilon$	TEST ERROR	LOWER BOUND		UPPER BOUND		
			PGD	MILP	MILP	LP-ALL	LP-GREEDY
ADV-MLP-B	0.03	1.53%	4.17%	4.18%	5.78%	10.04%	13.40%
ADV-MLP-B	0.05	1.62%	6.06%	6.11%	11.38%	23.29%	33.09%
ADV-MLP-B	0.1	3.33%	15.86%	16.25%	34.37%	61.59%	71.34%
ADV-MLP-A	0.1	4.18%	11.51%	14.36%	30.81%	60.14%	67.50%
NOR-MLP-B	0.02	2.05%	10.06%	10.16%	13.48%	26.41%	35.11%
NOR-MLP-B	0.03	2.05%	20.37%	20.43%	48.67%	65.70%	75.85%
NOR-MLP-B	0.05	2.05%	53.37%	53.37%	94.04%	97.95%	99.39%
LPD-MLP-B	0.1	4.09%	13.39%	14.45%	14.45%	17.24%	18.32%
LPD-MLP-B	0.2	15.72%	33.85%	36.33%	36.33%	37.50%	41.67%
LPD-MLP-B	0.3	39.22%	57.29%	59.85%	59.85%	60.17%	66.85%
LPD-MLP-B	0.4	67.97%	81.85%	83.17%	83.17%	83.62%	87.89%

## 6 Experiments

We conduct two experiments to assess the tightness of LP-ALL: 1) finding certified upper bounds on the robust error of several NN classifiers, 2) finding certified lower bounds on the minimum adversarial distortion  $\epsilon$  using different algorithms. All experiments are conducted on MNIST and/or CIFAR-10 datasets.

**Architectures.** We conduct experiments on a range of ReLU-activated feedforward networks. MLP-A and MLP-B refer to multilayer perceptrons: MLP-A has 1 hidden layer with 500 neurons, and MLP-B has 2 hidden layers with 100 neurons each. CNN-SMALL, CNN-WIDE-K, and CNN-DEEP-K are the ConvNet architectures used in Wong et al. (2018). Full details are in Appendix B.1.

**Training Modes.** We conduct experiments on networks trained with a regular cross-entropy (CE) loss function and networks trained to be robust. These networks are identified by a prefix corresponding to the method used to train them: **LPD** when the LP-relaxed dual formulation of Wong & Kolter (2018) is used for robust training, **ADV** when adversarial examples generated using PGD are used for robust training, as in Madry et al. (2017), and **NOR** when the network is normally trained using the CE loss function. Training details are in Appendix B.2.

**Experimental Setup.** We run experiments on a cluster with 1000 CPU-nodes. The total run time amounts to more than 22 CPU-years. Appendix C provides additional details about the computational resources and the scheduling scheme used, and Appendix D provides statistics of the verification time in these experiments.

Our code including pretrained models can be found at <https://github.com/Hadisalman/robust-verify-benchmark>.

### 6.1 Certified Bounds on the Robust Error

Table 1 presents the clean test errors and (upper and lower) bounds on the true robust errors for a range of classifiers trained with different procedures on MNIST. For both ADV- and LPD-trained networks, the  $\epsilon$  in Table 1 denotes the  $l_\infty$ -norm bound used for training *and* robust testing; for NORMALLY-trained networks,  $\epsilon$  is only used for the latter.

Lower bounds on the robust error are calculated by finding adversarial examples for inputs that are not robust. This is done by using PGD, a strong first-order attack, or using MILP (Tjeng et al., 2019). Upper bounds on the robust error are calculated by providing certificates of robustness for input that is robust. This is done using MILP, the dual formulation (LP-GREEDY) presented by Wong & Kolter (2018), or our LP-ALL algorithm.

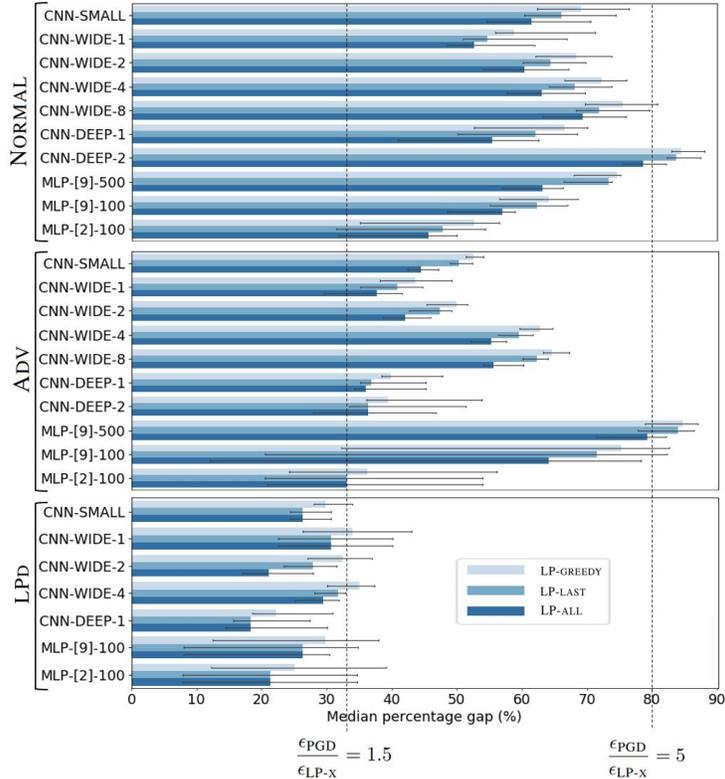


Figure 3: The median percentage gap between the convex-relaxed algorithms (LP-ALL, LP-LAST, and LP-GREEDY) and PGD estimates of the minimum adversarial distortion  $\epsilon$  on ten samples of MNIST. The error bars correspond to 95% confidence intervals. We highlight the  $1.5\times$  and  $5\times$  gaps between the  $\epsilon$  value estimated by PGD, and those estimated by the LP-relaxed algorithms. For more details, please refer to Table 2 in Appendix E.2.

For the MILP results, we use the code accompanying the paper by Tjeng et al. (2019). We run the code in parallel on a cluster with 1000 CPU-nodes, and set the MILP solver’s time limit to 3600 seconds. Note that this time limit is reached for ADV and NOR, and therefore the upper and lower bounds are separated by a gap that is especially large for some of the NORMally trained networks. On the other hand, for LPD-trained networks, the MILP solver finishes within the time limit, and thus the upper and lower bounds match.

For all NORMally and ADV-trained networks, we see that the certified upper bounds using LP-GREEDY and LP-ALL are very loose when we compare the gap between them to the lower bounds found by PGD and MILP. As a sanity check, note that LP-ALL gives a tighter bound than LP-GREEDY in each case, as one would expect. Yet this improvement is not significant enough to close the gap with the lower bounds.

This sanity check also passes for LPD-trained networks, where the LP-GREEDY-certified robust error upper bound is, as expected, much closer to the true error (given by MILP here) than for other networks. For  $\epsilon = 0.1$ , the improvement of LP-ALL-certified upper bound over LP-GREEDY is at most modest, and the PGD lower bound is tighter to the true error. For large  $\epsilon$ , the improvement is much more significant in relative terms, but the absolute improvement is only 4 – 7%. In this large  $\epsilon$  regime, however, both the clean and robust errors are quite large, so the tightness of LP-ALL is less useful.

## 6.2 Certified Bounds on the Minimum Adversarial Distortion $\epsilon$

We are interested in searching for the minimum adversarial distortion  $\epsilon$ , which is the radius of the largest  $l_\infty$  ball in which no adversarial examples can be crafted. An upper bound on  $\epsilon$  is calculated

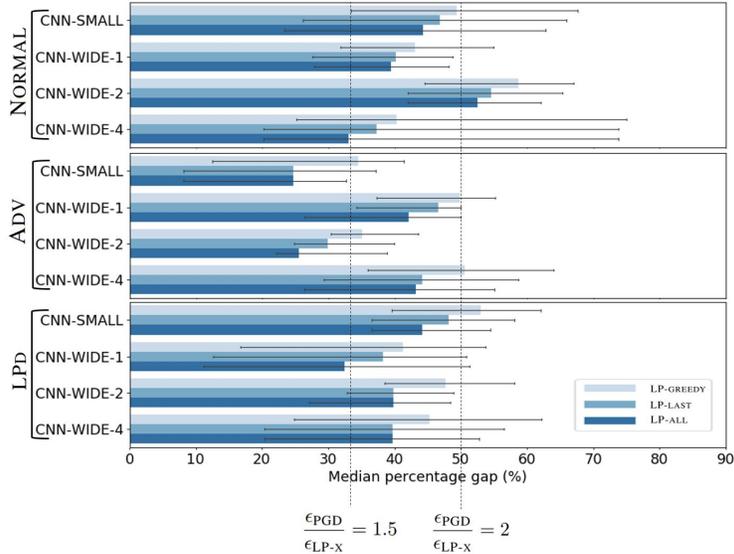


Figure 4: The median percentage gap of minimum adversarial distortion for CIFAR-10, in the same format as Fig. 3. For more details, please refer to Table 3 in Appendix E.2.

using PGD, and lower bounds are calculated using LP-GREEDY, LP-LAST, or our LP-ALL, all via binary search. Since solving LP-ALL is expensive, we find the  $\epsilon$ -bounds only for ten samples of the MNIST and CIFAR-10 datasets. In this experiment, both ADV- and LPD-networks are trained with an  $l_\infty$  maximum allowed perturbation of 0.1 and 8/255 on MNIST and CIFAR-10, respectively. See Appendix E.1 for details. Fig. 3 and 4 show the median percentage gap (defined in Appendix E.2) between the convex-relaxed algorithms and PGD bounds of  $\epsilon$  for MNIST and CIFAR, respectively. Details are reported in Tables 2 and 3 in Appendix E.2.

On MNIST, the results show that for all networks trained NORMally or via ADV, the certified lower bounds on  $\epsilon$  are 1.5 to 5 times smaller than the upper bound found by PGD; for LPD trained networks, below 1.5 times smaller. On CIFAR-10, the bounds are between 1.5 and 2 times smaller across all models. The smaller gap for LPD is of course as expected following similar observations in prior work (Wong & Kolter, 2018; Tjeng et al., 2019). Furthermore, the improvement of LP-ALL and LP-LAST over LP-GREEDY is not significant enough to close the gap with the PGD upper bound. Note that similar results hold as well for randomly initialized networks (no training). To avoid clutter, we report these in Appendix F.

## 7 Conclusions and Discussions

In this work we first presented a layer-wise convex relaxation framework that unifies all previous LP-relaxed verifiers, in both primal and dual spaces. Then we performed extensive experiments to show that even the optimal convex relaxation for ReLU networks in this framework cannot obtain tight bounds on robust error in all cases we consider here. Thus any method will face a *convex relaxation barrier* as soon as it can be described by our framework.

Note that different applications have different requirements for the tightness of the verification, so our barrier could be a problem for some but not for others. In so far as the ultimate goal of robustness verification is to construct a training method to lower certified error, this barrier is not necessarily problematic — some such method could still produce networks for which convex relaxation as described by our framework produces accurate robust error bounds. An example is the recent work of Goyal et al. (2018) which shows that interval bound propagation, which often leads to loose certification bounds, can still be used for verified training, and is able to achieve state-of-the-art verified accuracy when carefully tuned. However, without a doubt, in all cases, tighter estimates should lead to better results, and we reveal a definitive ceiling on most current methods.

**How to bypass the barrier?** The primal problem in our framework ( $\mathcal{C}$ ) has several possible sources of loosenes:

- (1) We relax the nonlinearity  $\sigma^{(l)}$  on a box domain  $\{\underline{z}^{(l)} \leq z \leq \bar{z}^{(l)}\}$ . This relaxation is simple to perform, but might come at the cost of losing some correlations between the coordinates of  $z$  and of obtaining a looser relaxation. Note that our framework does consider the correlations between coordinates of  $z^{(l)}$  to get bounds for all later layers, however it relies on  $\underline{z}^{(l)}$  and  $\bar{z}^{(l)}$  which are considered individually, without interactions within the same layer.
- (2) We solve for the bounds  $\underline{z}^{[l]}, \bar{z}^{[l]}$  recursively, and we incur some gap for every recursion; a loose bound in earlier layers will make bounds for later layers even looser. This can be problematic for very deep networks or recurrent networks.
- (3) In the specific case of ReLU, we lose a bit every time we relax over an unstable neuron; one possible future direction is to combine branch-and-bound with convex relaxation to strategically split the domains of unstable neurons.

Any method that improves on any of the above issues can possibly bypass the barrier. On the other hand, exact verifiers (Katz et al., 2017; Ehlers, 2017), SDP-based verifiers (Raghunathan et al., 2018b,a), Lipschitz-constant-based verifiers (Zhang et al., 2019), and hybrid approaches (Bunel et al., 2018; Singh et al., 2019a) do not fall under the purview of our framework. In general, none of them are strictly better than the convex relaxation approach, sacrificing either speed or accuracy. However, it would be fruitful to consider combinations of these methods in the future, as done in Singh et al. (2019a). We hope our work will foster much thought in the community toward new relaxation paradigms for tight neural network verification.

## References

- Anderson, R., Huchette, J., Tjandraatmadja, C., and Vielma, J. P. Strong convex relaxations and mixed-integer programming formulations for trained neural networks. *arXiv preprint arXiv:1811.01988*, 2018.
- Boyd, S. and Vandenberghe, L. *Convex optimization*. Cambridge university press, 2004.
- Bunel, R. R., Turkaslan, I., Torr, P., Kohli, P., and Mudigonda, P. K. A unified view of piecewise linear neural network verification. In *Advances in Neural Information Processing Systems*, pp. 4795–4804, 2018.
- Carlini, N. and Wagner, D. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pp. 39–57. IEEE, 2017.
- Carlini, N., Katz, G., Barrett, C., and Dill, D. L. Provably minimally-distorted adversarial examples. *arXiv preprint arXiv:1709.10207*, 2017.
- Chen, P.-Y., Zhang, H., Sharma, Y., Yi, J., and Hsieh, C.-J. Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, pp. 15–26. ACM, 2017.
- Cheng, C.-H., Nührenberg, G., and Ruess, H. Maximum resilience of artificial neural networks. In *International Symposium on Automated Technology for Verification and Analysis*, pp. 251–268. Springer, 2017.
- Diamond, S. and Boyd, S. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.
- Domahidi, A., Chu, E., and Boyd, S. ECOS: An SOCP solver for embedded systems. In *European Control Conference (ECC)*, pp. 3071–3076, 2013.
- Dutta, S., Jha, S., Sankaranarayanan, S., and Tiwari, A. Output range analysis for deep feedforward neural networks. In *NASA Formal Methods Symposium*, pp. 121–138. Springer, 2018.
- Dvijotham, K., Garnelo, M., Fawzi, A., and Kohli, P. Verification of deep probabilistic models. *NeurIPS Workshop on Security in Machine Learning*, 12 2018a.
- Dvijotham, K., Garnelo, M., Fawzi, A., and Kohli, P. Verification of deep probabilistic models. *CoRR*, abs/1812.02795, 2018b. URL <http://arxiv.org/abs/1812.02795>.
- Dvijotham, K., Gowal, S., Stanforth, R., Arandjelovic, R., O’Donoghue, B., Uesato, J., and Kohli, P. Training verified learners with learned verifiers. *arXiv preprint arXiv:1805.10265*, 2018c.
- Dvijotham, K., Stanforth, R., Gowal, S., Mann, T., and Kohli, P. A dual approach to scalable verification of deep networks. *UAI*, 2018d.
- Ehlers, R. Formal verification of piece-wise linear feed-forward neural networks. In *International Symposium on Automated Technology for Verification and Analysis*, pp. 269–286. Springer, 2017.
- Fischetti, M. and Jo, J. Deep neural networks as 0-1 mixed integer linear programs: A feasibility study. *arXiv preprint arXiv:1712.06174*, 2017.
- Gehr, T., Mirman, M., Drachler-Cohen, D., Tsankov, P., Chaudhuri, S., and Vechev, M. AI 2: Safety and robustness certification of neural networks with abstract interpretation. In *2018 IEEE Symposium on Security and Privacy (SP)*, 2018.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *ICLR*, 2015.
- Gowal, S., Dvijotham, K., Stanforth, R., Bunel, R., Qin, C., Uesato, J., Mann, T., and Kohli, P. On the effectiveness of interval bound propagation for training verifiably robust models. *arXiv preprint arXiv:1810.12715*, 2018.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

- Hein, M. and Andriushchenko, M. Formal guarantees on the robustness of a classifier against adversarial manipulation. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 2266–2276, 2017.
- Hinton, G., Deng, L., Yu, D., Dahl, G. E., Mohamed, A.-r., Jaitly, N., Senior, A., Vanhoucke, V., Nguyen, P., Sainath, T. N., et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal processing magazine*, 29(6):82–97, 2012.
- Katz, G., Barrett, C., Dill, D. L., Julian, K., and Kochenderfer, M. J. Reluplex: An efficient smt solver for verifying deep neural networks. In *International Conference on Computer Aided Verification*, pp. 97–117. Springer, 2017.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Lomuscio, A. and Maganti, L. An approach to reachability analysis for feed-forward relu neural networks. *arXiv preprint arXiv:1706.07351*, 2017.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- Mirman, M., Gehr, T., and Vechev, M. Differentiable abstract interpretation for provably robust neural networks. In *International Conference on Machine Learning*, pp. 3575–3583, 2018.
- Qin, C., Dvijotham, K. D., O’Donoghue, B., Bunel, R., Stanforth, R., Goyal, S., Uesato, J., Swirszcz, G., and Kohli, P. Verification of non-linear specifications for neural networks. *ICLR*, 2019.
- Raghunathan, A., Steinhardt, J., and Liang, P. Certified defenses against adversarial examples. *International Conference on Learning Representations (ICLR)*, *arXiv preprint arXiv:1801.09344*, 2018a.
- Raghunathan, A., Steinhardt, J., and Liang, P. S. Semidefinite relaxations for certifying robustness to adversarial examples. In *Advances in Neural Information Processing Systems*, pp. 10900–10910, 2018b.
- Rockafellar, R. T. *Convex analysis*. Princeton university press, 2015.
- Scheibler, K., Winterer, L., Wimmer, R., and Becker, B. Towards verification of artificial neural networks. In *MBMV*, pp. 30–40, 2015.
- Singh, G., Gehr, T., Mirman, M., Püschel, M., and Vechev, M. Fast and effective robustness certification. In *Advances in Neural Information Processing Systems*, pp. 10825–10836, 2018.
- Singh, G., Gehr, T., Püschel, M., and Vechev, M. Robustness certification with refinement. *ICLR*, 2019a.
- Singh, G., Gehr, T., Püschel, M., and Vechev, M. An abstract domain for certifying neural networks. *Proceedings of the ACM on Programming Languages*, 3(POPL):41, 2019b.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- Tjeng, V., Xiao, K. Y., and Tedrake, R. Evaluating robustness of neural networks with mixed integer programming. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=HyGIIdRqtm>.
- Wang, S., Chen, Y., Abdou, A., and Jana, S. Mixtrain: Scalable training of formally robust neural networks. *arXiv preprint arXiv:1811.02625*, 2018a.
- Wang, S., Pei, K., Whitehouse, J., Yang, J., and Jana, S. Efficient formal safety analysis of neural networks. In *Advances in Neural Information Processing Systems*, pp. 6369–6379, 2018b.
- Weng, T.-W., Zhang, H., Chen, H., Song, Z., Hsieh, C.-J., Boning, D., Dhillon, I. S., and Daniel, L. Towards fast computation of certified robustness for ReLU networks. In *International Conference on Machine Learning*, 2018.

- Wong, E. and Kolter, Z. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *International Conference on Machine Learning (ICML)*, pp. 5283–5292, 2018.
- Wong, E., Schmidt, F., Metzen, J. H., and Kolter, J. Z. Scaling provable adversarial defenses. *Advances in Neural Information Processing Systems (NIPS)*, 2018.
- Xiao, K. Y., Tjeng, V., Shafiullah, N. M. M., and Madry, A. Training for faster adversarial robustness verification via inducing reLU stability. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=BJfIVjAcKm>.
- Zhang, H., Weng, T.-W., Chen, P.-Y., Hsieh, C.-J., and Daniel, L. Efficient neural network robustness certification with general activation functions. In *Advances in Neural Information Processing Systems (NIPS)*, dec 2018.
- Zhang, H., Zhang, P., and Hsieh, C.-J. Recurjac: An efficient recursive algorithm for bounding jacobian matrix of neural networks and its applications. *AAAI Conference on Artificial Intelligence*, 2019.

## A Additional Theoretical Results and Proofs

### A.1 The optimal convex relaxation

In this section, we study the optimal convex relaxation of the nonlinear constraint  $x = \sigma(z)$ . If there is only one neuron, i.e., only a scalar output, the following proposition gives its optimal convex relaxation.

**Proposition A.1.** *Suppose the activation function  $\sigma : [\underline{z}, \bar{z}] \rightarrow \mathbb{R}$  is bounded from above and below. Let  $\underline{\sigma}$  and  $-\bar{\sigma}$  be the greatest closed convex functions majored by  $\sigma$  and  $-\sigma$ , respectively, i.e.,*

$$\underline{\sigma}(z) := \sup_{\{(\alpha, \gamma) : \alpha z' + \gamma \leq \sigma(z'), \forall z' \in [\underline{z}, \bar{z}]\}} \{\alpha z + \gamma\}, \quad (15)$$

$$-\bar{\sigma}(z) := \sup_{\{(\alpha, \gamma) : \alpha z' + \gamma \leq -\sigma(z'), \forall z' \in [\underline{z}, \bar{z}]\}} \{\alpha z + \gamma\}. \quad (16)$$

Then we have,

1. Both  $\underline{\sigma}$  and  $\bar{\sigma}$  are continuous in  $[\underline{z}, \bar{z}]$ .

2.

$$\{(z, x) : \underline{\sigma}(z) \leq x \leq \bar{\sigma}(z), \underline{z} \leq z \leq \bar{z}\} = \overline{\text{conv}}(\{(z, x) : x = \sigma(z), \underline{z} \leq z \leq \bar{z}\}).$$

where  $\overline{\text{conv}}$  denotes the closed convex hull.

*Proof.*

1. By the boundedness of  $\sigma$  on  $[\underline{z}, \bar{z}]$ , we know that the effective domain of  $\underline{\sigma}$  and  $-\bar{\sigma}$  is  $[\underline{z}, \bar{z}]$ . By definition (15) and (16),  $\underline{\sigma}$  and  $-\bar{\sigma}$  are closed convex functions. By Theorem 10.2 in Rockafellar (2015), we know that both  $\underline{\sigma}$  and  $-\bar{\sigma}$  are continuous in  $[\underline{z}, \bar{z}]$ , so is  $\bar{\sigma}$ .

2. We first decompose the left-hand-side into 3 terms:

$$\{(z, x) : \underline{\sigma}(z) \leq x \leq \bar{\sigma}(z), \underline{z} \leq z \leq \bar{z}\} = \{\underline{\sigma}(z) \leq x\} \cap \{x \leq \bar{\sigma}(z)\} \cap \{\underline{z} \leq z \leq \bar{z}\}.$$

Let  $\mathcal{F} = \{(\alpha, \gamma) : \alpha z' + \gamma \leq \sigma(z'), \forall z' \in [\underline{z}, \bar{z}]\}$  and  $\bar{\mathcal{F}} = \{(\alpha, \gamma) : \alpha z' + \gamma \geq \sigma(z'), \forall z' \in [\underline{z}, \bar{z}]\}$ . For the first term, by definition (15) we have

$$\begin{aligned} \{\underline{\sigma}(z) \leq x\} &= \bigcap_{\{(\alpha, \gamma) : \alpha z' + \gamma \leq \sigma(z'), \forall z' \in [\underline{z}, \bar{z}]\}} \{\alpha z + \gamma \leq x\} \\ &= \bigcap_{\{(\alpha, \beta, \gamma) : \beta < 0, \alpha z' + \beta \sigma(z') + \gamma \leq 0, \forall z' \in [\underline{z}, \bar{z}]\}} \{\alpha z + \beta x + \gamma \leq 0\}. \end{aligned}$$

For the second term, by definition (16) we have

$$\begin{aligned} \{x \leq \bar{\sigma}(z)\} &= \bigcap_{\{(\alpha, \gamma) : \alpha z' + \gamma \leq -\sigma(z'), \forall z' \in [\underline{z}, \bar{z}]\}} \{\alpha z + \gamma \leq -x\} \\ &= \bigcap_{\{(\alpha, \beta, \gamma) : \beta > 0, \alpha z' + \beta \sigma(z') + \gamma \leq 0, \forall z' \in [\underline{z}, \bar{z}]\}} \{\alpha z + \beta x + \gamma \leq 0\}. \end{aligned}$$

For the third term, we have

$$\begin{aligned} \{\underline{z} \leq z \leq \bar{z}\} &= \bigcap_{\{(\alpha, \gamma) : \alpha z' + \gamma \leq 0, \forall z' \in [\underline{z}, \bar{z}]\}} \{\alpha z + \gamma \leq 0\} \\ &= \bigcap_{\{(\alpha, \beta, \gamma) : \beta = 0, \alpha z' + \beta \sigma(z') + \gamma \leq 0, \forall z' \in [\underline{z}, \bar{z}]\}} \{\alpha z + \beta x + \gamma \leq 0\}. \end{aligned}$$

Combining the three terms, we conclude the proof by

$$\begin{aligned} \{(z, x) : \underline{\sigma}(z) \leq x \leq \bar{\sigma}(z), \underline{z} \leq z \leq \bar{z}\} &= \bigcap_{\{(\alpha, \beta, \gamma) : \alpha z' + \beta \sigma(z') + \gamma \leq 0, \forall z' \in [\underline{z}, \bar{z}]\}} \{\alpha z + \beta x + \gamma \leq 0\} \\ &= \overline{\text{conv}}(\{(z, x) : x = \sigma(z), \underline{z} \leq z \leq \bar{z}\}), \end{aligned}$$

where we use the definition of closed convex hull in the last identity.  $\square$

**Proposition 3.2** characterizes the optimal convex relaxation for a non-interactive nonlinear layer. Thanks to its non-interaction, Proposition 3.2 is a direct consequence of item 2 in Proposition A.1.

## A.2 Greedily solving the primal with linear bounds.

In this section, we show how to greedily solve (C) by over-relaxing the problem to give a lower bound directly. We give exactly one linear upper bound and exactly one linear lower bound for each activation function in (C):

$$\begin{aligned} \min_{(x^{[L+1]}, z^{[L]}) \in \mathcal{D}} \quad & c^\top x^{(L)} + c_0 \\ \text{s.t.} \quad & z^{(l)} = \mathbf{W}^{(l)} x^{(l)} + b^{(l)}, \quad l \in \{0, \dots, L-1\}, \\ & \underline{a}^{(l)} z^{(l)} + \underline{b}^{(l)} \leq x^{(l+1)} \leq \bar{a}^{(l)} z^{(l)} + \bar{b}^{(l)}, \quad l \in \{0, \dots, L-1\}, \end{aligned} \quad (17)$$

Typically, the selection of  $\underline{a}^{(l)}$ ,  $\bar{a}^{(l)}$ ,  $\bar{b}^{(l)}$ ,  $\underline{b}^{(l)}$  can depend on  $\bar{z}^{(l)}$  and  $\underline{z}^{(l)}$  to minimize the error between the upper/lower bound and the activation function. For element-wise activation functions, the linear upper and lower bounds are usually also element-wise. For example, for an unstable ReLU neuron with  $\bar{z}_i^{(l)} > 0$  and  $\underline{z}_i^{(l)} < 0$ , one upper bound is  $x_i^{(l+1)} \leq \frac{\bar{z}_i^{(l)}}{\bar{z}_i^{(l)} - \underline{z}_i^{(l)}} z_i^{(l)} - \frac{\bar{z}_i^{(l)} \underline{z}_i^{(l)}}{\bar{z}_i^{(l)} - \underline{z}_i^{(l)}}$ . According to Proposition A.1, this is the optimal convex relaxation for the upper bound. For the lower bound, the optimal convex relaxation  $(x_i^{(l+1)} \geq z_i^{(l)}) \cap (x_i^{(l+1)} \geq 0)$  is not achievable as one linear function; we use any over-relaxed bounds  $x_i^{(l+1)} \geq \underline{a}_i^{(l)} z_i^{(l)}$  with  $0 \leq \underline{a}_i^{(l)} \leq 1$  as the lower bound. This perspective covers Fast-Lin (Weng et al., 2018), DeepZ (Singh et al., 2018) and Neurify (Wang et al., 2018b), where the lower bound is fixed as  $\underline{a}_i^{(l)} = \bar{a}_i^{(l)} = \frac{\bar{z}_i^{(l)}}{\bar{z}_i^{(l)} - \underline{z}_i^{(l)}}$ ; this is referred as a ‘‘zonotope’’ relaxation in AI<sup>2</sup> (Gehr et al., 2018) and DeepZ. AI<sup>2</sup> is a general technique of using ‘‘abstract transformers’’ (sound relaxations of neural network elements) to verify neural networks, but it uses suboptimal relaxations for ReLU non-linearity; DeepZ further refines the transformers for ReLU and significantly outperforms AI<sup>2</sup> (Singh et al., 2018). Other activation functions can be linearly bounded as discussed in CROWN (Zhang et al., 2018), DeepZ and DeepPoly (Singh et al., 2019b); CROWN and DeepPoly are also more general and do not require  $\underline{a}_i^{(l)} = \bar{a}_i^{(l)}$  to allow a more flexible selection of bounds.

Assuming we have obtained the linear upper and lower bounds for  $x^{(l+1)}$  with respect to  $z^{(l)}$ ,  $\underline{z}_i^{(l+1)}$  can be formed greedily as a linear combination of these linear bounds: we greedily select the upper bound  $x_i^{(l+1)} \leq \underline{a}_i^{(l)} z_i^{(l)} + \underline{b}_i^{(l)}$  when  $\mathbf{W}_{i,k}^{(l+1)}$  is negative, and select the lower bound  $x_i^{(l+1)} \geq \bar{a}_i^{(l)} z_i^{(l)} + \bar{b}_i^{(l)}$  otherwise. This bound reflects the worst case scenario without considering any other neurons:

$$z_i^{(l+1)} \geq \underline{z}_i^{(l+1)} := \underline{A}_{i,:}^{(l)} z^{(l)} + \underline{b}_i^{\prime(l)} \quad (18)$$

where matrix  $\underline{A}_{i,k}^{(l)} = \begin{cases} \mathbf{W}_{i,k}^{(l+1)} \bar{a}_k^{(l)}, & \mathbf{W}_{i,k}^{(l+1)} < 0 \\ \mathbf{W}_{i,k}^{(l+1)} \underline{a}_k^{(l)}, & \mathbf{W}_{i,k}^{(l+1)} \geq 0 \end{cases}$  reflects the chosen upper or lower bound based on the sign of  $\mathbf{W}_{i,k}^{(l+1)}$ , and vector  $\underline{b}_i^{\prime(l)} = \sum_{k, \mathbf{W}_{i,k}^{(l+1)} \geq 0} \mathbf{W}_{i,k}^{(l+1)} \underline{b}_k^{(l)} + \sum_{k, \mathbf{W}_{i,k}^{(l+1)} < 0} \mathbf{W}_{i,k}^{(l+1)} \bar{b}_k^{(l)} + b_i^{(l)}$ .

The lower bound  $\bar{z}_i^{(l+1)}$  can also be formed similarly. Eventually, we get one linear upper bound and one linear lower bound for  $z_i^{(l+1)}$ , written as:

$$\underline{A}_{i,:}^{(l)} z^{(l)} + \underline{b}_i^{\prime(l)} \leq z_i^{(l+1)} \leq \bar{A}_{i,:}^{(l)} z^{(l)} + \bar{b}_i^{\prime(l)} \quad (19)$$

A sharp-eyed reader can notice that it is possible to also get a similar bound for each component of  $z^{(l)}$  and plug it into (19), thus obtaining a linear upper bound and a linear lower bound for  $z_i^{(l+1)}$  with respect to  $z^{(l-1)}$ . To do this, we first substitute  $z^{(l)} = \mathbf{W}^{(l)} x^{(l)} + b^{(l)}$  into Eq. (19), obtaining

$$\underline{A}_{i,:}^{(l)} (\mathbf{W}^{(l)} x^{(l)} + b^{(l)}) + \underline{b}_i^{\prime(l)} \leq z_i^{(l+1)} \leq \bar{A}_{i,:}^{(l)} (\mathbf{W}^{(l)} x^{(l)} + b^{(l)}) + \bar{b}_i^{\prime(l)}$$

Applying the bounds on  $x^{(l)}$  with respect to  $z^{(l-1)}$ , and using a similar technique as we did above to obtain (19), we get linear upper and lower bounds for  $z_i^{(l+1)}$  with respect to  $z^{(l-1)}$  in the following form:

$$\underline{A}_{i,:}^{(l-1)} z^{(l-1)} + \underline{b}_i^{\prime(l-1)} \leq z_i^{(l+1)} \leq \bar{A}_{i,:}^{(l-1)} z^{(l-1)} + \bar{b}_i^{\prime(l-1)} \quad (20)$$

where  $\underline{b}'_i^{(l-1)}$  and  $\bar{b}'_i^{(l-1)}$  collect all bias terms in the substitution process. Caution has to be taken when forming  $\underline{A}_{i,:}^{(l-1)}$  and  $\bar{A}_{i,:}^{(l-1)}$ , as we need to choose  $\underline{a}_k^{(l-1)}$  or  $\bar{a}_k^{(l-1)}$  based on the sign of  $\underline{A}_{i,:}^{(l)} \mathbf{W}_{:,k}^{(l)}$ , since the coefficients before each inequality now become  $\underline{A}_{i,:}^{(l)} \mathbf{W}^{(l)}$  rather than just  $\mathbf{W}^{(l)}$ :

$$\underline{A}_{i,k}^{(l-1)} = \begin{cases} \underline{A}_{i,:}^{(l)} \mathbf{W}_{:,k}^{(l)} \bar{a}_k^{(l-1)}, & \underline{A}_{i,:}^{(l)} \mathbf{W}_{:,k}^{(l)} < 0 \\ \underline{A}_{i,:}^{(l)} \mathbf{W}_{:,k}^{(l)} \underline{a}_k^{(l-1)}, & \underline{A}_{i,:}^{(l)} \mathbf{W}_{:,k}^{(l)} \geq 0 \end{cases}$$

An eagle-eyed reader can notice that we can continue this process until we have reached  $z^{(0)}$ , and obtain the following linear bounds:

$$\underline{A}_{i,:}^{(0)} z^{(0)} + \underline{b}'_i^{(0)} \leq z_i^{(l+1)} \leq \bar{A}_{i,:}^{(0)} z^{(0)} + \bar{b}'_i^{(0)} \quad (21)$$

where  $\underline{A}_{i,:}^{(0)}$ ,  $\bar{A}_{i,:}^{(0)}$ ,  $\underline{b}'_i^{(0)}$  and  $\bar{b}'_i^{(0)}$  can be formed similarly as above. Substituting  $z^{(0)} = \mathbf{W}^{(0)} x^{(0)} + b^{(0)}$  ( $x^{(0)} = x$  is the input of the neural network) simply yields:

$$\underline{A}_{i,:} x + \underline{b}'_i \leq z_i^{(l+1)} \leq \bar{A}_{i,:} x + \bar{b}'_i \quad (22)$$

where  $\underline{A}_{i,:} = \underline{A}_{i,:}^{(0)} \mathbf{W}^{(0)}$ ,  $\bar{A}_{i,:} = \bar{A}_{i,:}^{(0)} \mathbf{W}^{(0)}$  captures the products of  $\mathbf{W}$  of all layers and the chosen  $\underline{a}_k^{(l)}$  or  $\bar{a}_k^{(l)}$  for each layer;  $\underline{b}'_i, \bar{b}'_i$  collects all bias terms (we refer the readers to Theorem 3.2 in [Zhang et al. \(2018\)](#) for the exact form of  $\underline{A}, \bar{A}, \underline{b}', \bar{b}'$ ). This procedure beautifully works as the linear combination of linear bounds are still linear bounds. Eq. (22) is a remarkable result, as the output of a non-linear function (neural network) has been directly bounded linearly for all  $x$  close to  $x^{\text{nom}}$ . This allows us to immediately give upper and lower bounds of  $z_i^{(l+1)}$  by considering the worst case  $x \in \mathcal{S}_{in}(x^{\text{nom}})$ . When the set is an  $L_\infty$  normed ball, this is obvious,

$$- \epsilon \|\underline{A}_{i,:}\|_1 + \underline{A}_{i,:} x^{\text{nom}} + \underline{b}'_i \leq z_i^{(l+1)} \leq \epsilon \|\bar{A}_{i,:}\|_1 + \bar{A}_{i,:} x^{\text{nom}} + \bar{b}'_i, \quad (23)$$

The entire bound propagation process does not involve any LP solver, so it is efficient and can scale to quite large networks. The final objective  $c^\top x^{(L)} + c_0$  can be treated as an additional linear layer after  $z^{(L-1)}$ . Because to form the bounds for  $z^{(L-1)}$  we need to compute bounds for all  $z^{(l)}$ ,  $l \in [L-1]$  beforehand, each in  $O(l)$  time, the time complexity of this method is quadratic in  $L$ .

For each neuron, the selection of linear bounds are completely independent; this allows further improvements in this greedy algorithm. For example, the selection of  $\underline{a}_k^{(l)}$  can depend on  $\bar{z}_k^{(l)}$  and  $\underline{z}_k^{(l)}$  to adaptively minimize the error between the lower bound and ReLU function. CROWN ([Zhang et al., 2018](#)) and DeepPoly ([Singh et al., 2019b](#)) used this strategy to achieve tighter verification results than Fast-Lin ([Wang et al., 2018b](#)), DeepZ ([Singh et al., 2018](#)) and Neurify ([Wang et al., 2018b](#)). Note that although the bound propagation techniques used in these works can be viewed as using different linear relaxations and solve the primal problem greedily in our framework, each of the works has some unique features. For example, DeepPoly ([Singh et al., 2019b](#)) and DeepZ ([Singh et al., 2018](#)) carefully consider floating-point rounding during the computation; [Weng et al. \(2018\)](#) gives a theoretical hardness proof based on a reduction from the set-cover problem; Neurify ([Wang et al., 2018b](#)) combines the relaxed bound with a branch-and-bound search to give concrete instances of adversarial example if they exist, and also uses the bound for training ([Wang et al., 2018a](#)).

One the other hand, instead of propagating the bounds of  $z^{(l+1)}$  to  $z^{(l-k)}$  as shown above, we can decouple layer  $z^{(l-(k-1))}$  and  $z^{(l-k)}$  entirely: suppose we have obtained concrete upper and lower bounds for  $z^{(l-(k-1))}$ , we can treat  $z^{(l-(k-1))}$  as the input layer and only consider a  $k$ -layer network to compute the bounds of  $z^{(l+1)}$ . This leads to interval bounds propagation (IBP) ([Gowal et al., 2018](#)) ( $k = 1$ ) and ‘‘Hybrid Zonotope Domain’’ ([Mirman et al., 2018](#)) (general  $k$ ) which gives even looser bounds, but its computation cost is also greatly reduced to  $O(Lk)$ .

It is worth noting that the greedy algorithm in primal space is also closely connected to the greedy algorithm in dual space; the dual of (17) will recover a dual formulation with solution (49), and the closed form solution are related to the chosen slopes  $\bar{a}_i^{(l)}$  and  $\underline{a}_i^{(l)}$ . This explains the equivalence of Fast-Lin and the greedy algorithm to solve the dual problem presented in Algorithm 1 of [Wong & Kolter \(2018\)](#).

### A.3 Strong duality for (C): $p_C^* = d_C^*$

Consider the following perturbed version of problem (C):

$$\begin{aligned} \tilde{p}_C^* := \min_{(x^{[L+1]}, z^{[L]}) \in \mathcal{D}} \quad & c^\top x^{(L)} + c_0 \\ \text{s.t.} \quad & z^{(l)} = \mathbf{W}^{(l)} x^{(l)} + b^{(l)} + v^{(l)}, l \in [L], \\ & \underline{\sigma}^{(l)}(z^{(l)}) - \underline{u}^{(l)} \leq x^{(l+1)} \leq \bar{\sigma}^{(l)}(z^{(l)}) + \bar{u}^{(l)}, l \in [L]. \end{aligned} \quad (24)$$

In the following, we prove that the optimal value of the perturbed problem, i.e.,  $\tilde{p}_C^*$ , ‘‘smoothly’’ changes with the perturbations  $(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]})$ . Combined with convexity, this ensures the strong duality for problem (C).

To achieve this, we define  $X^{(0)} = \tilde{X}^{(0)} = \mathcal{S}_{in}(x^{\text{nom}})$ , and we also define  $Z^{(l)}$ ,  $\tilde{Z}^{(l)}$ ,  $X^{(l)}$  and  $\tilde{X}^{(l)}$  recursively as follows:

$$\begin{aligned} Z^{(l)} &= \{\mathbf{W}^{(l)} x^{(l)} + b^{(l)} : x^{(l)} \in X^{(l)}\} \cap [\underline{z}^{(l)}, \bar{z}^{(l)}], \\ \tilde{Z}^{(l)} &= \{\mathbf{W}^{(l)} \tilde{x}^{(l)} + b^{(l)} + v^{(l)} : \tilde{x}^{(l)} \in \tilde{X}^{(l)}\} \cap [\underline{z}^{(l)}, \bar{z}^{(l)}], \\ X^{(l+1)} &= \{x^{(l+1)} : \underline{\sigma}^{(l)}(z^{(l)}) \leq x^{(l+1)} \leq \bar{\sigma}^{(l)}(z^{(l)}), \\ &\quad z^{(l)} \in Z^{(l)}\}, \\ \tilde{X}^{(l+1)} &= \{\tilde{x}^{(l+1)} : \underline{\sigma}^{(l)}(\tilde{z}^{(l)}) - \underline{u}^{(l)} \leq \tilde{x}^{(l+1)} \\ &\quad \leq \bar{\sigma}^{(l)}(\tilde{z}^{(l)}) + \bar{u}^{(l)}, \tilde{z}^{(l)} \in \tilde{Z}^{(l)}\}. \end{aligned}$$

Intuitively,  $Z^{(l)}$ ,  $\tilde{Z}^{(l)}$ ,  $X^{(l)}$  and  $\tilde{X}^{(l)}$  are the set of activations that are achievable by the neural network given  $x^{(0)} \in \mathcal{S}_{in}(x^{\text{nom}})$  and  $z^{(l)} \in [\underline{z}^{(l)}, \bar{z}^{(l)}]$ .

**Lemma A.1.** *We assume that for each  $l \in [L]$ , both  $\underline{\sigma}^{(l)}$  and  $\bar{\sigma}^{(l)}$  have a finite Lipschitz constant in the domain  $[\underline{z}^{(l)}, \bar{z}^{(l)}]$ . Then for every  $l \in [L+1]$ , there exist positive constants  $C_x^{(l)}$  and  $C_z^{(l)}$  such that*

$$\sup_{\tilde{x}^{(l)} \in \tilde{X}^{(l)}} \text{dist}(\tilde{x}^{(l)}, X^{(l)}) \leq C_x^{(l)} \|(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]})\|_2, \quad (25)$$

$$\sup_{\tilde{z}^{(l)} \in \tilde{Z}^{(l)}} \text{dist}(\tilde{z}^{(l)}, Z^{(l)}) \leq C_z^{(l)} \|(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L+1]})\|_2, \quad (26)$$

where  $\text{dist}(s, \mathcal{S}) := \inf_{s' \in \mathcal{S}} \|s - s'\|_2$ .

There also exists a positive constant  $C_C$  such that

$$p_C^* - \tilde{p}_C^* \leq C_C \|(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]})\|_2 \quad (27)$$

*Proof.* When problem (24) is infeasible, i.e.,  $\tilde{X}^{(L)} = \emptyset$ ,  $\tilde{p}_C^* = +\infty$  and (27) naturally holds true. In the following, we prove (25), (26) and (27) when problem (24) is feasible. In this case, all  $Z^{(l)}$ ,  $\tilde{Z}^{(l)}$ ,  $X^{(l)}$  and  $\tilde{X}^{(l)}$  are non-empty sets.

Since  $X^{(l)} = \tilde{X}^{(l)} = \mathcal{S}_{in}(x^{\text{nom}})$ , we have that (25) holds true for  $l = 0$  with  $C_x^{(0)} = 0$ . In the following, we use mathematical induction to prove (26) for  $0 \leq l \leq L-1$  and (25) for  $1 \leq l \leq L$ .

First, suppose  $\text{dist}(\tilde{x}^{(l)}, X^{(l)}) \leq C_x^{(l)} \|(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]})\|_2$  holds true for any  $\tilde{x}^{(l)} \in \tilde{X}^{(l)}$ . Then for any  $\tilde{z}^{(l)} = \mathbf{W}^{(l)} \tilde{x}^{(l)} + b^{(l)} + v^{(l)} \in \tilde{Z}^{(l)}$ , we have

$$\begin{aligned} \text{dist}(\tilde{z}^{(l)}, Z^{(l)}) &= \inf_{z^{(l)} \in Z^{(l)}} \|\tilde{z}^{(l)} - z^{(l)}\| \\ &\leq \inf_{x^{(l)} \in X^{(l)}} \|\mathbf{W}^{(l)}(\tilde{x}^{(l)} - x^{(l)}) + v^{(l)}\| \\ &\leq \inf_{x^{(l)} \in X^{(l)}} \|\mathbf{W}^{(l)}\| \|\tilde{x}^{(l)} - x^{(l)}\| + \|v^{(l)}\| \\ &= \|\mathbf{W}^{(l)}\| \text{dist}(\tilde{x}^{(l)}, X^{(l)}) + \|v^{(l)}\| \\ &\leq \|\mathbf{W}^{(l)}\| C_x^{(l)} \|(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]})\| + \|v^{(l)}\| \\ &\leq \left( (C_x^{(l)})^2 \|\mathbf{W}^{(l)}\|^2 + 1 \right)^2 \|(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L+1]})\|. \end{aligned}$$

Therefore, (26) holds true with  $C_z^{(l)} = \left( (C_x^{(l)})^2 \|\mathbf{W}^{(l)}\|^2 + 1 \right)^2$ .

Then by definition, for any  $\tilde{x}^{(l+1)} \in \tilde{X}^{(l+1)}$ , there exists  $\tilde{z}^{(l)} \in \tilde{Z}^{(l)}$  such that

$$\underline{\sigma}^{(l)}(\tilde{z}^{(l)}) - \underline{u}^{(l)} \leq \tilde{x}^{(l+1)} \leq \bar{\sigma}^{(l)}(\tilde{z}^{(l)}) + \bar{u}^{(l)}.$$

By the induction assumption, there exists  $z^{(l)} \in Z^{(l)}$  such that

$$\text{dist}(\tilde{z}^{(l)}, z^{(l)}) \leq C_z^{(l)} \|(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L+1]})\|_2.$$

Thus, we have

$$\begin{aligned} \text{dist}(\tilde{x}^{(l+1)}, X^{(l+1)}) &= \inf_{x^{(l+1)} \in X^{(l+1)}} \|\tilde{x}^{(l+1)} - x^{(l+1)}\| \\ &\leq \inf \{ \|\tilde{x}^{(l+1)} - x^{(l+1)}\| : \underline{\sigma}^{(l)}(z^{(l)}) \leq x^{(l+1)} \leq \bar{\sigma}^{(l)}(z^{(l)}) \}. \end{aligned}$$

We re-parametrize  $\tilde{x}^{(l+1)}$  and  $x^{(l+1)}$  as

$$\begin{aligned} \tilde{x}^{(l+1)} &= \underline{\sigma}^{(l)}(\tilde{z}^{(l)}) + \tilde{t} \\ x^{(l+1)} &= \underline{\sigma}^{(l)}(z^{(l)}) + t, \end{aligned}$$

where

$$\begin{aligned} -\underline{u}^{(l)} \leq \tilde{t} &\leq \Delta\sigma^{(l)}(\tilde{z}^{(l)}) + \bar{u}^{(l)}, \\ 0 \leq t &\leq \Delta\sigma^{(l)}(z^{(l)}), \end{aligned}$$

and

$$\Delta\sigma^{(l)}(\tilde{z}^{(l)}) = \bar{\sigma}^{(l)}(z^{(l)}) - \underline{\sigma}^{(l)}(z^{(l)}).$$

It is easy to prove that if  $\underline{\sigma}^{(l)}$  and  $\bar{\sigma}^{(l)}$  have Lipschitz constant  $\underline{L}^{(l)}$  and  $\bar{L}^{(l)}$  respectively,  $\Delta\sigma^{(l)}$  has a Lipschitz constant  $\underline{L}^{(l)} + \bar{L}^{(l)}$ . Then we have

$$\begin{aligned} \text{dist}(\tilde{x}^{(l+1)}, X^{(l+1)}) &\leq \|\underline{\sigma}^{(l)}(\tilde{z}^{(l)}) - \underline{\sigma}^{(l)}(z^{(l)})\| + \inf_{t \in [0, \Delta\sigma^{(l)}(z^{(l)})]} \|\tilde{t} - t\| \\ &\leq \underline{L}^{(l)} \|\tilde{z}^{(l)} - z^{(l)}\| + \left( \sum_k \inf_{t_k \in [0, \Delta\sigma_k^{(l)}(z^{(l)})]} |\tilde{t}_k - t_k|^2 \right)^{1/2}. \end{aligned}$$

We have the entry-wise bound for  $\tilde{t} - t$ :

$$\begin{aligned} \inf_{t_k \in [0, \Delta\sigma_k^{(l)}(z^{(l)})]} |\tilde{t}_k - t_k|^2 &\leq \max(|\underline{u}_k^{(l)}|^2, |\Delta\sigma_k^{(l)}(\tilde{z}^{(l)}) - \Delta\sigma_k^{(l)}(z^{(l)}) + \bar{u}^{(l)}|^2) \\ &\leq 2 \left( |\Delta\sigma_k^{(l)}(\tilde{z}^{(l)}) - \Delta\sigma_k^{(l)}(z^{(l)})|^2 + |\underline{u}_k^{(l)}|^2 + |\bar{u}^{(l)}|^2 \right) \end{aligned}$$

Therefore, we get

$$\begin{aligned} \inf_{t \in [0, \Delta\sigma^{(l)}(z^{(l)})]} \|\tilde{t} - t\| &\leq \sqrt{2} \left( \|\Delta\sigma^{(l)}(\tilde{z}^{(l)}) - \Delta\sigma^{(l)}(z^{(l)})\|^2 + \|\underline{u}_k^{(l)}\|^2 + \|\bar{u}^{(l)}\|^2 \right)^{1/2} \\ &\leq \sqrt{2} \left( (\underline{L}^{(l)} + \bar{L}^{(l)})^2 \|\tilde{z}^{(l)} - z^{(l)}\|^2 + \|\underline{u}_k^{(l)}\|^2 + \|\bar{u}^{(l)}\|^2 \right)^{1/2} \\ &\leq \sqrt{2(\underline{L}^{(l)} + \bar{L}^{(l)})^2 (C_z^{(l)})^2 + 2} \|(\underline{u}^{[L+1]}, \bar{u}^{[L+1]}, v^{[L+1]})\|_2 \end{aligned}$$

Similarly, we have  $\underline{L}^{(l)} \|\tilde{z}^{(l)} - z^{(l)}\| \leq \underline{L}^{(l)} C_z^{(l)} \|(\underline{u}^{[L+1]}, \bar{u}^{[L+1]}, v^{[L+1]})\|_2$ . Therefore, we obtain

$$\text{dist}(\tilde{x}^{(l+1)}, X^{(l+1)}) \leq C_x^{(l+1)} \|(\underline{u}^{[L+1]}, \bar{u}^{[L+1]}, v^{[L+1]})\|_2,$$

where  $C_x^{(l+1)} = \underline{L}^{(l)} C_z^{(l)} + \sqrt{2(\underline{L}^{(l)} + \bar{L}^{(l)})^2 (C_z^{(l)})^2 + 2}$ .

Then by mathematical induction, we proved that (26) for  $0 \leq l \leq L-1$  and (25) for  $1 \leq l \leq L$ .

Now for any  $\tilde{x}^{(l)} \in \tilde{X}^{(l)}$ , there exists  $x^{(L)} \in X^{(L)}$  such that

$$\text{dist}(\tilde{x}^{(L)}, x^{(L)}) \leq C_x^{(L)} \|(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]})\|_2.$$

Then we have

$$\begin{aligned} p_C^* - (c^\top \tilde{x}^{(L)} + c_0) &\leq c^\top (x^{(L)} - \tilde{x}^{(L)}) \\ &\leq \|c\| \|\tilde{x}^{(L)} - x^{(L)}\| \\ &\leq C_x^{(L)} \|c\| \|(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]})\|_2. \end{aligned}$$

Taking the supremum over  $\tilde{x}^{(l)} \in \tilde{X}^{(l)}$ , we have proved (27) with  $C_C = C_x^{(L)} \|c\|$ .  $\square$

*Proof of Theorem 4.1.* The structure of the proof follows the proof of strong duality given the Slater's condition in (Boyd & Vandenberghe, 2004) (Section 5.3.2). However, we do not assume the Slater's condition in our result here. Let's define

$$\begin{aligned} \mathcal{A} = \{ &(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]}, t) : \exists (x^{[L+1]}, z^{[L]}) \in \mathcal{D}, \\ &\underline{\sigma}^{(l)}(z^{(l)}) - \underline{u}^{(l)} \leq x^{(l+1)} \leq \bar{\sigma}^{(l)}(z^{(l)}) + \bar{u}^{(l)}, \\ &z^{(l)} = \mathbf{W}^{(l)} x^{(l)} + b^{(l)} + v^{(l)}, \forall l \in [L], \\ &c^\top x^{(L)} + c_0 \leq t\}, \end{aligned}$$

and

$$\mathcal{B} = \{(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]}, t) : t < p_C^* - C_C \|(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]})\|_2\}.$$

$\mathcal{A}$  is convex because the problem (24) is convex.  $\mathcal{B}$  is convex by definition. The sets  $\mathcal{A}$  and  $\mathcal{B}$  do not intersect. To see this, suppose  $(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]}, t) \in \mathcal{A} \cap \mathcal{B}$ . Since  $(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]}, t) \in \mathcal{B}$ , we have  $t < p_C^* - C_C \|(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]})\|_2$ . Since  $(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]}, t) \in \mathcal{A}$ , there exists  $(x^{[L+1]}, z^{[L]}) \in \mathcal{D}$  such that it satisfies the constraints in problem (24), and  $t \geq c^\top x^{(L)} + c_0 \geq \tilde{p}_C^* \geq p_C^* - C_C \|(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]})\|_2$ , which is impossible since (27) bounds the optimal value of the perturbed problem (24).

By the separating hyperplane theorem, there exists  $(\underline{\lambda}^{[L]}, \bar{\lambda}^{[L]}, \mu^{[L]}, \nu) \neq 0$  and  $\alpha$  such that

$$(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]}, t) \in \mathcal{A} \Rightarrow \underline{\lambda}^{[L]\top} \underline{u}^{[L]} + \bar{\lambda}^{[L]\top} \bar{u}^{[L]} + \mu^{[L]\top} v^{[L]} + \nu t \geq \alpha, \quad (28)$$

and

$$(\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]}, t) \in \mathcal{B} \Rightarrow \underline{\lambda}^{[L]\top} \underline{u}^{[L]} + \bar{\lambda}^{[L]\top} \bar{u}^{[L]} + \mu^{[L]\top} v^{[L]} + \nu t \leq \alpha, \quad (29)$$

From (28), we conclude that  $\underline{\lambda}^{[L]} \geq 0$ ,  $\bar{\lambda}^{[L]} \geq 0$  and  $\nu \geq 0$ . Otherwise,  $\underline{\lambda}^{[L]\top} \underline{u}^{[L]} + \bar{\lambda}^{[L]\top} \bar{u}^{[L]} + \nu t$  is unbounded from below over  $\mathcal{A}$ , contradicting (28). Since  $(0, 0, 0, t) \in \mathcal{B}$  for any  $t > p_C^*$ , we have  $\nu t \leq \alpha$  for any  $t > p_C^*$  thanks to (29), and thus  $\nu p_C^* \leq \alpha$ . Together with (28), we conclude that for any  $(x^{[L+1]}, z^{[L]}) \in \mathcal{D}$ ,

$$\begin{aligned} \nu(c^\top x^{(L)} + c_0) + \sum_{l=0}^{L-1} \mu^{(l)\top} (z^{(l)} - \mathbf{W}^{(l)} x^{(l)} - b^{(l)}) \\ + \sum_{l=0}^{L-1} \underline{\lambda}^{(l)\top} (\underline{\sigma}^{(l)}(z^{(l)}) - x^{(l+1)}) \\ + \sum_{l=0}^{L-1} \bar{\lambda}^{(l)\top} (x^{(l+1)} - \bar{\sigma}^{(l)}(z^{(l)})) \\ \geq \alpha \geq \nu p_C^*. \end{aligned} \quad (30)$$

Assume that  $\nu > 0$ . In that case, we can divide (30) by  $\nu$  to obtain

$$L(x^{[L+1]}, z^{[L]}, \underline{\lambda}^{[L]}/\nu, \bar{\lambda}^{[L]}/\nu, \mu^{[L]}/\nu) \geq p_C^*$$

for all  $(x^{[L+1]}, z^{[L]}) \in \mathcal{D}$ , where  $L(\cdot)$ , defined in (8), is the Lagrangian of (C). Maximizing over  $(x^{[L+1]}, z^{[L]}) \in \mathcal{D}$ , we obtain  $g_C(\mu^{[L]}/\nu, \underline{\lambda}^{[L]}/\nu, \bar{\lambda}^{[L]}/\nu) \geq p_C^*$ . By weak duality, we have  $g_C(\mu^{[L]}/\nu, \underline{\lambda}^{[L]}/\nu, \bar{\lambda}^{[L]}/\nu) \leq p_C^*$ , so in fact  $g_C(\mu^{[L]}/\nu, \underline{\lambda}^{[L]}/\nu, \bar{\lambda}^{[L]}/\nu) = p_C^*$ . This shows that strong duality holds, and that the dual optimum is attained, at least in the case when  $\nu > 0$ .

Now we consider the case  $\nu = 0$ . From (30), we conclude that for any  $(x^{[L+1]}, z^{[L]}) \in \mathcal{D}$ ,

$$\begin{aligned} & \sum_{l=0}^{L-1} \mu^{(l)\top} (z^{(l)} - \mathbf{W}^{(l)} x^{(l)} - b^{(l)}) \\ & + \sum_{l=0}^{L-1} \underline{\lambda}^{(l)\top} (\underline{\sigma}^{(l)}(z^{(l)}) - x^{(l+1)}) \\ & + \sum_{l=0}^{L-1} \bar{\lambda}^{(l)\top} (x^{(l+1)} - \bar{\sigma}^{(l)}(z^{(l)})) \geq \alpha \geq 0. \end{aligned} \quad (31)$$

Taking any feasible point of problem (C), i.e.,  $(x^{[L+1]}, z^{[L]}) \in \mathcal{S}_C$  and combining with  $\underline{\lambda}^{[L]} \geq 0, \bar{\lambda}^{[L]} \geq 0$ , we know that the left-hand-side of (31) is non-positive, and thus  $\alpha = 0$ . Then from (29), we conclude that for any  $t \in \mathbb{R}$

$$\begin{aligned} \|\underline{u}^{[L]}, \bar{u}^{[L]}, v^{[L]}\|_2 & < \frac{p_C^* - t}{C_C} \Rightarrow \\ \underline{\lambda}^{[L]\top} \underline{u}^{[L]} + \bar{\lambda}^{[L]\top} \bar{u}^{[L]} + \mu^{[L]\top} v^{[L]} & \leq 0, \end{aligned}$$

which can only be possible when  $(\underline{\lambda}^{[L]}, \bar{\lambda}^{[L]}, \mu^{[L]}) = 0$ . Combined with  $\nu = 0$ , this contradicts with  $(\underline{\lambda}^{[L]}, \bar{\lambda}^{[L]}, \mu^{[L]}, \nu) \neq 0$ , and thus  $\nu$  cannot be 0.  $\square$

#### A.4 Equivalence of the dual problems: $d_C^* = d_{\mathcal{O}}^*$

**Lemma A.2.** *Suppose the activation function  $\sigma : [\underline{z}, \bar{z}] \rightarrow \mathbb{R}$  is bounded from above and below and that  $\underline{\sigma}(z) \leq \sigma(z) \leq \bar{\sigma}(z)$  for all  $z \in [\underline{z}, \bar{z}]$ . Define*

$$f_{\mathcal{O}}(\mu, \lambda) := \inf_{z \in [\underline{z}, \bar{z}]} \mu z - \lambda \sigma(z), \quad (32)$$

$$f_C(\mu, \underline{\lambda}, \bar{\lambda}) := \inf_{z \in [\underline{z}, \bar{z}]} \mu z + \underline{\lambda} \sigma(z) - \bar{\lambda} \bar{\sigma}(z). \quad (33)$$

For any  $\mu, \underline{\lambda} \geq 0$  and  $\bar{\lambda} \geq 0$ , we have

$$f_C(\mu, \underline{\lambda}, \bar{\lambda}) \leq f_C(\mu, -(\bar{\lambda} - \underline{\lambda})_-, (\bar{\lambda} - \underline{\lambda})_+), \quad (34)$$

where  $\lambda_+ = \max(\lambda, 0)$  and  $\lambda_- = \min(\lambda, 0)$ .

If  $\underline{\sigma}$  and  $\bar{\sigma}$  are the optimal convex relaxations defined in (15) and (16), we have that for any  $\mu$  and  $\lambda$

$$f_C(\mu, -\lambda_-, \lambda_+) = f_{\mathcal{O}}(\mu, \lambda). \quad (35)$$

*Proof.* First let's prove (34). For  $\underline{\lambda} \geq \bar{\lambda} \geq 0$ , we have

$$f_C(\mu, -(\bar{\lambda} - \underline{\lambda})_-, (\bar{\lambda} - \underline{\lambda})_+) = f_C(\mu, \underline{\lambda} - \bar{\lambda}, 0)$$

and

$$\begin{aligned} f_C(\mu, \underline{\lambda}, \bar{\lambda}) & = \inf_{z \in [\underline{z}, \bar{z}]} \mu z + \underline{\lambda} \sigma(z) - \bar{\lambda} \bar{\sigma}(z) \\ & = \inf_{z \in [\underline{z}, \bar{z}]} \mu z + (\underline{\lambda} - \bar{\lambda}) \sigma(z) - \bar{\lambda} (\bar{\sigma}(z) - \sigma(z)) \\ & \stackrel{(i)}{\leq} \sup_{z \in [\underline{z}, \bar{z}]} \mu z - (\underline{\lambda} - \bar{\lambda}) \sigma(z) = f_C(\mu, \underline{\lambda} - \bar{\lambda}, 0), \end{aligned}$$

where we use  $\bar{\lambda}(\bar{\sigma}(z) - \sigma(z)) \geq 0$  in (i). Similarly for  $\bar{\lambda} \geq \underline{\lambda} \geq 0$ , we have

$$f_C(\mu, -(\bar{\lambda} - \underline{\lambda})_-, (\bar{\lambda} - \underline{\lambda})_+) = f_C(\mu, 0, \bar{\lambda} - \underline{\lambda})$$

and

$$\begin{aligned} f_C(\mu, \underline{\lambda}, \bar{\lambda}) & = \inf_{z \in [\underline{z}, \bar{z}]} \mu z + \underline{\lambda} \sigma(z) - \bar{\lambda} \bar{\sigma}(z) \\ & = \sup_{z \in [\underline{z}, \bar{z}]} \mu z - (\bar{\lambda} - \underline{\lambda}) \bar{\sigma}(z) - \underline{\lambda} (\bar{\sigma}(z) - \sigma(z)) \\ & \stackrel{(i)}{\leq} \sup_{z \in [\underline{z}, \bar{z}]} \mu z - (\bar{\lambda} - \underline{\lambda}) \bar{\sigma}(z) = f_C(\mu, 0, \bar{\lambda} - \underline{\lambda}), \end{aligned}$$

where we use  $\underline{\lambda}(\bar{\sigma}(z) - \underline{\sigma}(z)) \geq 0$  in (i).

Then let's prove (35). For  $\lambda < 0$  ( $\lambda_+ = 0$  and  $\lambda_- = \lambda$ ), by definition we have

$$\begin{aligned} f_C(\mu, -\lambda, 0) &= \inf_{z \in [\underline{z}, \bar{z}]} \mu z - \lambda \underline{\sigma}(z) \\ &= \lambda \sup_{z \in [\underline{z}, \bar{z}]} \frac{\mu}{\lambda} z - \underline{\sigma}(z) \stackrel{(i)}{=} \lambda (\underline{\sigma})^* (\mu/\lambda) \stackrel{(ii)}{=} \lambda (\sigma)^* (\mu/\lambda) \\ &\stackrel{(iii)}{=} \inf_{z \in [\underline{z}, \bar{z}]} \mu z - \lambda \sigma(z) = f_{\mathcal{O}}(\mu, \lambda), \end{aligned}$$

where we use the definition of convex conjugate in (i) and (iii) and the Fenchel-Moreau theorem (Theorem 12.2 in Rockafellar (2015)) in (ii). For  $\lambda = 0$ , it is obvious. Similarly, for  $\lambda > 0$  ( $\lambda_+ = \lambda$  and  $\lambda_- = 0$ ), by definition we have

$$\begin{aligned} f_C(\mu, 0, \lambda) &= \inf_{z \in [\underline{z}, \bar{z}]} \mu z - \lambda \bar{\sigma}(z) \\ &= -\lambda \sup_{z \in [\underline{z}, \bar{z}]} -\frac{\mu}{\lambda} z - (-\bar{\sigma})(z) \stackrel{(i)}{=} -\lambda (-\bar{\sigma})^* (-\mu/\lambda) \\ &\stackrel{(ii)}{=} -\lambda (-\sigma)^* (-\mu/\lambda) \stackrel{(iii)}{=} \inf_{z \in [\underline{z}, \bar{z}]} \mu z - \lambda \sigma(z) = f_{\mathcal{O}}(\mu, \lambda), \end{aligned}$$

where we use the definition of convex conjugate in (i) and (iii) and the Fenchel-Moreau theorem in (ii), again.  $\square$

Now, let's prove Theorem 4.2.

*Proof of Theorem 4.2.* In the following, we prove (12) and (14) in sequel.

Let's first simplify the form of  $g_C(\mu^{[L]}, \underline{\lambda}^{[L]}, \bar{\lambda}^{[L]})$ . By definition (8), we have

$$\begin{aligned} g_C(\mu^{[L]}, \underline{\lambda}^{[L]}, \bar{\lambda}^{[L]}) &= g^{(0)}(\mu^{(0)}) \\ &\quad + \sum_{l=1}^{L-1} g^{(l)}(\mu^{(l)}, \bar{\lambda}^{(l-1)} - \underline{\lambda}^{(l-1)}) \\ &\quad + g^{(L)}(c, \bar{\lambda}^{(L-1)} - \underline{\lambda}^{(L-1)}) \\ &\quad + \sum_{l=0}^{L-1} \left( \tilde{g}_C^{(l)}(\mu^{(l)}, \underline{\lambda}^{(l)}, \bar{\lambda}^{(l)}) - b^{(l)\top} \mu^{(l)} \right), \end{aligned} \tag{36}$$

where

$$g^{(0)}(\mu^{(0)}) = \inf_{x^{(0)} \in \mathcal{S}_{\text{in}}(x^{\text{nom}})} \left( -\mathbf{W}^{(0)\top} \mu^{(0)} \right)^\top x^{(0)} \tag{37}$$

$$\begin{aligned} g^{(l)}(\mu^{(l)}, \lambda^{(l-1)}) &= \inf_{x^{(l)}} \left( \lambda^{(l-1)} - \mathbf{W}^{(l)\top} \mu^{(l)} \right)^\top x^{(l)} \\ &= \mathbf{1}_{\lambda^{(l-1)} = \mathbf{W}^{(l)\top} \mu^{(l)}}, \end{aligned} \tag{38}$$

$$\begin{aligned} g^{(L)}(c, \lambda^{(L-1)}) &= \inf_{x^{(L)}} \left( \lambda^{(L-1)} + c \right)^\top x^{(L)} + c_0 \\ &= \mathbf{1}_{\lambda^{(L-1)} = -c} + c_0, \end{aligned} \tag{39}$$

and

$$\tilde{g}_C^{(l)}(\mu^{(l)}, \underline{\lambda}^{(l)}, \bar{\lambda}^{(l)}) = \inf_{\underline{z}^{(l)} \leq z^{(l)} \leq \bar{z}^{(l)}} \left\{ \mu^{(l)\top} z^{(l)} + \underline{\lambda}^{(l)\top} \underline{\sigma}^{(l)}(z^{(l)}) - \bar{\lambda}^{(l)\top} \bar{\sigma}^{(l)}(z^{(l)}) \right\}. \tag{40}$$

Now, let's prove (12). For any  $\mu^{[L]}, \underline{\lambda}^{[L]} \geq 0$  and  $\bar{\lambda}^{[L]} \geq 0$ , we apply (34) in Lemma A.2 entry-wisely on (40), and obtain

$$\tilde{g}_C^{(l)}(\mu^{(l)}, -\lambda_-^{(l)}, \lambda_+^{(l)}) \leq \tilde{g}_C^{(l)}(\mu^{(l)}, \underline{\lambda}^{(l)}, \bar{\lambda}^{(l)}),$$

in which  $\lambda^{(l)} := \underline{\lambda}^{(l)} - \bar{\lambda}^{(l)}$ . After plugging  $\lambda^{(l)} = \lambda_+^{(l)} + \lambda_-^{(l)}$  into the left-hand-side and  $\lambda^{(l)} := \underline{\lambda}^{(l)} - \bar{\lambda}^{(l)}$  into the right-hand-side of (12), we see that the other three terms in  $g_C$  on the two sides are the same. Therefore, we have proved (12).

Similarly by definition (10) we have

$$\begin{aligned} g_{\mathcal{O}}(\mu^{[L]}, \lambda^{[L]}) &= g^{(0)}(\mu^{(0)}) + \sum_{l=1}^{L-1} g^{(l)}(\mu^{(l)}, \lambda^{(l-1)}) \\ &\quad + g^{(L)}(c, \lambda^{(L-1)}) + \sum_{l=0}^{L-1} \left( \tilde{g}_{\mathcal{O}}^{(l)}(\mu^{(l)}, \lambda^{(l)}) - b^{(l)\top} \mu^{(l)} \right) \end{aligned} \quad (41)$$

in which

$$\tilde{g}_{\mathcal{O}}^{(l)}(\mu^{(l)}, \lambda^{(l)}) = \inf_{\underline{z}^{(l)} \leq z^{(l)} \leq \bar{z}^{(l)}} \mu^{(l)\top} z^{(l)} - \lambda^{(l)\top} \sigma^{(l)}(z^{(l)}). \quad (42)$$

Now, let's prove (14). For any  $\mu^{[L]}$  and  $\lambda^{[L]}$ , since all the nonlinear layers are non-interactive, we apply (35) in Lemma A.2 entry-wisely on (40) and (42) and obtain

$$\tilde{g}_{\mathcal{O}}^{(l)}(\mu^{(l)}, -\lambda_-^{(l)}, \lambda_+^{(l)}) = \tilde{g}_{\mathcal{O}}^{(l)}(\mu^{(l)}, \lambda^{(l)}).$$

After plugging  $\lambda^{[L]} = \lambda_+^{[L]} + \lambda_-^{[L]}$  into  $g_C$ , we see that the other three terms in  $g_C$  and  $g_{\mathcal{O}}$  are the same. Therefore, we have proved (14).  $\square$

## A.5 Several useful results and a dual greedy algorithm

We provide the following useful results when solving (9) and (11). First, the dual problem (9) can be rewritten as an unconstrained optimization problem inspired by (13). We will define a two-argument function, reusing the name  $g_C$ , as

$$g_C(\mu^{[L]}, \lambda^{[L]}) := g_C(\mu^{[L]}, -\lambda_-^{[L]}, \lambda_+^{[L]}).$$

Then we have the following useful results.

**Proposition A.2.** Denote  $\lambda_+ = \max(\lambda, 0)$  and  $\lambda_- = \min(\lambda, 0)$ .

1. For dual of the convex relaxed problem (C) defined in (9), we have

$$d_C^* = \max_{\mu^{[L]}, \lambda^{[L]}} \left\{ g_C(\mu^{[L]}, \lambda^{[L]}) := c_0 + g^{(0)}(\mu^{(0)}) + \sum_{l=0}^{L-1} \left( \tilde{g}_C^{(l)}(\mu^{(l)}, \lambda^{(l)}) - b^{(l)\top} \mu^{(l)} \right) \right\}, \quad (43)$$

where

$$\begin{aligned} \lambda^{(L-1)} &= -c, \\ \lambda^{(l)} &= \mathbf{W}^{(l+1)\top} \mu^{(l+1)} \quad \forall l \in [L-1], \end{aligned} \quad (44)$$

$$g^{(0)}(\mu^{(0)}) = \inf_{x^{(0)} \in \mathcal{S}_{in}(x^{nom})} \left( -\mathbf{W}^{(0)\top} \mu^{(0)} \right)^\top x^{(0)}, \quad (45)$$

and

$$\tilde{g}_C^{(l)}(\mu^{(l)}, \lambda^{(l)}) = \inf_{\underline{z}^{(l)} \leq z^{(l)} \leq \bar{z}^{(l)}} \left\{ \mu^{(l)\top} z^{(l)} - \lambda_-^{(l)\top} \underline{\sigma}^{(l)}(z^{(l)}) - \lambda_+^{(l)\top} \bar{\sigma}^{(l)}(z^{(l)}) \right\}.$$

2. For the dual of the original nonconvex problem (O) defined in (11), we have

$$d_{\mathcal{O}}^* := \max_{\mu^{[L]}, \lambda^{[L]}} \left\{ g_{\mathcal{O}}(\mu^{[L]}, \lambda^{[L]}) = c_0 + g^{(0)}(\mu^{(0)}) + \sum_{l=0}^{L-1} \left( \tilde{g}_{\mathcal{O}}^{(l)}(\mu^{(l)}, \lambda^{(l)}) - b^{(l)\top} \mu^{(l)} \right) \right\}, \quad (46)$$

where (44) still holds true and

$$\tilde{g}_{\mathcal{O}}^{(l)}(\mu^{(l)}, \lambda^{(l)}) = \inf_{\underline{z}^{(l)} \leq z^{(l)} \leq \bar{z}^{(l)}} \mu^{(l)\top} z^{(l)} - \lambda^{(l)\top} \sigma^{(l)}(z^{(l)}).$$

3. Suppose that a nonlinear neuron  $x_j^{(l+1)} = \sigma^{(l)}(z_{I_j}^{(l)})$  is effectively linear within the input domain  $\mathcal{S}_{in}(x^{nom})$ , i.e., there exists a linear relation  $x_j^{(l+1)} = V_j^{(l)} z_{I_j}^{(l)} + d_j^{(l)}$  for all  $x^{(0)} \in \mathcal{S}_{in}(x^{nom})$ , then we can simplify the convex relaxed problem (C) by setting

$$\underline{\sigma}_i^{(l)}(z^{(l)}) = \bar{\sigma}_i^{(l)}(z^{(l)}) = V_j^{(l)} z_{I_j}^{(l)} + d_j^{(l)},$$

or simplify the original convex problem

$$\sigma_i^{(l)}(z^{(l)}) = V_j^{(l)} z_{I_j}^{(l)} + d_j^{(l)}.$$

If this neuron does not interact with other neurons in the same layer, i.e.,  $z_{I_j}^{(l)}$  is not the input of  $x_k^{(l+1)}$  for any  $k \neq j$ . Then for any optimal point for both dual problems, we have

$$\mu_{I_j}^{(l)} = V_j^{(l)\top} \lambda_j^{(l)}. \quad (47)$$

Similar results have been obtained in several previous works (Wong & Kolter, 2018; Dvijotham et al., 2018d; Wong et al., 2018; Qin et al., 2019).

*Proof.*

1. (43) is a straightforward rewriting of (36) with (38), (39), (37) and (40).
2. (46) is a straightforward rewriting of (41) with (38), (39), (37) and (42).
3. This can be proved with the same treatment of linear layers in the two items above.

□

**Greedly solving the dual with linear bounds.** Suppose the relaxed bounds  $\underline{\sigma}$  and  $\bar{\sigma}$  are linear, i.e.,

$$\underline{\sigma}^{(l)}(z^{(l)}) := \underline{a}^{(l)} z^{(l)} + \underline{b}^{(l)}, \quad \bar{\sigma}^{(l)}(z^{(l)}) := \bar{a}^{(l)} z^{(l)} + \bar{b}^{(l)}.$$

In this case, in the dual problem (43) we have

$$d_{\mathcal{C}}^* = \max_{\mu^{[L]}, \lambda^{[L]}} \left\{ g_{\mathcal{C}}(\mu^{[L]}, \lambda^{[L]}) := c_0 + g^{(0)}(\mu^{(0)}) + \sum_{l=0}^{L-1} \left( \tilde{g}_{\mathcal{C}}^{(l)}(\mu^{(l)}, \lambda^{(l)}) - b^{(l)\top} \mu^{(l)} \right) \right\},$$

where

$$\tilde{g}_{\mathcal{C}}^{(l)}(\mu^{(l)}, \lambda^{(l)}) = \inf_{\underline{z}^{(l)} \leq z^{(l)} \leq \bar{z}^{(l)}} \left\{ \left( \mu^{(l)} - \lambda_+^{(l)} \bar{a}^{(l)} - \lambda_-^{(l)} \underline{a}^{(l)} \right) z^{(l)} + \left( \lambda_+^{(l)} \bar{b}^{(l)} - \lambda_-^{(l)} \underline{b}^{(l)} \right) \right\}.$$

In the following, we propose a dual greedy algorithm to *greedly* (approximately) solve the dual problem (9) and/or its simplified version (43). Let  $\lambda^{[L]}$  be determined by (44) and  $\mu^{[L]}$ , for stable neurons, be determined by (47). Both of these are optimal. For *unstable* neurons ( $\underline{z}_i^{(l)} \leq 0 \leq \bar{z}_i^{(l)}$ ), a suboptimal  $\mu^{[L]}$  can be obtained by

$$\mu^{(l)} = \arg \max_{\mu^{(l)}} \tilde{g}_{\mathcal{C}}^{(l)}(\mu^{(l)}, \lambda^{(l)}), \quad (48)$$

which has a closed form solution

$$\mu_i^{(l)} = \bar{a}_i^{(l)} \left( \lambda_i^{(l)} \right)_+ + \underline{a}_i^{(l)} \left( \lambda_i^{(l)} \right)_-. \quad (49)$$

The corresponding lower bound is

$$g_{\mathcal{C}}(\mu^{[L]}, \lambda^{[L]}) = c_0 + g^{(0)}(\mu^{(0)}) - \sum_{l=0}^{L-1} b^{(l)\top} \mu^{(l)} + \sum_{\text{unstable}} \left( \left( \lambda_i^{(l)} \right)_+ \bar{b}^{(l)} - \left( \lambda_i^{(l)} \right)_- \underline{b}^{(l)} \right).$$

We point out that the algorithm above is exactly that proposed in Theorem 1 in Wong & Kolter (2018). Their  $\nu$  is our  $\mu$  and their  $\hat{\nu}$  is our  $\lambda$ .

## B Additional Experimental Details

### B.1 Neural Networks Used

Here is a list of the network architectures that we use in this paper along with their references if applicable.

#### MNIST robust error experiment

- MLP-A: a multilayer perceptron consisting of 1 hidden layer with 500 neurons (Tjeng et al., 2019).
- MLP-B: a multilayer perceptron consisting of 2 hidden layers with 100 neurons each.

#### MNIST $\epsilon$ -search experiment

- CNN-SMALL: ConvNet architecture with two convolutional layers with 16 and 32 filters respectively (size  $4 \times 4$  and stride of 2 in both), followed by two fully-connected layers with 100 and 10 units respectively (Wong et al., 2018).
- CNN-WIDE-K: ConvNet architecture with two convolutional layers of  $4 \times k$  and  $8 \times k$  filters (size  $4 \times 4$  and stride of 2 in both) followed by a  $128 \times k$  fully connected layer followed by two fully-connected layers of sizes  $128 \times k$  and 10 respectively. The parameter  $k$  is used to control the width of the network (Wong et al., 2018).
- CNN-DEEP-K: ConvNet architecture with  $k$  convolutional layers with 8 filters followed by  $k$  convolutional filters with 16 filters followed by two fully-connected layers of sizes  $100 \times k$  and 10 respectively. The parameter  $k$  is used to control the depth of the network (Wong et al., 2018).
- MLP-[9]-500: a multilayer perceptron consisting of 9 hidden layer with 500 neurons each.
- MLP-[9]-100: a multilayer perceptron consisting of 9 hidden layer with 100 neurons each.
- MLP-[2]-100: a multilayer perceptron consisting of 2 hidden layer with 100 neurons each.

#### CIFAR-10 $\epsilon$ -search experiment

- CNN-SMALL: ConvNet architecture with two convolutional layers with 16 and 32 filters respectively (size  $4 \times 4$  and stride of 2 in both), followed by two fully-connected layers with 100 and 10 units respectively.
- CNN-WIDE-K: ConvNet architecture with two convolutional layers of  $4 \times k$  and  $8 \times k$  filters (size  $4 \times 4$  and stride of 2 in both) followed by a  $128 \times k$  fully connected layer followed by two fully-connected layers of sizes  $128 \times k$  and 10 respectively. The parameter  $k$  is used to control the width of the network.

### B.2 Training Modes

In this paper, we use only one pre-trained network from the literature, and we train the rest from scratch.

#### Pre-trained Networks

- ADV-MLP-A: this is a multilayer perceptron with 1 hidden layer having 500 units. It is trained using PGD with  $l_\infty$  perturbation of  $\epsilon = 0.1$ , and is used in Tjeng et al. (2019) and Raghunathan et al. (2018a). It can be found at [https://github.com/vtjeng/MIPVerify\\_data/tree/master/weights/mnist/RSL18a](https://github.com/vtjeng/MIPVerify_data/tree/master/weights/mnist/RSL18a).

**Networks Trained from Scratch.** We train all models in parallel on a GPU-cluster with P100 GPUs.

- All networks in the paper that have the prefix or training mode ADV are trained with PGD using the code available at [https://github.com/locuslab/convex\\_adversarial/blob/master/examples/mnist.py](https://github.com/locuslab/convex_adversarial/blob/master/examples/mnist.py).

- All networks in the paper that have the prefix or training mode LPD are trained with the robust training method of Wong et al. (2018) using the code available at [https://github.com/locuslab/convex\\_adversarial/blob/master/examples/mnist.py](https://github.com/locuslab/convex_adversarial/blob/master/examples/mnist.py).
- All networks in the paper that have the prefix NOR or training mode NORMAL are trained the regular cross-entropy loss using the code available at [https://github.com/locuslab/convex\\_adversarial/blob/master/examples/mnist.py](https://github.com/locuslab/convex_adversarial/blob/master/examples/mnist.py).
- All the CIFAR-10 networks in the paper have the same naming convention as above, but are trained using the code available at [https://github.com/locuslab/convex\\_adversarial/blob/master/examples/cifar.py](https://github.com/locuslab/convex_adversarial/blob/master/examples/cifar.py).

## C Parallel Computation Details

**Why do we need parallel computing to solve LP-ALL?** The nature of our LP-ALL algorithm requires solving a number of LP that scales with the number of neurons in the network we are verifying. For example, if we want to verify a network with 10k neurons on ten samples the MNIST dataset. We need to solve roughly  $10k \text{ LPs/sample} \times 10 \text{ samples} = 100k \text{ LPs}$ .

The average time for solving an LP varies with the size of the network (see Fig. 5 and 6). It also varies depending on which layer in the network the neuron, for which we are solving the LP, is in (see Fig. 7). Let us say on average the duration for solving one LP is 10 sec on the CPUs we use, which is reasonable for networks that we consider in this paper. *Therefore, for verifying one network, we need around 1 million sec which is roughly 11 days.*

Doing this for all the models in the paper and for more samples would take years. This is why parallelizing the computation was crucial. Therefore we conduct all the experiments on a cluster with **1000 CPU-nodes**. Another key point here was to make sure that the scheduling pipeline on the cluster has **very low latency**, because we need to solve around 100 million jobs in total in the paper, each of which is on the order of seconds. So any latency in the pipeline can cause significant overhead. The details of the scheduling pipeline are beyond the scope of this paper.

**CPU specifications.** Each CPU-node we used has 2 virtual CPUs with a 2.4 GHz Intel(R) Xeon(R) E5-2673 v3 (Haswell) processor and 7GB of RAM.

**Linear programming (LP) solver used.** We construct all the LP models in python using CVXPY (Diamond & Boyd, 2016), and the models are solved using an open-source solver, ECOS (Domahidi et al., 2013). We found this solver to be the fastest among other open-source solvers for our application.

## D Computational Time for Solving LP-ALL

The solve time of the LP in (C) depends mainly on the size and the training method of a neural network. It also depends on the input-space dimension.

**Dependence on architecture and training mode.** Fig. 5 and 6 shows the average solve time of the LP in (C) for various networks and training methods that are used in the paper on MNIST and CIFAR-10 datasets, respectively. This averaging is over all the neurons in each network, and over ten samples of each dataset. Note how the solve time increases as the network becomes wider or deeper. This is because the number of decision variables and constraints in the LP increases as the network becomes wider or deeper. Another observation is that, in contrast to MILP (Tjeng et al., 2019), the solve time for robustly trained networks seems to be larger than those which are trained using the regular cross-entropy loss or those which are randomly initialized. This is possibly due to the fact that we are not exploiting the stability of neurons in our implementation of the LP as opposed to what is done in the MILP implementation of Tjeng et al. (2019).

**Dependence on which layer we are solving for.** Fig. 7 shows the average solve time per neuron per layer of the LP in (C) for each of the networks that are used in the paper on the CIFAR-10 dataset. Notice how the solve time of the LP increases as we go deeper into the network.

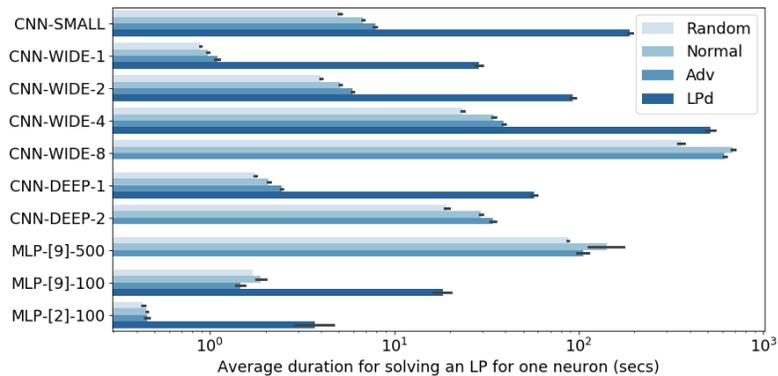


Figure 5: Average duration for solving the LPs for each model (averaged over the neurons in the model and over 10 samples of the MNIST dataset).

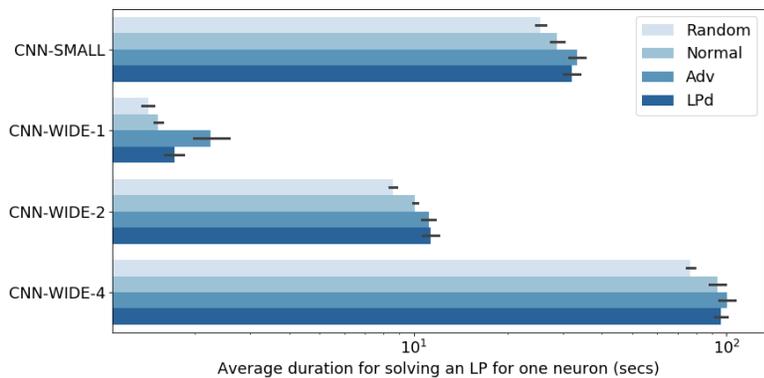


Figure 6: Average duration for solving the LPs for each model averaged over the neurons in the model and over 10 samples of the CIFAR-10 dataset.

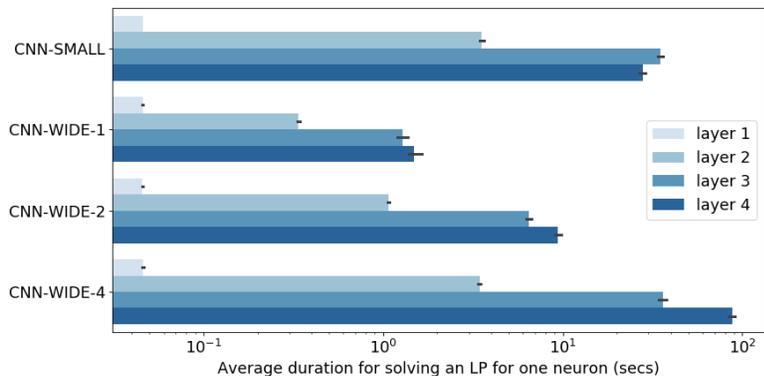


Figure 7: Average duration for solving the LPs per layer for each model averaged over the neurons in the model and over 10 samples of the CIFAR-10 dataset.

## E Full Results of Certified Bounds on the Minimum Adversarial Distortion Experiment

### E.1 Implementation details

In this experiment, we are interested in searching for the minimum adversarial distortion  $\epsilon$ , which is the  $l_\infty$  radius of largest  $l_\infty$  ball in which no adversarial examples can be crafted.

An upper bound on  $\epsilon$  can be calculated by using PGD in a binary search setting: given an initial guess of  $\epsilon$ , PGD can be used to find an adversarial example. If successful, divide  $\epsilon$  by 2; else multiply  $\epsilon$  by 2; and repeat until the change in  $\epsilon$  is below some tolerance ( $10^{-5}$  in our case).

Lower bounds on  $\epsilon$  are calculated using LP-GREEDY, LP-LAST, or our LP-ALL algorithm in a binary search setting; given an initial guess of  $\epsilon$ , any of these algorithms can be used to check whether the network is robust within  $\epsilon$ -perturbation of the input. If robust, multiply  $\epsilon$  by 2; else divide  $\epsilon$  by 2; and repeat until the change in  $\epsilon$  is below a tolerance. The tolerances used in the paper are:

- $\text{tol}(\epsilon_{\text{LP-GREEDY}}) = 10^{-5}$  because LP-GREEDY is computationally very cheap.
- $\text{tol}(\epsilon_{\text{LP-LAST}}) = 5\% \times \epsilon_{\text{LP-GREEDY}}$  because LP-LAST is computationally expensive.
- $\text{tol}(\epsilon_{\text{LP-ALL}}) = 5\% \times \epsilon_{\text{LP-GREEDY}}$  because LP-ALL is computationally expensive.

Since solving LP-ALL is really expensive, we find the  $\epsilon$ -bounds only for ten samples of the MNIST and CIFAR-10 datasets. In this experiment, both ADV- and LPD-networks are trained with an  $l_\infty$  maximum allowed perturbation of 0.1 and 8/255 on MNIST and CIFAR-10, respectively. The full results are reported in Tables 2 and 3 respectively.

### E.2 Results

Tables 2 and 3 both report, for ten samples of MNIST and CIFAR-10 respectively, for a wide range of networks :

1. The training mode, whether the network is trained using regular CE loss (Normal), using adversarial examples generated by PGD (ADV), or using the robust loss in [Wong & Kolter \(2018\)](#) (LPD).
2. Mean lower bounds on  $\epsilon$  found by LP-GREEDY, LP-LAST, and LP-ALL. Note that naturally

$$\epsilon_{\text{LP-GREEDY}} \leq \epsilon_{\text{LP-LAST}} \leq \epsilon_{\text{LP-ALL}}$$

3. A mean upper bound on  $\epsilon$  found by PGD.
4. The median percentage gap between PGD and the three LP-relaxed algorithms. The percentage gap is defined as

$$\% \text{gap} = \frac{(\epsilon_{\text{PGD}} - \epsilon_{\text{LP-X}})}{\epsilon_{\text{PGD}}} \times 100.$$

It is also easy to see that naturally,

$$\% \text{gap}_{\text{LP-GREEDY}} \geq \% \text{gap}_{\text{LP-LAST}} \geq \% \text{gap}_{\text{LP-ALL}}$$

The results of both tables show that for all networks, the certified lower bounds on  $\epsilon$  using LP-GREEDY, LP-LAST, or LP-ALL are 1.5 to 5 times smaller than the upper bound found by PGD on MNIST, and 1.5 to 2 times smaller than the upper bound found by PGD on CIFAR-10. This gap can also clearly be observed in Fig. 3 and Fig. 4 for MNIST and CIFAR-10, respectively.

Therefore, the improvement that we get using LP-ALL and LP-LAST over LP-GREEDY is not significant and doesn't close the gap with the PGD upper bound.

## F Results on Randomly Initialized Networks

In this section, we report additional results for the  $\epsilon$ -search experiment because they might be of interest as a comparison. The results are reported in Table 4. The results are in accordance to what

Table 2: Certified bounds on the minimum adversarial distortion  $\epsilon$  for ten random samples from the test set of MNIST.

NETWORK	TRAINING MODE	MEAN LOWER BOUND ( $\times 10^{-3}$ )			MEAN UPPER BOUND ( $\times 10^{-3}$ )	MEDIAN PERCENTAGE GAP (%)		
		LP-GREEDY	LP-LAST	LP-ALL	PGD	LP-GREEDY	LP-LAST	LP-ALL
CNN-SMALL	NORMAL	14.98	16.29	18.87	52.70	69.12	66.03	61.40
	ADV	73.42	77.09	85.94	155.16	52.52	50.14	44.42
	LPD	153.17	160.83	160.83	226.72	29.72	26.21	26.21
CNN-WIDE-1	NORMAL	14.09	15.76	16.92	39.61	58.84	54.69	52.66
	ADV	81.52	86.25	91.76	142.89	43.59	40.77	37.58
	LPD	116.72	122.55	122.55	183.66	33.90	30.59	30.59
CNN-WIDE-2	NORMAL	13.29	14.83	16.82	43.95	68.34	64.40	60.42
	ADV	91.50	96.08	104.02	179.86	49.83	47.32	41.98
	LPD	148.07	156.78	169.77	221.67	32.45	27.76	21.03
CNN-WIDE-4	NORMAL	12.84	14.37	16.45	47.23	72.23	68.06	63.06
	ADV	67.64	72.34	79.90	178.01	62.72	59.37	55.17
	LPD	142.30	149.41	155.23	217.64	34.92	31.67	29.34
CNN-WIDE-8	NORMAL	10.82	11.72	13.35	47.75	75.49	71.85	69.36
	ADV	62.57	67.42	77.66	181.09	64.45	62.17	55.57
	LPD	N.A	N.A	N.A	N.A	N.A	N.A	N.A
CNN-DEEP-1	NORMAL	15.21	16.78	19.58	44.79	66.50	62.04	55.44
	ADV	94.68	99.41	100.20	166.38	39.81	36.80	35.93
	LPD	136.09	142.89	142.89	184.23	22.10	18.20	18.20
CNN-DEEP-2	NORMAL	6.12	6.42	8.76	43.32	84.47	83.69	78.65
	ADV	102.47	107.60	112.82	185.70	39.35	36.32	36.32
	LPD	N.A	N.A	N.A	N.A	N.A	N.A	N.A
MLP-[9]-500	NORMAL	12.64	13.27	16.84	45.84	74.57	73.30	63.14
	ADV	20.77	21.99	28.50	129.45	84.60	83.83	79.05
	LPD	N.A	N.A	N.A	N.A	N.A	N.A	N.A
MLP-[9]-100	NORMAL	11.35	11.92	14.23	31.37	64.13	62.34	57.03
	ADV	19.41	21.12	25.41	94.57	75.15	71.42	63.96
	LPD	68.25	71.51	73.96	103.87	29.79	26.28	26.28
MLP-[2]-100	NORMAL	14.19	15.11	15.83	28.14	52.66	47.82	45.56
	ADV	41.68	43.76	43.76	81.22	36.23	33.04	33.04
	LPD	81.50	85.33	85.33	118.10	25.01	21.26	21.26

Table 3: Certified bounds on the minimum adversarial distortion  $\epsilon$  for ten random samples from the test set of CIFAR-10.

NETWORK	TRAINING MODE	MEAN LOWER BOUND ( $\times 10^{-3}$ )			MEAN UPPER BOUND ( $\times 10^{-3}$ )	MEDIAN PERCENTAGE GAP (%)		
		LP-GREEDY	LP-LAST	LP-ALL	PGD	LP-GREEDY	LP-LAST	LP-ALL
CNN-SMALL	NORMAL	7.48	7.86	8.46	20.13	49.40	46.87	44.23
	ADV	24.33	26.53	27.59	37.90	34.50	24.67	24.67
	LPD	67.34	72.27	77.84	157.01	52.94	48.13	43.13
CNN-WIDE-1	NORMAL	6.97	7.32	7.56	14.57	43.01	40.16	39.39
	ADV	58.52	63.26	67.84	115.47	49.83	46.63	42.15
	LPD	57.03	62.51	65.83	122.00	41.22	38.29	32.40
CNN-WIDE-2	NORMAL	8.27	8.86	9.46	22.16	58.66	54.53	52.46
	ADV	42.05	45.99	49.09	74.13	35.10	29.85	25.54
	LPD	73.19	81.75	87.38	157.03	47.64	39.78	39.78
CNN-WIDE-4	NORMAL	4.14	4.35	4.63	10.97	40.27	37.28	33.03
	ADV	29.11	32.84	35.45	71.57	50.59	44.21	43.18
	LPD	41.62	47.17	48.51	104.49	45.19	39.67	39.67

was discussed in Seciton 6.2 i.e. for all networks and both datasets, the certified lower bounds on  $\epsilon$  using LP-GREEDY, LP-LAST, or LP-ALL are 2 to 3 times smaller than the upper bound found by PGD. Furthermore, the improvement that we get using LP-ALL and LP-LAST over LP-GREEDY is not significant and doesn't close the gap with the PGD upper bound.

Table 4: Certified bounds on the minimum adversarial distortion  $\epsilon$  for ten random samples from the test set of MNIST and CIFAR-10 on randomly initialized networks (no training).

NETWORK	TRAINING MODE	MEAN LOWER BOUND ( $\times 10^{-3}$ )			MEAN UPPER BOUND ( $\times 10^{-3}$ )	MEDIAN PERCENTAGE GAP (%)		
		LP-GREEDY	LP-LAST	LP-ALL	PGD	LP-GREEDY	LP-LAST	LP-ALL
<b>MNIST</b>								
CNN-SMALL	RANDOM	5.79	6.08	6.25	14.86	51.37	48.94	48.94
CNN-WIDE-1	RANDOM	10.42	10.94	11.98	33.77	67.09	65.45	62.16
CNN-WIDE-2	RANDOM	8.12	8.53	9.34	29.43	72.54	71.17	68.42
CNN-WIDE-4	RANDOM	8.68	9.12	9.99	45.26	78.65	77.59	75.45
CNN-DEEP-1	RANDOM	11.12	11.81	12.79	42.28	72.76	71.40	68.67
MLP-[2]-100	RANDOM	4.69	5.16	5.25	15.71	64.85	58.53	57.83
<b>CIFAR-10</b>								
CNN-SMALL	RANDOM	8.77	10.01	10.13	24.50	62.61	57.04	57.04
CNN-WIDE-1	RANDOM	5.61	5.89	6.09	11.33	45.27	42.53	41.46
CNN-WIDE-2	RANDOM	2.83	3.31	3.31	6.24	50.60	46.13	46.13
CNN-WIDE-4	RANDOM	8.93	8.52	9.00	28.69	69.63	68.11	68.11