

# UNDERSTANDING CERTIFIED TRAINING WITH INTERVAL BOUND PROPAGATION

Yuhao Mao, Mark Niklas Müller, Marc Fischer & Martin Vechev

Department of Computer Science, ETH Zürich, Switzerland

{yuhao.mao, mark.mueller, marc.fischer, martin.vechev}@inf.ethz.ch

## ABSTRACT

As robustness verification methods are becoming more precise, training certifiably robust neural networks is becoming ever more relevant. To this end, certified training methods compute and then optimize an upper bound on the worst-case loss over a robustness specification. Curiously, training methods based on the imprecise interval bound propagation (IBP) consistently outperform those leveraging more precise bounds. Still, we lack a theoretical understanding of the mechanisms making IBP so successful. In this work, we investigate these mechanisms by leveraging a novel metric measuring the tightness of IBP bounds. We first show theoretically that, for deep linear models (DLNs), tightness decreases with width and depth at initialization, but improves with IBP training. We, then, derive sufficient and necessary conditions on weight matrices for IBP bounds to become exact and demonstrate that these impose strong regularization, providing an explanation for the observed robustness-accuracy trade-off. Finally, we show how these results on DLNs transfer to ReLU networks, before conducting an extensive empirical study, (i) confirming this transferability and yielding state-of-the-art certified accuracy, (ii) finding that while all IBP-based training methods lead to high tightness, this increase is dominated by the size of the propagated input regions rather than the robustness specification, and finally (iii) observing that non-IBP-based methods do not increase tightness. Together, these results help explain the success of recent certified training methods and may guide the development of new ones.

## 1 INTRODUCTION

The increasing deployment of deep-learning-based systems in safety-critical domains has made their trustworthiness and especially formal robustness guarantees against adversarial examples (Biggio et al., 2013; Szegedy et al., 2014) an ever more important topic. As significant progress has been made on neural network certification (Zhang et al., 2022; Ferrari et al., 2022), the focus in the field is increasingly shifting to the development of specialized training methods that improve certifiable robustness while minimizing the accompanying reduction in standard accuracy.

**Certified Training** These certified training methods aim to compute and then optimize approximations of the network’s worst-case loss over an input region defined by an adversary specification. To this end, they compute an over-approximation of the network’s reachable set using symbolic bound propagation methods (Singh et al., 2018; 2019b; Gowal et al., 2018). Surprisingly, training methods based on the least precise bounds, obtained via interval bound propagation (IBP), empirically yield the best performance (Shi et al., 2021). Jovanović et al. (2022) investigate this surprising observation theoretically and find that more precise bounding methods induce harder optimization problems.

As a result, *all* methods obtaining state-of-the-art performance leverage IBP bounds either directly (Shi et al., 2021), as regularizer (Palma et al., 2022), or to precisely but unsoundly approximate the worst-case loss (Müller et al., 2022b; Mao et al., 2023; Palma et al., 2023). However, while IBP is crucial to their success, none of these works develop a theoretical understanding of what makes IBP training so effective and how it affects bound tightness and network regularization.

**This Work** We take a first step towards building a deeper understanding of the mechanisms making IBP training so successful and thereby pave the way for further advances in certified training. To this end, we derive necessary and sufficient conditions on a network’s weights under which IBP

bounds become tight, a property we call *propagation invariance*, and prove that it implies an extreme regularization, agreeing well with the empirically observed trade-off between certifiable robustness and accuracy (Tsipras et al., 2019; Müller et al., 2022b). To investigate how close real networks are to full propagation invariance, we introduce the metric *propagation tightness* as the ratio of optimal and IBP bounds, and show how to efficiently compute it globally for deep linear networks (DLNs) and locally for ReLU networks.

This novel metric enables us to theoretically investigate the effects of model architecture, weight initialization, and training methods on IBP bound tightness for deep linear networks (DLNs). We show that (i) at initialization, tightness decreases with width (polynomially) and depth (exponentially), (ii) tightness is increased by IBP training, and (iii) sufficient width becomes crucial for trained networks.

Conducting an extensive empirical study, we confirm the predictiveness of our theoretical results for deep ReLU networks and observe that: (i) increasing network width but not depth improves state-of-the-art certified accuracy, (ii) IBP training significantly increases tightness, almost to the point of propagation invariance, (iii) unsound IBP-based training methods increase tightness to a smaller degree, determined by the size of the propagated input region and the weight of the IBP-loss, but yield better performance, and (iv) non-IBP-based training methods barely increase tightness, leading to higher accuracy but worse robustness. These findings suggest that while IBP-based training methods improve robustness by increasing tightness at the cost of standard accuracy, high tightness is not generally necessary for robustness. This observation explains the recent success of unsound IBP-based methods and, in combination with the theoretical and practical insights developed here, promises to be a key step toward constructing novel and more effective certified training methods.

## 2 BACKGROUND

Here, we provide a background on adversarial and certified robustness. We consider a classifier  $f: \mathbb{R}^{d_{\text{in}}} \mapsto \mathbb{R}^c$  predicting a numerical score  $\mathbf{y} := f(\mathbf{x})$  per class given an input  $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^{d_{\text{in}}}$ .

**Adversarial Robustness** describes the property of a classifier  $f$  to consistently predict the target class  $t$  for all perturbed inputs  $\mathbf{x}'$  in an  $\ell_p$ -norm ball  $\mathcal{B}_p^{\epsilon_p}(\mathbf{x})$  of radius  $\epsilon_p$ . As we focus on  $\ell_\infty$  perturbations in this work, we henceforth drop the subscript  $p$  for notational clarity. More formally, we define *adversarial robustness* as:

$$\arg \max_j f(\mathbf{x}')_j = t, \quad \forall \mathbf{x}' \in \mathcal{B}_p^{\epsilon_p}(\mathbf{x}) := \{\mathbf{x}' \in \mathcal{X} \mid \|\mathbf{x} - \mathbf{x}'\|_p \leq \epsilon_p\}. \quad (1)$$

**Neural Network Certification** can be used to formally prove the robustness of a classifier  $f$  for a given input region  $\mathcal{B}^\epsilon(\mathbf{x})$ . Interval bound propagation (IBP) (Gowal et al., 2018; Mirman et al., 2018) is a simple but popular such certification method. It is based on propagating an input region  $\mathcal{B}^\epsilon(\mathbf{x})$  through a neural network by computing BOX over-approximations (each dimension is described as an interval) of the hidden state after every layer until we reach the output space. There, it is checked whether all points in the resulting over-approximation of the network’s reachable set yield the correct classification. As an example, consider an  $L$ -layer network  $f = \mathbf{h}_L \circ \sigma \circ \mathbf{h}_{L-2} \circ \dots \circ \mathbf{h}_1$ , with linear layers  $\mathbf{h}_i$  and ReLU activation functions  $\sigma$ . We first over-approximate the input region  $\mathcal{B}^\epsilon(\mathbf{x})$  as BOX with radius  $\delta^0 = \epsilon$  and center  $\hat{\mathbf{x}}^0 = \mathbf{x}$ , such that we have the  $i^{\text{th}}$  dimension of the input  $x_i^0 \in [\underline{x}_i, \bar{x}_i] := [\hat{x}_i^0 - \delta_i^0, \hat{x}_i^0 + \delta_i^0]$ . Propagating such a BOX through the linear layer  $\mathbf{h}_i(\mathbf{x}^{i-1}) = \mathbf{W}\mathbf{x}^{i-1} + \mathbf{b} =: \mathbf{x}^i$ , we obtain the output hyperbox with centre  $\hat{\mathbf{x}}^i = \mathbf{W}\hat{\mathbf{x}}^{i-1} + \mathbf{b}$  and radius  $\delta^i = |\mathbf{W}|\delta^{i-1}$ , where  $|\cdot|$  denotes the element-wise absolute value. To propagate a BOX through the ReLU activation  $\text{ReLU}(\mathbf{x}^{i-1}) := \max(0, \mathbf{x}^{i-1})$ , we propagate the lower and upper bound separately, resulting in an output BOX with  $\hat{\mathbf{x}}^i = \frac{\hat{\mathbf{x}}^i + \mathbf{x}^i}{2}$  and  $\delta^i = \frac{\hat{\mathbf{x}}^i - \mathbf{x}^i}{2}$  where  $\mathbf{x}^i = \text{ReLU}(\hat{\mathbf{x}}^{i-1} - \delta^{i-1})$  and  $\bar{\mathbf{x}}^i = \text{ReLU}(\hat{\mathbf{x}}^{i-1} + \delta^{i-1})$ . We proceed this way for all layers obtaining first lower and upper bounds on the network’s output  $\mathbf{y}$  and then an upper bound  $\bar{\mathbf{y}}^\Delta$  on the logit difference  $y_i^\Delta := y_i - y_t$ . Showing that  $\bar{y}_i^\Delta < 0, \forall i \neq t$  is then equivalent to proving adversarial robustness on the considered input region.

We illustrate this propagation process for a two-layer network in Figure 1. There, we show the exact propagation of the input region  $\blacksquare$  in blue, its optimal box approximation  $\blacksquare$  in green, and the IBP approximation as dashed boxes  $\square$ . Note how after the first linear and ReLU layer (third column), the box approximations (both optimal  $\blacksquare$  and IBP  $\square$ ) contain already many points outside the reachable set




Figure 1: Comparison of exact (■), optimal box (■), and IBP (□) propagation through a one layer network. We show the concrete points maximizing the logit difference  $y_2 - y_1$  as a black  $\times$  and the corresponding relaxation as a red  $\times$ .

■, despite it being the smallest hyper-box containing the exact region. These so-called approximation errors accumulate quickly when using IBP, leading to an increasingly imprecise abstraction, as can be seen by comparing the optimal box ■ and IBP □ approximation after an additional linear layer (rightmost column). To verify that this network classifies all inputs in  $[-1, 1]^2$  to class 1, we have to show the upper bound of the logit difference  $y_2 - y_1$  to be less than 0. While the concrete maximum of  $-0.3 \geq y_2 - y_1$  (black  $\times$ ) is indeed less than 0, showing that the network is robust, IBP □ only yields  $0.6 \geq y_2 - y_1$  (red  $\times$ ) and is thus too imprecise to prove it. In contrast, the optimal box ■ yields a precise approximation of the true reachable set, sufficient to prove robustness.

**Training for Robustness** is required to obtain (certifiably) robust neural networks. For a data distribution  $(x, t) \sim \mathcal{D}$ , standard training optimizes the network parametrization  $\theta$  to minimize the expected cross-entropy loss:

$$\theta_{\text{std}} = \arg \min_{\theta} \mathbb{E}_{\mathcal{D}} [\mathcal{L}_{\text{CE}}(\mathbf{f}_{\theta}(x), t)], \quad \text{with} \quad \mathcal{L}_{\text{CE}}(\mathbf{y}, t) = \ln \left( 1 + \sum_{i \neq t} \exp(y_i - y_t) \right). \quad (2)$$

To train for robustness, we, instead, aim to minimize the expected *worst-case loss* for a given robustness specification, leading to a min-max optimization problem:

$$\theta_{\text{rob}} = \arg \min_{\theta} \mathbb{E}_{\mathcal{D}} \left[ \max_{x' \in \mathcal{B}^{\epsilon}(x)} \mathcal{L}_{\text{CE}}(\mathbf{f}_{\theta}(x'), t) \right]. \quad (3)$$

As computing the worst-case loss by solving the inner maximization problem is generally intractable, it is commonly under- or over-approximated, yielding adversarial and certified training, respectively.

**Adversarial Training** optimizes a lower bound on the inner optimization objective in Equation (3). It first computes concrete examples  $x' \in \mathcal{B}^{\epsilon}(x)$  that approximately maximize the loss term  $\mathcal{L}_{\text{CE}}$  and then optimizes the network parameters  $\theta$  for these examples. While networks trained this way typically exhibit good empirical robustness, they remain hard to formally certify and are sometimes vulnerable to stronger attacks (Tramèr et al., 2020; Croce & Hein, 2020).

**Certified Training** typically optimizes an upper bound on the inner maximization objective in Equation (3). The resulting robust cross-entropy loss  $\mathcal{L}_{\text{CE,rob}}(\mathcal{B}^{\epsilon}(x), t) = \mathcal{L}_{\text{CE}}(\bar{\mathbf{y}}^{\Delta}, t)$  is obtained by first computing an upper bound  $\bar{\mathbf{y}}^{\Delta}$  on the logit differences  $\mathbf{y}^{\Delta} := \mathbf{y} - y_t$  with a bound propagation method as described above and then plugging it into the standard cross-entropy loss.

Surprisingly, the imprecise IBP bounds (Mirman et al., 2018; Gowal et al., 2018; Shi et al., 2021) consistently yield better performance than methods based on tighter approximations (Wong et al., 2018; Zhang et al., 2020; Balunovic & Vechev, 2020). Jovanović et al. (2022) trace this back to the optimization problems induced by the more precise methods becoming intractable to solve.

However, the heavy regularization that makes IBP trained networks amenable to certification also severely reduces their standard accuracy. To alleviate the resulting robustness-accuracy trade-off, *all* current state-of-the-art certified training methods combine IBP and adversarial training by using IBP bounds only for regularization (IBP-R (Palma et al., 2022)), by only propagating small, adversarially selected regions (SABR (Müller et al., 2022b)), using IBP bounds only for the first layers and PGD bounds for the remainder of the network (TAPS (Mao et al., 2023)), or combining losses over adversarial samples and IBP bounds (CC-IBP, MTL-IBP (Palma et al., 2023)).

In light of this surprising dominance of IBP-based training methods, understanding the regularization IBP induces and its effect on tightness promises to be a key step towards developing novel and more effective certified training methods.

### 3 UNDERSTANDING IBP TRAINING

In this section, we theoretically investigate the relationship between the box bounds obtained by layer-wise propagation, *i.e.*, IBP, and optimal propagation. We illustrate both in Figure 1 and note that the latter are sufficient for exact robustness certification (see Lemma 3.1). First, we formally define layer-wise (IBP) and optimal box propagation, before deriving sufficient and necessary conditions under which the resulting bounds become identical. Then, we show that these conditions induce strong regularization, motivating us to introduce the propagation tightness  $\tau$  as a relaxed measure of bound precision, which can be efficiently computed globally for deep linear (DLN) and locally for ReLU networks. Based on these results, we first investigate how tightness depends on network architecture at initialization, before showing that it improves with IBP training. Finally, we demonstrate that even linear dimensionality reduction is inherently imprecise for both optimal and IBP propagation, making sufficient network width key for tight box bounds. We defer all proofs to App. B.

**Setting** We focus our theoretical analysis on deep linear networks (DLNs), *i.e.*,  $\mathbf{f}(x) = \prod_{i=1}^L \mathbf{W}^{(i)}x$ , popular for theoretical discussion of neural networks (Saxe et al., 2014; Ji & Telgarsky, 2019; Wu et al., 2019). While such a reduction of a ReLU network to an overall linear function may seem restrictive, it preserves many interesting properties and allows for theoretical insights, while ReLU networks are theoretically unwieldy. As ReLU networks become linear for fixed activation patterns, the DLN approximation becomes exact for robustness analysis at infinitesimal perturbation magnitudes. Further, DLNs retain the layer-wise structure and joint non-convexity in the weights of different layers of ReLU networks, making them a widely popular analysis tool (Ribeiro et al., 2016). After proving key results on DLNs, we will show how they transfer to ReLU networks.

#### 3.1 LAYER-WISE AND OPTIMAL BOX PROPAGATION

We define the optimal hyper-box approximation  $\text{Box}^*(\mathbf{f}, \mathcal{B}^\epsilon(x))$  as the smallest hyper-box  $[\underline{z}, \bar{z}]$  such that it contains the image  $\mathbf{f}(x')$  of all points  $x'$  in  $\mathcal{B}^\epsilon(x)$ , *i.e.*,  $\mathbf{f}(x') \in [\underline{z}, \bar{z}], \forall x' \in \mathcal{B}^\epsilon(x)$ . Similarly, we define the layer-wise box approximation as the result of sequentially applying the optimal approximation to every layer individually:  $\text{Box}^\dagger(\mathbf{f}, \mathcal{B}^\epsilon(x)) := \text{Box}^*(\mathbf{W}_L, \text{Box}^*(\dots, \text{Box}^*(\mathbf{W}^{(1)}, \mathcal{B}^\epsilon(x))))$ . We write their upper- and lower-bounds as  $[\underline{z}^*, \bar{z}^*]$  and  $[\underline{z}^\dagger, \bar{z}^\dagger]$ , respectively. We note that optimal box bounds on the logit differences  $\mathbf{y}^\Delta := \mathbf{y} - y_t$  (instead of on the logits  $\mathbf{y}$  as shown in Figure 1) are sufficient for exact robustness verification:

**Lemma 3.1.** *Any  $C^0$  continuous classifier  $\mathbf{f}$ , computing the logit difference  $y_i^\Delta := y_i - y_t, \forall i \neq t$ , is robustly correct on  $\mathcal{B}^\epsilon(x)$  if and only if  $\text{Box}^*(\mathbf{f}, \mathcal{B}^\epsilon(x)) \subseteq \mathbb{R}_{<0}^{c-1}$ , *i.e.*,  $\bar{y}_i^{\Delta*} < 0, \forall i \neq t$ .*

For DLNs, we can efficiently compute both optimal  $\text{Box}^*$  and layerwise  $\text{Box}^\dagger$  box bounds as follows:

**Theorem 3.2 (Box Propagation).** *For an  $L$ -layer DLN  $\mathbf{f} = \prod_{k=1}^L \mathbf{W}^{(k)}$ , we obtain the box centres  $\underline{z}^* = \underline{z}^\dagger = \mathbf{f}(x)$  and the radii*

$$\frac{\bar{z}^* - \underline{z}^*}{2} = \left| \prod_{k=1}^L \mathbf{W}^{(k)} \right| \epsilon, \quad \text{and} \quad \frac{\bar{z}^\dagger - \underline{z}^\dagger}{2} = \left( \prod_{k=1}^L \left| \mathbf{W}^{(k)} \right| \right) \epsilon. \quad (4)$$

Comparing the radius computation of the optimal and layer-wise approximations, we observe that the main difference lies in where the element-wise absolute value  $|\cdot|$  of the weight matrix is taken. For the optimal box, we first multiply all weight matrices before taking the absolute value  $|\prod_{k=1}^L \mathbf{W}^{(k)}|$ , thus allowing for cancellations of terms of opposite signs. For the layer-wise approximation, in contrast, we first take the absolute value of each weight matrix before multiplying them together  $\prod_{k=1}^L |\mathbf{W}^{(k)}|$ , thereby losing all relational information between variables. Let us now investigate under which conditions layer-wise and optimal bounds become identical.

#### 3.2 PROPAGATION INVARIANCE AND IBP BOUND TIGHTNESS

**Propagation Invariance** We call a network (globally) *propagation invariant* (PI) if the layer-wise and optimal box over-approximations are identical for every input box. Clearly, non-negative weight matrices lead to PI networks (Lin et al., 2022), as the absolute value in Theorem 3.2 loses its effect. However, non-negative weights significantly reduce network expressiveness and performance

(Chorowski & Zurada, 2014), raising the question of whether they are a necessary condition. We show that they are not, by deriving a sufficient *and* necessary condition for a two-layer DLN:

**Lemma 3.3** (Propagation Invariance). *A two-layer DLN  $\mathbf{f} = \mathbf{W}^{(2)}\mathbf{W}^{(1)}$  is propagation invariant if and only if for every fixed  $(i, j)$ , we have  $\left| \sum_k W_{i,k}^{(2)} \cdot W_{k,j}^{(1)} \right| = \sum_k |W_{i,k}^{(2)} \cdot W_{k,j}^{(1)}|$ , i.e., either  $W_{i,k}^{(2)} \cdot W_{k,j}^{(1)} \geq 0$  for all  $k$  or  $W_{i,k}^{(2)} \cdot W_{k,j}^{(1)} \leq 0$  for all  $k$ .*

**Conditions for Propagation Invariance** To see how strict the condition described by Lemma 3.3 is, we observe that propagation invariance requires the sign of the last element in any two-by-two block in  $\mathbf{W}^{(2)}\mathbf{W}^{(1)}$  to be determined by the signs of the other three elements:

**Theorem 3.4** (Non-Propagation Invariance). *Assume  $\exists i, i', j, j'$ , such that  $W_{i,j}^{(1)}$ ,  $W_{i',j'}^{(1)}$ ,  $W_{i,j}^{(2)}$  and  $W_{i',j'}^{(2)}$  are all non-zero. If  $(\mathbf{W}^{(2)}\mathbf{W}^{(1)})_{i,j} \cdot (\mathbf{W}^{(2)}\mathbf{W}^{(1)})_{i,j'} \cdot (\mathbf{W}^{(2)}\mathbf{W}^{(1)})_{i',j} \cdot (\mathbf{W}^{(2)}\mathbf{W}^{(1)})_{i',j'} < 0$ , then  $\mathbf{f} = \mathbf{W}^{(2)}\mathbf{W}^{(1)}$  is not propagation invariant.*

To obtain a propagation invariant network with weights  $\mathbf{W}^{(2)}\mathbf{W}^{(1)} \in \mathcal{R}^{d \times d}$ , we can thus only choose  $2d - 1$  (e.g., one row and one column) of the  $d^2$  signs freely (see Corollary A.1 in App. A).

The statements of Lemma 3.3 and Theorem 3.4 naturally extend to DLNs with more than two layers  $L > 2$ . However, the conditions within Theorem 3.4 become increasingly complex and strict as more and more terms need to yield the same sign. Thus, we focus our analysis on  $L = 2$  for clarity.

**IBP Bound Tightness** To analyze the tightness of IBP bounds for networks that do not satisfy the strict conditions for propagation invariance, we relax it to introduce *propagation tightness* as the ratio between the optimal and layer-wise box radii, simply referred to as *tightness* in this paper.

**Definition 3.5.** *Given a DLN  $\mathbf{f}$ , we define the global propagation tightness  $\tau$  as the ratio between optimal  $\text{Box}^*(\mathbf{f}, \mathcal{B}^\epsilon(\mathbf{x}))$  and layer-wise  $\text{Box}^\dagger(\mathbf{f}, \mathcal{B}^\epsilon(\mathbf{x}))$  approximation radius, i.e.,  $\tau = \frac{\underline{z}^* - \bar{z}^*}{\underline{z}^\dagger - \bar{z}^\dagger}$ .*

Intuitively, tightness measures how much smaller the exact dimension-wise  $\text{Box}^*$  bounds are, compared to the layer-wise approximation  $\text{Box}^\dagger$ , thus quantifying the gap between IBP certified and true adversarial robustness. When tightness equals 1, the network is propagation invariant and can be certified exactly with IBP; when tightness is close to 0, IBP bounds become arbitrarily imprecise. We highlight that this is orthogonal to the box diameter  $\Delta = \bar{z}^\dagger - \underline{z}^\dagger$ , considered by Shi et al. (2021).

**ReLU Networks** The nonlinearity of ReLU networks leads to locally varying tightness and makes the computation of optimal box bounds intractable. However, for infinitesimal perturbation magnitudes  $\epsilon$ , the activation patterns of ReLU networks remain stable, making them locally linear. We thus introduce a local version of tightness around concrete inputs.

**Definition 3.6.** *For an  $L$ -layer ReLU network with weight matrices  $\mathbf{W}^{(k)}$  and activation pattern  $\mathbf{d}^{(k)}(\mathbf{x}) = \mathbb{1}_{\mathbf{x}^{(k-1)} > 0} \in \{0, 1\}^{d_k}$  (1 for active and 0 for inactive ReLUs), depending on the input  $\mathbf{x}$ , we define its local tightness as*

$$\tau = \frac{\frac{d}{d\epsilon}(\underline{z}^* - \bar{z}^*)|_{\epsilon=0}}{\frac{d}{d\epsilon}(\underline{z}^\dagger - \bar{z}^\dagger)|_{\epsilon=0}} = \frac{|\prod_{k=1}^L \text{diag}(\mathbf{d}^{(k)})\mathbf{W}^{(k)}| \mathbf{1}}{(\prod_{k=1}^L \text{diag}(\mathbf{d}^{(k)})|\mathbf{W}^{(k)})| \mathbf{1}}.$$

In Definition 3.6, we calculate tightness as the ratio of box size growth rates, evaluated for an infinitesimal input box size  $\epsilon$ . In this setting, the ReLU network will not have any unstable neurons, making our analysis exact. Only when considering larger perturbation magnitudes will neurons become unstable, making our analysis an approximation of the tightness at that  $\epsilon$ . However, for the networks and perturbation magnitudes typically considered in the literature, only a very small fraction ( $\approx 1\%$ ) of neurons are unstable (Müller et al., 2022b). To assess the estimation quality of our local tightness, we show its mean relative error compared to the exact tightness computed with MILP for a small CNN3 in Figure 2 for MNIST and in Figure 14 for CIFAR-10. We find that for perturbations smaller than those used during training ( $\epsilon \leq 0.05$ ) relative errors are extremely small ( $< 0.5\%$ ), and only increase slowly after, reaching 2.2% at  $\epsilon = 0.1$ .




Figure 2: Mean relative error between local tightness (Definition 3.6) and true tightness computed with MILP for a CNN3 trained with PGD or IBP at  $\epsilon = 0.05$  on MNIST.

### 3.3 TIGHTNESS AT INITIALIZATION

We first investigate the (expected) tightness  $\tau = \frac{\mathbb{E}_{\mathcal{D}_\theta}(\mathbf{z}^* - \bar{\mathbf{z}}^*)}{\mathbb{E}_{\mathcal{D}_\theta}(\bar{\mathbf{z}}^\dagger - \mathbf{z}^\dagger)}$  (independent of the dimension due to symmetry) at initialization, *i.e.*, w.r.t. a weight distribution  $\mathcal{D}_\theta$ . Let us consider a two-layer DLN at initialization, *i.e.*, with i.i.d. weights following a zero-mean Gaussian distribution  $\mathcal{N}(0, \sigma^2)$  with an arbitrary but fixed variance  $\sigma^2$  (Glorot & Bengio, 2010; He et al., 2015).

**Lemma 3.7** (Initialization Tightness w.r.t. Width). *Given a 2-layer DLN with weight matrices  $\mathbf{W}^{(1)} \in \mathcal{R}^{d_1 \times d_0}$ ,  $\mathbf{W}^{(2)} \in \mathcal{R}^{d_2 \times d_1}$  with i.i.d. entries from  $\mathcal{N}(0, \sigma_1^2)$  and  $\mathcal{N}(0, \sigma_2^2)$  (together denoted as  $\theta$ ), we obtain the expected tightness  $\tau(d_1) = \frac{\mathbb{E}_\theta(\mathbf{z}^* - \bar{\mathbf{z}}^*)}{\mathbb{E}_\theta(\bar{\mathbf{z}}^\dagger - \mathbf{z}^\dagger)} = \frac{\sqrt{\pi} \Gamma(\frac{1}{2}(d_1+1))}{d_1 \Gamma(\frac{1}{2}d_1)} \in \Theta(\frac{1}{\sqrt{d_1}})$ .*

Tightness at initialization, thus, decreases quickly with internal width ( $\Theta(\frac{1}{\sqrt{d_1}})$ ), *e.g.*, by a factor of  $\tau(500) \approx 0.056$  for the penultimate layer of the popular CNN7 (Gowal et al., 2018; Zhang et al., 2020). It, further, follows directly that tightness will decrease exponentially w.r.t. network depth.

**Corollary 3.8** (Initialization Tightness w.r.t. Depth). *The expected tightness of an  $L$ -layer DLN  $f$  with minimum internal dimension  $d_{\min}$  is at most  $\tau \leq \tau(d_{\min})^{\lfloor \frac{L}{2} \rfloor}$  at initialization.*

This result is independent of the variance  $\sigma_1^2, \sigma_2^2$ . Thus, tightness at initialization can not be increased by scaling  $\sigma^2$ , as proposed by Shi et al. (2021) to achieve constant box radius over network depth.

**ReLU Networks** We extend Lemma 3.7 to two-layer ReLU networks (Corollary A.2 in App. A), obtain an expected tightness of  $\sqrt{2}\tau(d_1)$ , and empirically validate it in Section 4.1.

### 3.4 IBP TRAINING INCREASES TIGHTNESS

We now show theoretically that IBP training increases this tightness. To this end, we again consider a DLN with layer-wise propagation matrix  $\mathbf{W}^\dagger = \prod_{i=1}^L |\mathbf{W}^{(i)}|$  and optimal propagation matrix  $\mathbf{W}^* = |\prod_{i=1}^L \mathbf{W}^{(i)}|$ , yielding the expected risk for IBP training as  $R(\epsilon) = \mathbb{E}_{\mathbf{x}, y} \mathcal{L}(\text{Box}^\dagger(\mathbf{f}, \mathcal{B}^\epsilon(\mathbf{x})), y)$ .

**Theorem 3.9** (IBP Training Increases Tightness). *Assume homogenous tightness, *i.e.*,  $\mathbf{W}^* = \tau \mathbf{W}^\dagger$ , and  $\frac{\|\nabla_\theta \mathbf{W}_{ij}^*\|_2}{\mathbf{W}_{ij}^*} \leq \frac{1}{2} \frac{\|\nabla_\theta \mathbf{W}_{ij}^\dagger\|_2}{\mathbf{W}_{ij}^\dagger}$  for all  $i, j$ , then, the gradient difference between the IBP and standard loss is aligned with an increase in tightness, *i.e.*,  $\langle \nabla_\theta(R(\epsilon) - R(0)), \nabla_\theta \tau \rangle \leq 0$  for all  $\epsilon > 0$ .*

### 3.5 NETWORK WIDTH AND TIGHTNESS AFTER TRAINING

Many high-dimensional computer vision datasets were shown to have low intrinsic data dimensionality (Pope et al., 2021). Thus, we study the reconstruction loss of a linear embedding into a low-dimensional subspace as a proxy for performance and find that tightness decreases with the width  $w$  of a bottleneck layer as long as it is smaller than the data-dimensionality  $d$ , *i.e.*,  $w \ll d$ . Further, while reconstruction becomes lossless for points as soon as the width  $w$  reaches the intrinsic dimension  $k$  of the data, even optimal box propagation requires a width of at least the original data dimension  $d$  to achieve loss-less reconstruction. For a  $k$ -dimensional data distribution, linearly embedded into a  $d$  dimensional space with  $d \gg k$ , the data matrix  $X$  has a low-rank eigendecomposition  $\text{Var}(X) = U \Lambda U^\top$  with  $k$  non-zero eigenvalues. The optimal reconstruction  $\hat{X} = U_k U_k^\top X$  is exact by choosing  $U_k$  as the  $k$  columns of  $U$  with non-zero eigenvalues. Yet, box propagation is imprecise:

**Theorem 3.10** (Box Reconstruction Error). *Consider the linear embedding and reconstruction  $\hat{\mathbf{x}} = U_k U_k^\top \mathbf{x}$  of a  $d$  dimensional data distribution  $\mathbf{x} \sim \mathcal{X}$  into a  $k$  dimensional space with  $d \gg k$  and eigenmatrices  $U$  drawn uniformly at random from the orthogonal group. Propagating the input box  $\mathcal{B}^\epsilon(\mathbf{x})$  layer-wise and optimally, thus, yields  $\mathcal{B}^{\delta^\dagger}(\hat{\mathbf{x}})$ , and  $\mathcal{B}^{\delta^*}(\hat{\mathbf{x}})$ , respectively. Then, we have, (i)  $\mathbb{E}(\delta_i/\epsilon) = ck \in \Theta(k)$  for a positive constant  $c$  depending solely on  $d$  and  $c \rightarrow \frac{2}{\pi} \approx 0.64$  for large  $d$ ; and (ii)  $\mathbb{E}(\delta_i^*/\epsilon) \rightarrow \frac{2}{\sqrt{\pi}} \frac{\Gamma(\frac{1}{2}(k+))}{\Gamma(\frac{1}{2}k)} \in \Theta(\sqrt{k})$ .*

Intuitively, Theorem 3.10 implies that, while input points can be embedded into and reconstructed from a  $k$  dimensional space losslessly, box propagation will yield a box growth of  $\Theta(\sqrt{k})$  for optimal and  $\Theta(k)$  for layer-wise propagation. However, with  $k = d$ , we can choose  $U_k$  to be an identity matrix, thus obtaining lossless "reconstruction", highlighting the importance of network width.




Figure 3: Dependence of tightness at initialization on width (left) and depth (right) for a CNN7 and CIFAR-10.




Figure 4: Box reconstruction error over bottleneck width  $w$ .

## 4 EMPIRICAL EVALUATION ANALYSIS

Now, we conduct an empirical study of IBP-based certified training, leveraging our novel tightness metric and specifically its local variant (see Definition 3.6) to gain a deeper understanding of these methods and confirm the applicability of our theoretical analysis to ReLU networks. For certification, we use MN-BAB (Ferrari et al., 2022), a state-of-the-art verifier, and defer further details to App. C.

### 4.1 NETWORK ARCHITECTURE AND TIGHTNESS

First, we confirm the predictiveness of our theoretical results on the effect of network width and depth on tightness at initialization and after training. In Figure 3, we visualize tightness at initialization, depending on network width and depth for DLNs and ReLU networks. As predicted by Lemma 3.7 and Corollary 3.8, tightness decreases polynomially with width (see Figure 3 left) and exponentially with depth (see Figure 3 right), both for DLNs and ReLU networks. We confirm our results on the inherent hardness of linear reconstruction in Figure 4, where we plot the ratio of recovered and original box radii, given a bottleneck layer of width  $w$  and synthetic data with intrinsic dimensionality  $k = w$ . As predicted by Theorem 3.10, IBP propagation yields linear and Box\* sublinear growth.




Figure 5: Effect of network depth (top) and width (bottom) on tightness and *training set* IBP-certified accuracy.

Next, we study the interaction of network architecture and IBP training. To this end, we train CNNs with 3 to 13 layers on CIFAR-10 for  $\epsilon = 2/255$ , visualizing results in Figure 5 (right). To quantify the regularizing effect of propagation tightness, we report training set IBP-certified accuracy as a measure of the goodness of fit. Generally, we would expect increased depth to increase capacity and thus decrease the robust training loss and increase training set accuracy. However, we only observe such an increase in accuracy until a depth of 7 layers before accuracy starts to drop. We can explain this by analyzing the corresponding tightness. As expected, tightness is high for shallow networks but decreases quickly with depth, reaching a minimum for 7 layers. From there, tightness increases again, indicating significant regularization, and thereby decreasing accuracy. This is in line with the popularity of the 7-layer CNN7 in the certified training literature (Shi et al., 2021; Müller et al., 2022b).

Continuing our study of architecture effects, we train networks with 0.5 to 16 times the width of a standard CNN7 using IBP training and visualize the resulting IBP certified train set accuracy and tightness in Figure 5 (left). We observe that increasing capacity via width instead of depth yields a monotone although diminishing increase in accuracy as tightness decreases gradually. The different trends for width and depth agree well with our theoretical results, predicting

Table 1: Certified and standard accuracy depending on network width.

Dataset	$\epsilon$	Method	Width	Accuracy	Certified
MNIST	0.1	IBP	1×	98.83	98.10
			4×	98.86	98.23
		SABR	1×	98.99	98.20
	0.3	IBP	1×	97.44	93.26
			4×	97.66	93.35
		SABR	1×	<b>98.82</b>	93.38
CIFAR-10	$\frac{2}{255}$	IBP	1×	67.93	55.85
			2×	68.06	56.18
		IBP-R	1×	78.43	60.87
	$\frac{8}{255}$	IBP	1×	<b>80.46</b>	62.03
			2×	79.24	62.84
		SABR	1×	79.89	<b>63.28</b>
TinyImageNet	$\frac{1}{255}$	IBP	1×	47.35	34.17
			2×	47.83	33.98
		SABR	1×	50.78	34.12
	0.5	IBP	1×	<b>51.56</b>	<b>34.95</b>
			2×	24.47	18.76
		SABR	1×	25.33	19.46
1	IBP	1×	25.40	19.92	
		2×	27.56	20.54	
	SABR	1×	28.63	21.21	
			2×	<b>28.97</b>	<b>21.36</b>




Figure 7: Tightness, standard, and certified accuracy for CNN3 on CIFAR-10, depending on training method and perturbation magnitude  $\epsilon$  used for training and evaluation.

that sufficient network width is essential for trained networks (see Theorem 3.10). It can further be explained by the observation that increasing depth, at initialization, reduces tightness exponentially, while increasing width only reduces it polynomially. Intuitively, this suggests that less regularization is required to offset the tightness penalty of increasing network width rather than depth.

As these experiments indicate that optimal architectures for IBP-based training have only moderate depth but large width, we train wider versions of the popular CNN7 using IBP, SABR, and IBP-R, showing results in Table 1 and Figure 6. We observe that this width increase improves certified accuracy in all settings. We note that, while these improvements might seem marginal, they are of similar magnitude as multiple years of progress on certified training methods, see Figure 6 where CROWN-IBP (Zhang et al., 2020) and MTL-IBP (Palma et al., 2023) (the previous SOTA on MNIST) are shown for reference.




Figure 6: Effect of a 4-fold width increase on certified and standard accuracy for MNIST at  $\epsilon = 0.3$ .

#### 4.2 CERTIFIED TRAINING INCREASES TIGHTNESS

To assess how different training methods affect tightness, we train a CNN3 on CIFAR-10 for a wide range of perturbation magnitudes ( $\epsilon \in [10^{-5}, 5 \cdot 10^{-2}]$ ) using IBP, PGD, and SABR training and illustrate the resulting tightness and accuracies in Figure 7. Recall, that while IBP computes and optimizes a sound over-approximation of the worst-case loss over the whole input region, SABR propagates only a small subregion with IBP, thus yielding an unsound but generally more precise approximation of the worst-case loss. PGD, in contrast, does not use IBP at all but rather trains with samples that approximately maximize the worst-case loss. We observe that training with either IBP-based method increases tightness with perturbation magnitude until networks become almost propagation invariant with  $\tau = 0.98$  (see Figure 7, right). This confirms our theoretical results, showing that IBP training increases tightness with  $\epsilon$  (see Theorem 3.9). In contrast, training with PGD barely influences tightness. Further, the regularization required for such high tightness comes at the cost of standard accuracies being severely reduced (see Figure 7, left). However, while this reduced standard accuracy translates to smaller certified accuracies for very small perturbation magnitudes ( $\epsilon \leq 5 \cdot 10^{-3}$ ), the increased tightness improves certifiability sufficiently to yield higher certified accuracies for larger perturbation magnitudes ( $\epsilon \geq 10^{-2}$ ).

We further investigate this dependency between (certified) robustness and tightness by varying the subselection ratio  $\lambda$  when training with SABR. Recall that  $\lambda$  controls the size of the propagated regions for a fixed perturbation magnitude  $\epsilon$ , recovering IBP for  $\lambda = 1$  and PGD for  $\lambda = 0$ . Plotting results in Figure 8, we observe that while decreasing  $\lambda$ , severely reduces tightness and thus regularization, it not only leads to increasing natural but also certified accuracies until tightness falls below  $\tau < 0.5$  at  $\lambda = 0.4$ . We observe similar trends when varying the regularization level for other unsound certified training methods, discussed in App. D.1. In Figure 9, we vary the perturbation size  $\epsilon$  for three different  $\lambda$  and show tightness over the size of the propagation region  $\xi = \lambda\epsilon$  for a CNN3 and CIFAR-10. Here, we observe that tightness is dominated by the size of the propagation region  $\xi$  and not the robustness specification  $\epsilon$ , indicating that while training with IBP-bounds increases tightness, the resulting high levels of tightness and thus regularization are not generally necessary for robustness. This helps to explain SABR’s success and highlights the potential for developing novel certified training methods that reduce tightness while maintaining sufficient certifiability.








Figure 8: Accuracies and tightness of a CNN7 for CIFAR-10  $\epsilon = \frac{2}{255}$  depending on regularization strength with SABR.

Figure 9: Tightness over propagation region size  $\xi$  for SABR.

To study how certified training methods that do not use IBP-bounds at all (COLT) or only as a regularizer with very small weight (IBP-R) affect tightness, we compare tightness, certified, and standard accuracies on a 4-layer CNN (used by COLT and IBP-R) in Table 2. We observe that the orderings of tightness and accuracy are exactly inverted, highlighting the accuracy penalty of a strong regularization for tightness. While both COLT and IBP-R affect a much smaller increase in tightness than SABR or IBP, they still yield networks an order of magnitude tighter than PGD, suggesting that slightly increased tightness might be desirable for certified robustness. This is further corroborated by the more heavily regularizing SABR outperforming IBP-R at larger  $\epsilon$  while being outperformed at smaller  $\epsilon$ .

Table 2: Tightness and accuracies for various training methods on CIFAR-10.

Method	$\epsilon$	Accuracy	Tightness	Certified
PGD	2/255	81.2	0.001	-
	8/255	69.3	0.007	-
COLT	2/255	78.4*	0.009	60.7*
	8/255	51.7*	0.057	26.7*
IBP-R	2/255	78.2*	0.033	62.0*
	8/255	51.4*	0.124	27.9*
SABR	2/255	75.6	0.182	57.7
	8/255	48.2	0.950	31.2
IBP	2/255	63.0	0.803	51.3
	8/255	42.2	0.977	31.0

\* Literature result.

## 5 RELATED WORK

Baader et al. (2020) show that continuous functions can be approximated by IBP-certifiable ReLU networks up to arbitrary precision. Wang et al. (2022b) extend this result to more activation functions and characterize the hardness of such a construction. Wang et al. (2022a) find that IBP-training converges to a global optimum with high probability for sufficient width. Mirman et al. (2022) show that functions with points of non-invertibility can not be precisely approximated with IBP. Zhu et al. (2022) show that width is advantageous while depth is not for approximate average case robustness.

Shi et al. (2021) define *tightness* as the size of the layerwise  $\text{Box}^\dagger$ , i.e.,  $\Delta = \bar{z}^\dagger - \underline{z}^\dagger$ , rather than its ratio  $\tau$  to the size of the optimal box (Definition 3.6). They thus study the size of the approximation irrespective of the size of the ground truth, while we study the quality of the approximation. This leads to significantly different insights, e.g., propagation tightness  $\tau$  remains the same under scaling of the network weights, while the abstraction size  $\Delta$  is scaled proportionally.

Wu et al. (2021) study the relation between empirical adversarial robustness and network width. They observe that in this setting, increased width actually hurts perturbation stability and thus potentially empirical robustness while improving natural accuracy. In contrast, we have shown theoretically and empirically that width is beneficial for certified robustness when training with IBP-based methods.

## 6 CONCLUSION

Motivated by the recent and surprising dominance of IBP-based certified training methods, we investigate its underlying mechanisms and trade-offs. By quantifying the relationship between IBP and optimal BOX bounds with our novel propagation tightness metric, we are able to predict the influence of architecture choices on deep linear networks at initialization and after training. We experimentally confirm the applicability of these results to ReLU networks and show that wider networks improve the performance of state-of-the-art methods, while deeper networks do not. Finally, we show that IBP-based training methods increase propagation tightness, depending on the size of the propagated region, at the cost of strong regularization. This observation not only helps explain the success of recent certified training methods but, in combination with the novel metric of propagation tightness, might constitute a key step towards developing novel training methods, balancing certifiability and the (over-)regularization resulting from propagation tightness.

## REPRODUCIBILITY STATEMENT

We publish our code, trained models, and detailed instructions on how to reproduce our results at <https://github.com/eth-sri/ibp-propagation-tightness>. Additionally, we provide detailed descriptions of all hyper-parameter choices, data sets, and preprocessing steps in App. C.

## ACKNOWLEDGEMENTS

We would like to thank our anonymous reviewers for their constructive comments and insightful questions.

This work has been done as part of the EU grant ELSA (European Lighthouse on Secure and Safe AI, grant agreement no. 101070617) and the SERI grant SAFEAI (Certified Safe, Fair and Robust Artificial Intelligence, contract no. MB22.00088). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or European Commission. Neither the European Union nor the European Commission can be held responsible for them.

The work has received funding from the Swiss State Secretariat for Education, Research and Innovation (SERI).

## REFERENCES

- Maximilian Baader, Matthew Mirman, and Martin T. Vechev. Universal approximation with certified networks. In *Proc. of ICLR*, 2020.
- Mislav Balunovic and Martin T. Vechev. Adversarial training and provable defenses: Bridging the gap. In *Proc. of ICLR*, 2020.
- Battista Biggio, Iginio Corona, Davide Maiorca, Blaine Nelson, Nedim Srndic, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. Evasion attacks against machine learning at test time. In *Proc of ECML PKDD*, volume 8190, 2013. doi: 10.1007/978-3-642-40994-3\\_25.
- Christopher Brix, Mark Niklas Müller, Stanley Bak, Taylor T. Johnson, and Changliu Liu. First three years of the international verification of neural networks competition (VNN-COMP). *CoRR*, abs/2301.05815, 2023. doi: 10.48550/arXiv.2301.05815.
- Rudy Bunel, Jingyue Lu, Ilker Turkaslan, Philip H. S. Torr, Pushmeet Kohli, and M. Pawan Kumar. Branch and bound for piecewise linear neural network verification. *J. Mach. Learn. Res.*, 21, 2020.
- Jan Chorowski and Jacek M Zurada. Learning understandable neural networks with nonnegative weight constraints. *IEEE transactions on neural networks and learning systems*, 26(1), 2014. doi: 10.1109/TNNLS.2014.2310059.
- JM Cook. Rational formulae for the production of a spherically symmetric probability distribution. *Mathematics of Computation*, 11(58), 1957.
- Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *Proc. of ICML*, volume 119, 2020.
- Claudio Ferrari, Mark Niklas Müller, Nikola Jovanovic, and Martin T. Vechev. Complete verification via multi-neuron relaxation guided branch-and-bound. In *Proc. of ICLR*, 2022.
- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proc. of AISTATS*, volume 9, 2010.
- Sven Gowal, Krishnamurthy Dvijotham, Robert Stanforth, Rudy Bunel, Chongli Qin, Jonathan Uesato, Relja Arandjelovic, Timothy A. Mann, and Pushmeet Kohli. On the effectiveness of interval bound propagation for training verifiably robust models. *ArXiv preprint*, abs/1810.12715, 2018.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proc. of ICCV*, 2015. doi: 10.1109/ICCV.2015.123.
- Ziwei Ji and Matus Telgarsky. Gradient descent aligns the layers of deep linear networks. In *Proc. of ICLR*, 2019.
- Nikola Jovanović, Mislav Balunović, Maximilian Baader, and Martin Vechev. On the paradox of certified training. In *Transactions on Machine Learning Research*, 2022.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.
- Vivian Lin, Radoslav Ivanov, James Weimer, Oleg Sokolsky, and Insup Lee. T4V: exploring neural network architectures that improve the scalability of neural network verification. In *Principles of Systems Design - Essays Dedicated to Thomas A. Henzinger on the Occasion of His 60th Birthday*, volume 13660, 2022. doi: 10.1007/978-3-031-22337-2\\_28.
- Yuhao Mao, Mark Niklas Müller, Marc Fischer, and Martin Vechev. Taps: Connecting certified and adversarial training, 2023.
- George Marsaglia. Choosing a point from the surface of a sphere. *The Annals of Mathematical Statistics*, 43(2), 1972.

- Matthew Mirman, Timon Gehr, and Martin T. Vechev. Differentiable abstract interpretation for provably robust neural networks. In *Proc. of ICML*, volume 80, 2018.
- Matthew B Mirman, Maximilian Baader, and Martin Vechev. The fundamental limits of neural networks for interval certified robustness. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856.
- Mark Niklas Müller, Christopher Brix, Stanley Bak, Changliu Liu, and Taylor T. Johnson. The third international verification of neural networks competition (VNN-COMP 2022): Summary and results. *CoRR*, abs/2212.10376, 2022a. doi: 10.48550/arXiv.2212.10376.
- Mark Niklas Müller, Franziska Eckert, Marc Fischer, and Martin T. Vechev. Certified training: Small boxes are all you need. *CoRR*, abs/2210.04871, 2022b.
- Mark Niklas Müller, Gleb Makarchuk, Gagandeep Singh, Markus Püschel, and Martin T. Vechev. PRIMA: general and precise neural network certification via scalable convex hull approximations. *Proc. ACM Program. Lang.*, 6(POPL), 2022c. doi: 10.1145/3498704.
- Alessandro De Palma, Rudy Bunel, Krishnamurthy Dvijotham, M. Pawan Kumar, and Robert Stanforth. IBP regularization for verified adversarial robustness via branch-and-bound. *ArXiv preprint*, abs/2206.14772, 2022.
- Alessandro De Palma, Rudy Bunel, Krishnamurthy Dvijotham, M. Pawan Kumar, Robert Stanforth, and Alessio Lomuscio. Expressive losses for verified robustness via convex combinations. *CoRR*, abs/2305.13991, 2023. doi: 10.48550/arXiv.2305.13991. URL <https://doi.org/10.48550/arXiv.2305.13991>.
- Iosif Pinelis and Raymond Molzon. Optimal-order bounds on the rate of convergence to normality in the multivariate delta method, 2016.
- Phillip Pope, Chen Zhu, Ahmed Abdelkader, Micah Goldblum, and Tom Goldstein. The intrinsic dimension of images and its impact on learning. In *Proc. of ICLR*, 2021.
- Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. "why should I trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, 2016. doi: 10.1145/2939672.2939778.
- Andrew M. Saxe, James L. McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. In *Proc. of ICLR*, 2014.
- Zhouxing Shi, Yihan Wang, Huan Zhang, Jinfeng Yi, and Cho-Jui Hsieh. Fast certified robust training with short warmup. In *Proc. of NeurIPS*, 2021.
- Gagandeep Singh, Timon Gehr, Matthew Mirman, Markus Püschel, and Martin T. Vechev. Fast and effective robustness certification. In *Proc. of NeurIPS*, 2018.
- Gagandeep Singh, Rupanshu Ganvir, Markus Püschel, and Martin T. Vechev. Beyond the single neuron convex barrier for neural network certification. In *Proc. of NeurIPS*, 2019a.
- Gagandeep Singh, Timon Gehr, Markus Püschel, and Martin T. Vechev. An abstract domain for certifying neural networks. *Proc. ACM Program. Lang.*, 3(POPL), 2019b.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *Proc. of ICLR*, 2014.
- Vincent Tjeng, Kai Y. Xiao, and Russ Tedrake. Evaluating robustness of neural networks with mixed integer programming. In *Proc. of ICLR*, 2019.
- Florian Tramèr, Nicholas Carlini, Wieland Brendel, and Aleksander Madry. On adaptive attacks to adversarial example defenses. In *Proc. of NeurIPS*, 2020.
- Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. In *Proc. of ICLR*, 2019.

- Yihan Wang, Zhouxing Shi, Quanquan Gu, and Cho-Jui Hsieh. On the convergence of certified robust training with interval bound propagation. In *Proc. of ICLR*, 2022a.
- Zi Wang, Aws Albarghouthi, Gautam Prakriya, and Somesh Jha. Interval universal approximation for neural networks. *Proc. ACM Program. Lang.*, 6(POPL), 2022b. doi: 10.1145/3498675.
- Eric Wong, Frank R. Schmidt, Jan Hendrik Metzen, and J. Zico Kolter. Scaling provable adversarial defenses. In *Proc. of NeurIPS*, 2018.
- Boxi Wu, Jinghui Chen, Deng Cai, Xiaofei He, and Quanquan Gu. Do wider neural networks really help adversarial robustness? In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 7054–7067, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/3937230de3c8041e4da6ac3246a888e8-Abstract.html>.
- Lei Wu, Qingcan Wang, and Chao Ma. Global convergence of gradient descent for deep linear residual networks. In *Proc. of NeurIPS*, 2019.
- Huan Zhang, Hongge Chen, Chaowei Xiao, Sven Gowal, Robert Stanforth, Bo Li, Duane S. Boning, and Cho-Jui Hsieh. Towards stable and efficient training of verifiably robust neural networks. In *Proc. of ICLR*, 2020.
- Huan Zhang, Shiqi Wang, Kaidi Xu, Linyi Li, Bo Li, Suman Jana, Cho-Jui Hsieh, and J. Zico Kolter. General cutting planes for bound-propagation-based neural network verification. *ArXiv preprint*, abs/2208.05740, 2022.
- Zhenyu Zhu, Fanghui Liu, Grigorios Chrysos, and Volkan Cevher. Robustness in deep learning: The good (width), the bad (depth), and the ugly (initialization). In *NeurIPS*, 2022. URL [http://papers.nips.cc/paper\\_files/paper/2022/hash/ea5a63f7ddb82e58623693fd1f4933f7-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/ea5a63f7ddb82e58623693fd1f4933f7-Abstract-Conference.html).

## A ADDITIONAL THEORETICAL RESULTS

Below we present a corollary, formalizing the intuitions we provided in Section 3.2.

**Corollary A.1.** *Assume all elements of  $\mathbf{W}^{(1)}$ ,  $\mathbf{W}^{(2)}$  and  $\mathbf{W}^{(2)}\mathbf{W}^{(1)}$  are non-zero and  $\mathbf{W}^{(2)}\mathbf{W}^{(1)}$  is propagation invariant. Then choosing the signs of one row and one column of  $\mathbf{W}^{(2)}\mathbf{W}^{(1)}$  fixes all signs of  $\mathbf{W}^{(2)}\mathbf{W}^{(1)}$ .*

*Proof.* For notational reasons, we define  $W := W^{(2)}W^{(1)}$ . Without loss of generality, assume we know the signs of the first row and the first column, i.e.,  $W_{1,\cdot}$  and  $W_{\cdot,1}$ . We prove via a construction of the signs of all elements. The construction is given by the following: whenever  $\exists i, j$ , such that we know the sign of  $W_{i,j}$ ,  $W_{i,j+1}$  and  $W_{i+1,j}$ , we fix the sign of  $W_{i+1,j+1}$  to be positive if there are an odd number of positive elements among  $W_{i,j}$ ,  $W_{i,j+1}$  and  $W_{i+1,j}$ , otherwise negative.

By Theorem 3.4, propagation invariance requires us to fix the sign of the last element in the  $W_{i:i+1,j:j+1}$  block in this way. We only need to prove that when this process terminates, we fix the signs of all elements. We show this via recursion.

When  $i = 1$  and  $j = 1$ , we have known the signs of  $W_{i,j}$ ,  $W_{i,j+1}$  and  $W_{i+1,j}$ , thus the sign of  $W_{i+1,j+1}$  is fixed. Continuing towards the right, we gradually fix the sign of  $W_{2,j+1}$  for  $j = 1, \dots, d-1$ . Continuing downwards, we gradually fix the sign of  $W_{i+1,2}$  for  $i = 1, \dots, d-1$ . Therefore, all signs of the elements of the second row and the second column are fixed. By recursion, we would finally fix all the rows and the columns, thus the whole matrix.  $\square$

We extend our result in Section 3.3 to two-layer ReLU networks. The intuition is that when the input data is symmetric around zero, ReLU status (activated or not) is independent to weights and the probability of activation is exactly 0.5.

**Corollary A.2.** *Assume the input distribution is symmetric around zero, i.e.,  $p_X(x) = p_X(-x)$  for all  $x > 0$ , and  $P(X = 0) = 0$ . Then for a two-layer ReLU network  $\mathbf{f} = \mathbf{W}^{(2)} \text{ReLU}(\mathbf{W}^{(1)}\mathbf{x})$  initialized with i.i.d. Gaussian, the expected local tightness  $\tau' \sim \sqrt{2}\tau$ , where  $\tau$  is the expected tightness of corresponding deep linear network.*

*Proof.* Since the input  $X$  is symmetric around 0, the distribution of  $\mathbf{W}^{(1)}\mathbf{x}$  is symmetric around 0 as well, regardless of the initialized weights. By assumption on the input and weight distribution,  $P(\mathbf{W}^{(1)}\mathbf{x} = 0) = 0$ , thus  $P(\text{ReLU}(\mathbf{W}^{(1)}\mathbf{x}) = 0) = 0.5$ . In addition, the status of activation is independent to the initialized weights. Thus, the effect can be viewed as randomly setting rows of  $\mathbf{W}^{(1)}$  to zero with probability 0.5. Following Equation (8) and Equation (9), we get that the size of  $\text{Box}^\dagger$  is scaled by 0.5, and the size of  $\text{Box}^*$  is scaled by  $\mathbb{E}(\sqrt{\chi^2(d_1/2)})/\mathbb{E}(\sqrt{\chi^2(d_1)}) \sim \frac{\sqrt{2}}{2}$ . Therefore,  $\tau' \sim \sqrt{2}\tau$ .  $\square$

We perform a Monte-Carlo estimation of the ratio  $\tau'/\tau$  with a two-layer fully connected network and a two-layer convolutional network on MNIST. The estimation is  $1.4167 \pm 0.0059$  and  $1.4228 \pm 0.0368$ , respectively, which is close to the theoretical value  $\sqrt{2} \approx 1.4142$ . This confirms the correctness of our theoretical analysis and its generalization even to convolutional networks which do not fully satisfy the assumption.

## B DEFERRED PROOFS

### Proof of Lemma 3.1

Here we prove Lemma 3.1, restated below for convenience.

**Lemma 3.1.** *Any  $C^0$  continuous classifier  $\mathbf{f}$ , computing the logit difference  $y_i^\Delta := y_i - y_t, \forall i \neq t$ , is robustly correct on  $\mathcal{B}^\epsilon(\mathbf{x})$  if and only if  $\text{Box}^*(\mathbf{f}, \mathcal{B}^\epsilon(\mathbf{x})) \subseteq \mathbb{R}_{<0}^{c-1}$ , i.e.,  $\bar{y}_i^{\Delta*} < 0, \forall i \neq t$ .*

*Proof.* On the one hand, assume  $y_i - y_{\text{true}} < 0$  for all  $i$ . Then for the  $i^{\text{th}}$  output dimension, the optimal bounding box is  $\max y_i - y_{\text{true}}$ . Since the classifier is continuous,  $\mathbf{f}(\mathcal{B}(\mathbf{x}, \epsilon))$  is a closed and

bounded set. Therefore, by extreme value theorem,  $\exists \eta \in \mathcal{B}(\mathbf{x}, \epsilon)$  such that  $\eta = \arg \max y_i - y_{\text{true}}$ , thus  $\max y_i - y_{\text{true}} < 0$ . Since this holds for every  $i$ ,  $\text{Box}^*(\mathbf{f}, \mathcal{B}(\mathbf{x}, \epsilon)) \subseteq \mathcal{R}_{<0}^{K-1}$ .

On the other hand, assume  $\text{Box}^*(\mathbf{f}, \mathcal{B}(\mathbf{x}, \epsilon)) \subseteq \mathcal{R}_{<0}^{K-1}$ . Since  $\mathbf{f}(\mathcal{B}(\mathbf{x}, \epsilon)) \subseteq \text{Box}^*(\mathbf{f}, \mathcal{B}(\mathbf{x}, \epsilon)) \subseteq \mathcal{R}_{<0}^{K-1}$ , we get  $y_i - y_{\text{true}} < 0$  for all  $i$ .  $\square$

### Proof of Theorem 3.2

We first prove Theorem 3.2 for a 2-layer DLN as Lemma B.1.

**Lemma B.1.** *For a two-layer DLN  $\mathbf{f} = \mathbf{W}^{(2)}\mathbf{W}^{(1)}$ ,  $(\bar{\mathbf{z}}^* - \underline{\mathbf{z}}^*)/2 = |\mathbf{W}^{(2)}\mathbf{W}^{(1)}| \epsilon$  and  $(\bar{\mathbf{z}}^\dagger - \underline{\mathbf{z}}^\dagger)/2 = |\mathbf{W}^{(2)}| |\mathbf{W}^{(1)}| \epsilon$ . In addition,  $\text{Box}^*$  and  $\text{Box}^\dagger$  have the same center  $\mathbf{f}(\mathbf{x})$ .*

*Proof.* First, assume  $W^{(1)} \in \mathcal{R}^{d_1 \times d_0}$ ,  $W^{(2)} \in \mathcal{R}^{d_2 \times d_1}$  and  $B_i = [-1, 1]^{d_i}$  for  $i = 0, 1, 2$ , where  $d_i \in \mathcal{Z}_+$  are some positive integers. The input box can be represented as  $\text{diag}(\epsilon_0)B_0 + b$  for  $\epsilon_0 = \epsilon$ .

For a single linear layer, the box propagation yields

$$\begin{aligned} \text{Box}(W^{(1)}(\text{diag}(\epsilon_0)B_0 + b)) &= \text{Box}(W^{(1)} \text{diag}(\epsilon_0)B_0) + W^{(1)}b \\ &= \text{diag} \left( \sum_{j=1}^{d_0} |W_{i,j}^{(1)}| \epsilon_0[j] \right) B_1 + W^{(1)}b \\ &:= \text{diag}(\epsilon_1)B_1 + W^{(1)}b. \end{aligned} \quad (5)$$

Applying Equation (5) iteratively, we get the explicit formula of layer-wise propagation for two-layer linear network:

$$\begin{aligned} &\text{Box}(W^{(2)} \text{Box}(W^{(1)}(\text{diag}(\epsilon_0)B_0 + b))) \\ &= \text{Box} \left( W^{(2)}(\text{diag}(\epsilon_1)B_1 + W^{(1)}b) \right) \\ &= \text{diag} \left( \sum_{k=1}^{d_1} |W_{i,k}^{(2)}| \epsilon_1[k] \right) B_2 + W^{(2)}W^{(1)}b \\ &= \text{diag} \left( \sum_{j=1}^{d_0} \epsilon_0[j] \left( \sum_{k=1}^{d_1} |W_{i,k}^{(2)} W_{k,j}^{(1)}| \right) \right) B_2 + W^{(2)}W^{(1)}b. \end{aligned} \quad (6)$$

Applying Equation (5) on  $W := W^{(2)}W^{(1)}$ , we get the explicit formula of the tightest box:

$$\begin{aligned} &\text{Box}(W^{(2)}W^{(1)}(\text{diag}(\epsilon_0)B_0 + b)) \\ &= \text{diag} \left( \sum_{j=1}^{d_0} |(W^{(2)}W^{(1)})_{i,j}| \epsilon_0[j] \right) B_2 + W^{(2)}W^{(1)}b \\ &= \text{diag} \left( \sum_{j=1}^{d_0} \epsilon_0[j] \left| \sum_{k=1}^{d_1} W_{i,k}^{(2)} W_{k,j}^{(1)} \right| \right) B_2 + W^{(2)}W^{(1)}b. \end{aligned} \quad (7)$$

$\square$

Now, we use induction and Lemma B.1 to prove Theorem 3.2, restated below for convenience. The key insight is that a multi-layer DLN is equivalent to a single-layer linear network. Thus, we can group layers together and view general DLNs as two-layer DLNs.

**Theorem 3.2 (Box Propagation).** *For an  $L$ -layer DLN  $\mathbf{f} = \prod_{k=1}^L \mathbf{W}^{(k)}$ , we obtain the box centres  $\bar{\mathbf{z}}^* = \underline{\mathbf{z}}^* = \mathbf{f}(\mathbf{x})$  and the radii*

$$\frac{\bar{\mathbf{z}}^* - \underline{\mathbf{z}}^*}{2} = \left| \prod_{k=1}^L \mathbf{W}^{(k)} \right| \epsilon, \quad \text{and} \quad \frac{\bar{\mathbf{z}}^\dagger - \underline{\mathbf{z}}^\dagger}{2} = \left( \prod_{k=1}^L \left| \mathbf{W}^{(k)} \right| \right) \epsilon. \quad (4)$$

*Proof.* For  $L = 2$ , by Lemma B.1, the result holds. Assume for  $L \leq m$ , the result holds. Therefore, for  $L = m + 1$ , we group the first  $m$  layers as a single layer, resulting in a “two” layer equivalent network. Thus,  $(\bar{\mathbf{z}}^* - \mathbf{z}^*)/2 = |\mathbf{W}^{(m+1)} \prod_{k=1}^m \mathbf{W}^{(k)}| \boldsymbol{\epsilon} = |\prod_{k=1}^L \mathbf{W}^{(k)}| \boldsymbol{\epsilon}$ . Similarly, by Equation (5), we can prove  $(\bar{\mathbf{z}}^* - \mathbf{z}^*)/2 = (|\mathbf{W}^{(m+1)}| \prod_{k=1}^m |\mathbf{W}^{(k)}|) \boldsymbol{\epsilon} = (\prod_{k=1}^L |\mathbf{W}^{(k)}|) \boldsymbol{\epsilon}$ . The claim about center follows by induction similarly.  $\square$

### Proof of Lemma 3.3

Here, we prove Lemma 3.3, restated below for convenience.

**Lemma 3.3** (Propagation Invariance). *A two-layer DLN  $\mathbf{f} = \mathbf{W}^{(2)} \mathbf{W}^{(1)}$  is propagation invariant if and only if for every fixed  $(i, j)$ , we have  $|\sum_k W_{i,k}^{(2)} \cdot W_{k,j}^{(1)}| = \sum_k |W_{i,k}^{(2)} \cdot W_{k,j}^{(1)}|$ , i.e., either  $W_{i,k}^{(2)} \cdot W_{k,j}^{(1)} \geq 0$  for all  $k$  or  $W_{i,k}^{(2)} \cdot W_{k,j}^{(1)} \leq 0$  for all  $k$ .*

*Proof.* We prove the statement via comparing the box bounds. By Lemma B.1, we need  $|\sum_{k=1}^{d_1} W_{i,k}^{(2)} W_{k,j}^{(1)}| = \sum_{k=1}^{d_1} |W_{i,k}^{(2)} W_{k,j}^{(1)}|$ . The triangle inequality of absolute function says this holds if and only if  $W_{i,k}^{(2)} W_{k,j}^{(1)} \geq 0$  for all  $k$  or  $W_{i,k}^{(2)} W_{k,j}^{(1)} \leq 0$  for all  $k$ .  $\square$

### Proof of Theorem 3.4

Here, we prove Theorem 3.4, restated below for convenience.

**Theorem 3.4** (Non-Propagation Invariance). *Assume  $\exists i, i', j, j'$ , such that  $W_{i,j}^{(1)}, W_{i',j'}^{(1)}, W_{i,j}^{(2)}$  and  $W_{i',j'}^{(2)}$  are all non-zero. If  $(\mathbf{W}^{(2)} \mathbf{W}^{(1)})_{i,j} \cdot (\mathbf{W}^{(2)} \mathbf{W}^{(1)})_{i',j'} \cdot (\mathbf{W}^{(2)} \mathbf{W}^{(1)})_{i',j} \cdot (\mathbf{W}^{(2)} \mathbf{W}^{(1)})_{i,j'} < 0$ , then  $\mathbf{f} = \mathbf{W}^{(2)} \mathbf{W}^{(1)}$  is not propagation invariant.*

*Proof.* The assumption  $(\mathbf{W}^{(2)} \mathbf{W}^{(1)})_{i,j} \cdot (\mathbf{W}^{(2)} \mathbf{W}^{(1)})_{i',j'} \cdot (\mathbf{W}^{(2)} \mathbf{W}^{(1)})_{i',j} \cdot (\mathbf{W}^{(2)} \mathbf{W}^{(1)})_{i,j'} < 0$  implies three elements are of the same sign while the other element has a different sign. Without loss of generality, assume  $(\mathbf{W}^{(2)} \mathbf{W}^{(1)})_{i',j'} < 0$  and the rest three are all positive.

Assume  $\mathbf{W}^{(2)} \mathbf{W}^{(1)}$  is propagation invariant. By Lemma 3.3, this means  $W_{i,j}^{(2)} \cdot \text{sign} = W_{i',j'}^{(1)} \cdot \text{sign}$ ,  $W_{i,j}^{(2)} \cdot \text{sign} = W_{i',j'}^{(1)} \cdot \text{sign}$ ,  $W_{i',j}^{(2)} \cdot \text{sign} = W_{i,j}^{(1)} \cdot \text{sign}$  and  $W_{i',j'}^{(2)} \cdot \text{sign} = -W_{i,j}^{(1)} \cdot \text{sign}$ . Therefore, we have  $-W_{i',j'}^{(1)} \cdot \text{sign} = W_{i,j}^{(1)} \cdot \text{sign}$ , which implies all elements of  $W_{i',j'}^{(1)}$  must be zero. However, this results in  $(\mathbf{W}^{(2)} \mathbf{W}^{(1)})_{i,j'} = 0$ , a contradiction.  $\square$

### Proof of Lemma 3.7

Here, we prove Lemma 3.7, restated below for convenience.

**Lemma 3.7** (Initialization Tightness w.r.t. Width). *Given a 2-layer DLN with weight matrices  $\mathbf{W}^{(1)} \in \mathcal{R}^{d_1 \times d_0}$ ,  $\mathbf{W}^{(2)} \in \mathcal{R}^{d_2 \times d_1}$  with i.i.d. entries from  $\mathcal{N}(0, \sigma_1^2)$  and  $\mathcal{N}(0, \sigma_2^2)$  (together denoted as  $\boldsymbol{\theta}$ ), we obtain the expected tightness  $\tau(d_1) = \frac{\mathbb{E}_{\boldsymbol{\theta}}(\bar{\mathbf{z}}^* - \mathbf{z}^*)}{\mathbb{E}_{\boldsymbol{\theta}}(\bar{\mathbf{z}}^\dagger - \mathbf{z}^\dagger)} = \frac{\sqrt{\pi} \Gamma(\frac{1}{2}(d_1+1))}{d_1 \Gamma(\frac{1}{2} d_1)} \in \Theta(\frac{1}{\sqrt{d_1}})$ .*

*Proof.* We first compute the size of the layer-wisely propagated box. From Equation (6), we get that for the  $i$ -th dimension,

$$\begin{aligned} \mathbb{E}(u_i - l_i) &= \mathbb{E} \left( \sum_{j=1}^{d_0} \epsilon_0[j] \left( \sum_{k=1}^{d_1} |W_{i,k}^{(2)} W_{k,j}^{(1)}| \right) \right) \\ &= \sum_{j=1}^{d_0} \epsilon_0[j] \left( \sum_{k=1}^{d_1} \mathbb{E}(|W_{i,k}^{(2)}|) \cdot \mathbb{E}(|W_{k,j}^{(1)}|) \right) \\ &= \sigma_1 \sigma_2 \sum_{j=1}^{d_0} \epsilon_0[j] \left( \sum_{k=1}^{d_1} \mathbb{E}(|\mathcal{N}(0, 1)|)^2 \right). \end{aligned}$$



Since  $\mathbb{E}(|\mathcal{N}(0, 1)|) = \sqrt{\frac{2}{\pi}}$ , we have

$$\mathbb{E}(u_i - l_i) = \frac{2}{\pi} \sigma_1 \sigma_2 d_1 \|\epsilon_0\|_1. \quad (8)$$

Now we compute the size of the tightest box. From Equation (7), we get that for the  $i$ -th dimension,

$$\mathbb{E}(u_i^* - l_i^*) = \mathbb{E} \left( \sum_{j=1}^{d_0} \epsilon_0[j] \left| \sum_{k=1}^{d_1} W_{i,k}^{(2)} W_{k,j}^{(1)} \right| \right) = \sigma_1 \sigma_2 \sum_{j=1}^{d_0} \epsilon_0[j] \mathbb{E} \left( \left| \sum_{k=1}^{d_1} X_k Y_k \right| \right),$$

where  $X_k$  and  $Y_k$  are i.i.d. standard Gaussian random variables. Using the law of total expectation, we have

$$\begin{aligned} \mathbb{E} \left( \left| \sum_{k=1}^{d_1} X_k Y_k \right| \right) &= \mathbb{E} \left( \mathbb{E} \left( \left| \sum_{k=1}^{d_1} X_k Y_k \right| \middle| Y_k \right) \right) \\ &= \mathbb{E} \left( \mathbb{E} \left( \left| \mathcal{N}(0, \sum_{k=1}^{d_1} Y_k^2) \right| \middle| Y_k \right) \right) \\ &= \sqrt{\frac{2}{\pi}} \mathbb{E} \left( \sqrt{\sum_{k=1}^{d_1} Y_k^2} \right) \\ &= \sqrt{\frac{2}{\pi}} \mathbb{E}(\sqrt{\chi^2(d_1)}). \end{aligned}$$

Since  $\mathbb{E}(\sqrt{\chi^2(d_1)}) = \sqrt{2} \Gamma(\frac{1}{2}(d_1 + 1)) / \Gamma(\frac{1}{2}d_1)$ ,<sup>2</sup> we have

$$\mathbb{E}(u_i^* - l_i^*) = \frac{2}{\sqrt{\pi}} \sigma_1 \sigma_2 \|\epsilon_0\|_1 \Gamma(\frac{1}{2}(d_1 + 1)) / \Gamma(\frac{1}{2}d_1). \quad (9)$$

Combining Equation (8) and Equation (9), we have:

$$\frac{\mathbb{E}(u_i - l_i)}{\mathbb{E}(u_i^* - l_i^*)} = \frac{d_1 \Gamma(\frac{1}{2}d_1)}{\sqrt{\pi} \Gamma(\frac{1}{2}(d_1 + 1))}. \quad (10)$$

To see the asymptotic behavior, use  $\Gamma(x + \alpha) / \Gamma(x) \sim x^\alpha$ ,<sup>3</sup> we have

$$\frac{\mathbb{E}(u_i - l_i)}{\mathbb{E}(u_i^* - l_i^*)} \sim \frac{1}{\sqrt{\pi}} d_1^{\frac{1}{2}}. \quad (11)$$

To establish the bounds on the minimum expected slackness, we use Lemma B.2.  $\square$

**Lemma B.2.** Let  $g(n) := \frac{n \Gamma(\frac{1}{2}n)}{\sqrt{\pi} \Gamma(\frac{1}{2}(n+1))}$ .  $g(n)$  is monotonically increasing for  $n \geq 1$ . Thus, for  $n \geq 2$ ,  $g(n) \geq g(2) > 1.27$ .

*Proof.* Using polygamma function  $\psi^{(0)}(z) = \Gamma'(z) / \Gamma(z)$ ,<sup>4</sup> we have


$$g'(n) \propto 1 + \frac{1}{2}n \left( \psi^{(0)} \left( \frac{1}{2}n \right) - \psi^{(0)} \left( \frac{1}{2}(n+1) \right) \right).$$

<sup>1</sup>[https://en.wikipedia.org/wiki/Half-normal\\_distribution](https://en.wikipedia.org/wiki/Half-normal_distribution)

<sup>2</sup>[https://en.wikipedia.org/wiki/Chi\\_distribution](https://en.wikipedia.org/wiki/Chi_distribution)

<sup>3</sup>[https://en.wikipedia.org/wiki/Gamma\\_function#Stirling's\\_formula](https://en.wikipedia.org/wiki/Gamma_function#Stirling's_formula)

<sup>4</sup>[https://en.wikipedia.org/wiki/Polygamma\\_function](https://en.wikipedia.org/wiki/Polygamma_function)

Figure 10:  $g(n)$  and  $g^2(n)$  visualized.

Using the fact that  $\psi^{(0)}(z)$  is monotonically increasing for  $z > 0$  and  $\psi^{(0)}(z+1) = \psi^{(0)}(z) + \frac{1}{z}$ , we have

$$\begin{aligned}
 & 1 + \frac{1}{2}n \left( \psi^{(0)}\left(\frac{1}{2}n\right) - \psi^{(0)}\left(\frac{1}{2}(n+1)\right) \right) \\
 & > 1 + \frac{1}{2}n \left( \psi^{(0)}\left(\frac{1}{2}n\right) - \psi^{(0)}\left(\frac{1}{2}n+1\right) \right) \\
 & = 1 + \frac{1}{2}n \left( -\frac{2}{n} \right) \\
 & = 0.
 \end{aligned}$$

Therefore,  $g'(n)$  is strictly positive for  $n \geq 1$ , and thus  $g(n)$  is monotonically increasing for  $n \geq 1$ .  $\square$

As a final comment, we visualize  $g(n)$  in Figure 10. As expected,  $g(n)$  is monotonically increasing in the order of  $O(\sqrt{n})$ .

### Proof of Corollary 3.8

Here, we prove Corollary 3.8, restated below for convenience.

**Corollary 3.8** (Initialization Tightness w.r.t. Depth). *The expected tightness of an  $L$ -layer DLN  $\mathbf{f}$  with minimum internal dimension  $d_{\min}$  is at most  $\tau \leq \tau(d_{\min})^{\lfloor \frac{L}{2} \rfloor}$  at initialization.*

*Proof.* This is pretty straightforward and only requires a coarse application of Lemma 3.7. Without loss of generality, we assume  $L$  is even. If  $L$  is odd, then we simply discard the slackness introduced by the last layer, *i.e.*, assume the last layer does not introduce additional slackness.

We group the  $2i-1$ -th and  $2i$ -th layer as a new layer. By Lemma 3.7, these  $L/2$  subnetworks all introduce an additional slackness factor of  $\tau$ . Note that Equation (8) implies that the size of the output box is proportional to the size of the input box. Therefore, the layer-wisely propagated box of  $\prod_{i=1}^L W_i$  is  $\tau^{L/2}$  looser than the layer-wisely propagated box of  $\prod_{j=1}^{L/2} (W_{2j-1}W_{2j})$ . In addition, the size of the tightest box for  $\prod_{i=1}^L W_i$  is upper bounded by layer-wisely propagating  $\prod_{j=1}^{L/2} (W_{2j-1}W_{2j})$ . Therefore, the minimum expected slackness is lower bounded by  $\tau^{L/2}$ .  $\square$

### Proof of Theorem 3.9

Here, we prove Theorem 3.9, restated below for convenience.

**Theorem 3.9** (IBP Training Increases Tightness). *Assume homogenous tightness, *i.e.*,  $\mathbf{W}^* = \tau\mathbf{W}^\dagger$ , and  $\frac{\|\nabla_\theta \mathbf{W}_{ij}^*\|_2}{\mathbf{W}_{ij}^*} \leq \frac{1}{2} \frac{\|\nabla_\theta \mathbf{W}_{ij}^\dagger\|_2}{\mathbf{W}_{ij}^\dagger}$  for all  $i, j$ , then, the gradient difference between the IBP and standard loss is aligned with an increase in tightness, *i.e.*,  $\langle \nabla_\theta (R(\epsilon) - R(0)), \nabla_\theta \tau \rangle \leq 0$  for all  $\epsilon > 0$ .*

*Proof.* We prove a stronger claim:  $\langle \nabla_\theta (R(\epsilon + \Delta\epsilon) - R(\epsilon)), \nabla_\theta \tau \rangle \leq 0$  for all  $\epsilon \geq 0$  and  $\Delta\epsilon > 0$ . Let  $\epsilon = \mathbf{0}$  yields the theorem.

We prove the claim for  $\Delta\epsilon \rightarrow 0$ . For large  $\Delta\epsilon$ , we can break it into  $R(\epsilon + \Delta\epsilon) - R(\epsilon) = \sum_{i=1}^n R(\epsilon + \frac{i}{n}\Delta\epsilon) - R(\epsilon + \frac{i-1}{n}\Delta\epsilon)$ , thus proving the claim since each summand satisfies the theorem.

Let  $L_1 = R(\epsilon)$  and  $L_2 = R(\epsilon + \Delta\epsilon)$ . By Taylor expansion, we have  $L_2 = L_1 + \Delta\epsilon^\top \mathbf{W}^\dagger \nabla_{\mathbf{u}} g = L_1 + \frac{1}{\tau} \Delta\epsilon^\top \mathbf{W}^* \nabla_{\mathbf{u}} g$ , where  $\nabla_{\mathbf{u}} g = \nabla_{\mathbf{u}} g(\mathbf{u})$  evaluated at  $\mathbf{u} = \mathbf{W}^\dagger \epsilon$ . Note that the increase of  $\epsilon$  would increase the risk, thus  $\nabla_{\mathbf{u}} g \geq \mathbf{0}$ .

For the  $i^{\text{th}}$  parameter  $\theta_i$ ,  $\nabla_{\theta_i}(L_2 - L_1) \nabla_{\theta_i} \tau = \frac{1}{\tau^2} \Delta\epsilon^\top (\tau \nabla_{\theta_i} \mathbf{W}^* - \mathbf{W}^* \nabla_{\theta_i} \tau) \nabla_{\mathbf{u}} g \nabla_{\theta_i} \tau$ . Thus,  $\langle \nabla_{\theta}(L_2 - L_1), \nabla_{\theta} \tau \rangle = \frac{1}{\tau^2} \Delta\epsilon^\top (\tau \sum_i \nabla_{\theta_i} \tau \cdot \nabla_{\theta_i} \mathbf{W}^* - \mathbf{W}^* \|\nabla_{\theta} \tau\|_2^2) \nabla_{\mathbf{u}} g$ . Since  $\Delta\epsilon > \mathbf{0}$  and  $\nabla_{\mathbf{u}} g \geq \mathbf{0}$ , it suffices to prove that  $\tau \sum_i \nabla_{\theta_i} \tau \cdot \nabla_{\theta_i} \mathbf{W}^* - \mathbf{W}^* \|\nabla_{\theta} \tau\|_2^2$  is nonpositive, *i.e.*,  $\tau \langle \nabla_{\theta} \tau, \nabla_{\theta} \mathbf{W}_{ij}^* \rangle - \mathbf{W}_{ij}^* \|\nabla_{\theta} \tau\|_2^2$  is nonpositive for every  $i, j$ .

Since  $\|\mathbf{u}\|_2 \|\mathbf{v}\|_2 \geq \langle \mathbf{u}, \mathbf{v} \rangle$ , we have

$$\begin{aligned} \frac{\|\nabla_{\theta} \mathbf{W}_{ij}^*\|_2}{\mathbf{W}_{ij}^*} &\leq \frac{1}{2} \frac{\|\nabla_{\theta} \mathbf{W}_{ij}^\dagger\|_2}{\mathbf{W}_{ij}^\dagger} \\ \Rightarrow \|\nabla_{\theta} \log \mathbf{W}^\dagger\|_2 &\geq 2 \|\nabla_{\theta} \log \mathbf{W}^*\|_2 \\ \Rightarrow \|\nabla_{\theta} \log \mathbf{W}^\dagger\|_2^2 &\geq 2 \langle \nabla_{\theta} \log \mathbf{W}^\dagger, \nabla_{\theta} \log \mathbf{W}^* \rangle \end{aligned}$$

Therefore,  $\|\nabla_{\theta} \log \tau\|_2^2 = \|\nabla_{\theta}(\log \mathbf{W}_{ij}^* - \log \mathbf{W}_{ij}^\dagger)\|_2^2 = \|\nabla_{\theta} \log \mathbf{W}_{ij}^*\|_2^2 - 2 \langle \nabla_{\theta} \log \mathbf{W}^\dagger, \nabla_{\theta} \log \mathbf{W}^* \rangle + \|\nabla_{\theta} \log \mathbf{W}^\dagger\|_2^2 \geq \|\nabla_{\theta} \log \mathbf{W}_{ij}^*\|_2^2$ . This means  $\frac{\|\nabla_{\theta} \mathbf{W}_{ij}^*\|_2}{\mathbf{W}_{ij}^*} \leq \frac{\|\nabla_{\theta} \tau\|_2}{\tau}$ , thus  $\mathbf{W}_{ij}^* \|\nabla_{\theta} \tau\|_2^2 \geq \tau \|\nabla_{\theta} \tau\|_2 \|\nabla_{\theta} \mathbf{W}_{ij}^*\|_2 \geq \tau \langle \nabla_{\theta} \tau, \nabla_{\theta} \mathbf{W}_{ij}^* \rangle$ , which fulfills our goal.  $\square$

### Proof of Theorem 3.10

Here, we prove Theorem 3.10, restated below for convenience.

**Theorem 3.10** (Box Reconstruction Error). *Consider the linear embedding and reconstruction  $\hat{\mathbf{x}} = U_k U_k^\top \mathbf{x}$  of a  $d$  dimensional data distribution  $\mathbf{x} \sim \mathcal{X}$  into a  $k$  dimensional space with  $d \gg k$  and eigenmatrices  $U$  drawn uniformly at random from the orthogonal group. Propagating the input box  $\mathcal{B}^\epsilon(\mathbf{x})$  layer-wise and optimally, thus, yields  $\mathcal{B}^{\delta_1}(\hat{\mathbf{x}})$ , and  $\mathcal{B}^{\delta^*}(\hat{\mathbf{x}})$ , respectively. Then, we have, (i)  $\mathbb{E}(\delta_i/\epsilon) = ck \in \Theta(k)$  for a positive constant  $c$  depending solely on  $d$  and  $c \rightarrow \frac{2}{\pi} \approx 0.64$  for large  $d$ ; and (ii)  $\mathbb{E}(\delta_i^*/\epsilon) \rightarrow \frac{2}{\sqrt{\pi}} \frac{\Gamma(\frac{1}{2}(k+))}{\Gamma(\frac{1}{2}k)} \in \Theta(\sqrt{k})$ .*

*Proof.* Since box propagation for linear functions maps the center of the input box to the center of the output box, the center of the output box is exactly  $\hat{X}$ . By Lemma B.1, we have  $\delta = |U_k| |U_k|^\top \epsilon \mathbf{1}$ . For notational simplicity, let  $V = |U_k|$ , thus

$$\delta_i = \sum_{j=1}^k V_{ij} \left( \sum_{p=1}^d V_{jp}^\top \epsilon \right) = \epsilon \sum_{p=1}^d \sum_{j=1}^k V_{ij} V_{pj} = \epsilon \sum_{j=1}^k V_{ij} \|V_{:j}\|_1.$$

Therefore,  $\mathbb{E} \delta_i / \epsilon = \sum_{j=1}^k \mathbb{E}(V_{ij} \|V_{:j}\|_1) = ck$ , where  $c = \mathbb{E}(V_{ij} \|V_{:j}\|_1)$ . Since  $V_{:j}$  is the absolute value of a column of the orthogonal matrix uniformly drawn,  $V_{:j}$  itself is the absolute value of a vector drawn uniformly from the unit hyper-ball. By Cook (1957) and Marsaglia (1972),  $V_{:j}$  is equivalent in distribution to *i.i.d.* draw samples from the standard Gaussian for each dimension and then normalize it by its  $L_2$  norm. For notational simplicity, let  $V_{:j} \stackrel{d}{=} v = |u|$ , where  $u = \hat{u} / \|\hat{u}\|_2$  and all dimensions of  $\hat{u}$  are *i.i.d.* drawn from the standard Gaussian distribution, thus  $c = \mathbb{E}(v_1 \|v\|_1)$ .

Expanding  $\|v\|_1$ , we have  $c = \mathbb{E}(v_1^2) + \sum_{i=2}^d \mathbb{E}(v_1 v_i) = \frac{1}{d} \mathbb{E}(\|v\|_2^2) + (d-1) \mathbb{E}(v_1 v_2) = \frac{1}{d} + (d-1) \mathbb{E}(v_1 v_2)$ . From page 20 of Pinelis & Molzon (2016), we know each entry in  $u$  converges to  $\mathcal{N}(0, 1/d)$  at  $O(1/d)$  speed in Kolmogorov distance. In addition,  $\mathbb{E}(v_1 v_2) = \mathbb{E}(\mathbb{E}(v_2 | v_1) \cdot v_1) = \mathbb{E}(v_1 \sqrt{1-v_1^2}) \mathbb{E}(v_2')$ , where  $v'$  is the absolute value of a random vector uniformly drawn from the  $d-1$  dimensional sphere. Therefore, for large  $d$ ,  $c = (d-1) \mathbb{E}(v_1 \sqrt{1-v_1^2}) \mathbb{E}(v_2') = (d-1) \mathbb{E}(v_1) \mathbb{E}(v_2') = (d-1) \mathbb{E}(|\mathcal{N}(0, 1/d)|) \mathbb{E}(|\mathcal{N}(0, 1/(d-1))|) = \frac{2}{\pi}$ .




Figure 11: Monte-Carlo estimations of Theorem 3.10. Result bases on 10000 samples for each  $d$ . Left:  $c$  plotted against  $d$  in log scale. Right:  $\mathbb{E}(\delta_i^*)$  plotted against  $k$  for  $d = 2000$  (blue), together with the theoretical predictions (orange).

To show how good the asymptotic result is, we run Monte-Carlo to get the estimation of  $c$ . As shown in the left of Figure 11, the Monte-Carlo result is consistent to this theorem. In addition, it converges very quickly, *e.g.*, stabilizing at 0.64 when  $d \geq 100$ .

Now we start proving (2). By Lemma B.1, we have  $\delta^* = |U_k U_k^\top| \epsilon \mathbf{1}$ . Thus,

$$\mathbb{E}(\delta_i^* / \epsilon) = \sum_{j=1}^d \mathbb{E} \left| \sum_{p=1}^k U_{ip} U_{jp} \right| = \sum_{j \neq i} \mathbb{E} \left| \sum_{p=1}^k U_{ip} U_{jp} \right| + \mathbb{E} \left( \sum_{p=1}^k U_{ip}^2 \right) = (d-1) \mathbb{E} \left| \sum_{p=1}^k U_{ip} U_{jp} \right| + \frac{k}{d}.$$

In addition, we have

$$\begin{aligned} (d-1) \mathbb{E} \left| \sum_{p=1}^k U_{ip} U_{jp} \right| &= (d-1) \mathbb{E}_{U_i} \left( \mathbb{E}_{U_j} \left( \left| \sum_{p=1}^k U_{ip} U_{jp} \right| \middle| U_i \right) \right) \\ &\rightarrow (d-1) \mathbb{E}_{U_i} \left( \mathbb{E} \left| \mathcal{N} \left( 0, \frac{\sum_{p=1}^k U_{ip}^2}{d-1} \right) \right| \right) \\ &= (d-1) \sqrt{\frac{2}{\pi(d-1)}} \mathbb{E} \sqrt{\sum_{p=1}^k U_{ip}^2} \\ &= \sqrt{\frac{2(d-1)}{\pi}} \mathbb{E} \sqrt{\frac{1}{d} \chi^2(k)} \\ &\rightarrow \frac{2}{\sqrt{\pi}} \frac{\Gamma(\frac{1}{2}(k+1))}{\Gamma(\frac{1}{2}k)}, \end{aligned}$$

where we use again that for large  $d$ , the entries of a column tends to Gaussian. This proves (2). The expected tightness follows by definition, *i.e.*, dividing the result of (1) and (2).  $\square$

The right of Figure 11 plots the Monte-Carlo estimations against our theoretical results. Clearly, this confirms our result.

## C EXPERIMENTAL DETAILS

### C.1 DATASET

We use the MNIST (LeCun et al., 2010) and CIFAR-10 (Krizhevsky et al., 2009) datasets for our experiments. Both are open-source and freely available. For MNIST, we do not apply any preprocessing or data augmentation. For CIFAR-10, we normalize images with their mean and standard deviation and, during training, first apply 2-pixel zero padding and then random cropping to  $32 \times 32$ .

### C.2 MODEL ARCHITECTURE

We follow previous works (Shi et al., 2021; Müller et al., 2022b; Mao et al., 2023) and use a 7-layer convolutional network CNN7 in most experiments. We also use a simplified 3-layer convolutional network CNN3 in Section 4.2. Details about them can be found in the released code.




Figure 12: Accuracies and tightness of a CNN7 for CIFAR-10  $\epsilon = \frac{2}{255}$  depending on regularization strength with STAPS.

### C.3 TRAINING

Following previous works (Müller et al., 2022b; Mao et al., 2023), we use the initialization, warm-up regularization, and learning schedules introduced by Shi et al. (2021). Specifically, for MNIST, the first 20 epochs are used for  $\epsilon$ -scheduling, increasing  $\epsilon$  smoothly from 0 to the target value. Then, we train an additional 50 epochs with two learning rate decays of 0.2 at epochs 50 and 60, respectively. For CIFAR-10, we use 80 epochs for  $\epsilon$ -annealing, after training models with standard training for 1 epoch. We continue training for 80 further epochs with two learning rate decays of 0.2 at epochs 120 and 140, respectively. The initial learning rate is  $5 \times 10^{-3}$  and the gradients are clipped to an  $L_2$  norm of at most 10.0 before every step.

### C.4 CERTIFICATION

We apply MN-BAB (Ferrari et al., 2022), a state-of-the-art (Brix et al., 2023; Müller et al., 2022a) verifier based on multi-neuron constraints (Müller et al., 2022c; Singh et al., 2019a) and the branch-and-bound paradigm (Bunel et al., 2020) to certify all models. MN-BAB is a state-of-the-art complete certification method built on multi-neuron relaxations. For Table 1, we use the same hyperparameters for MN-BAB as Müller et al. (2022b) and set the timeout to 1000 seconds. For other experiments, we use the same hyperparameters but reduce timeout to 200 seconds for efficiency reasons.

## D EXTENDED EMPIRICAL EVALUATION

### D.1 STAPS-TRAINING AND REGULARIZATION LEVEL

To confirm our observations on the interaction of regularization level, accuracies, and propagation tightness from Section 4.2, we extend our experiments to STAPS (Mao et al., 2023), an additional state-of-the-art certified training method beyond SABR (Müller et al., 2022b). Recall that STAPS combines SABR with adversarial training as follows. The model is first (conceptually) split into a feature extractor and classifier. Then, during training IBP is used to propagate the input region through the feature extractor yielding box bounds in the model’s latent space. Then, adversarial training with PGD is conducted over the classifier using these box bounds as input region. As IBP leads to an over-approximation while PGD leads to an under-approximation, STAPS induces more regularization as fewer (ReLU) layers are included in the classifier.

We visualize the result of thus varying regularization levels by changing the number of ReLU layers in the classifier in Figure 12. We observe very similar trends as for SABR in Figure 8, although to a lesser extent, as 0 ReLU layers in the classifier still recovers SABR and not standard IBP. Again, decreasing regularization (increasing the number of ReLU layers in the classifier) leads to reducing tightness and increasing standard and certified accuracies.

### D.2 TIGHTNESS AND PROPAGATION REGION SIZE

We repeat the experiment illustrated in Figure 9 for CIFAR-10 on MNIST using a CNN3 in Figure 13. We again observe the propagation region size  $\xi$  dominating the tightness (except for very large perturbation sizes of  $\epsilon > 0.2$ ), and smaller perturbation magnitudes leading to slightly larger tightness.

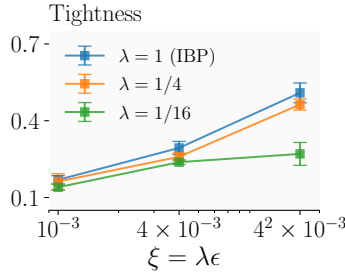


Figure 13: Tightness over propagation region size  $\xi$  for SABR and MNIST.

### D.3 TIGHTNESS APPROXIMATION ERROR

To investigate the approximation quality of our local tightness as defined in Definition 3.6, we compare it against the true tightness computed using MILP (Tjeng et al., 2019). We confirm our results from Figure 2 on MNIST on CIFAR-10 in Figure 14, where we again observe small approximation errors across a wide range of perturbation magnitudes. Interestingly, the effect of the chosen perturbation magnitude on the approximation error is less pronounced than on MNIST, remaining low even for large perturbation magnitudes ( $\epsilon = 0.01 > 8/255$ ). While the approximation error remains below 0.3% for a PGD-trained net, our approximation exhibits a consistent bias for the IBP-trained network, overestimating tightness by approximately 1.2%.




Figure 14: Mean relative error between local tightness (Definition 3.6) and true tightness computed with MILP for a CNN3 trained with PGD or IBP at  $\epsilon = 0.005$  on CIFAR-10.

### D.4 COMPARING TIGHTNESS TO (INVERSE) ROBUST CROSS-ENTROPY LOSS

To investigate to what extent our novel tightness metric is complimentary to the (inverse) robust cross-entropy loss (see Section 2) computed with IBP, we repeat the key experiments confirming our theoretical insights with the inverse IBP-loss and observe significantly different, partially opposite trends.

**IBP Loss at Initialization** We repeat our experiments on the dependence of tightness at initialization on network depth and width, illustrated in Figure 5, and additionally report the inverse IBP loss in Figure 15. For all experiments, we use the initialization of Shi et al. (2021) which has become the de-facto standard for IBP-based training methods. While we (theoretically and empirically) observe an exponential reduction in tightness with increasing depth, the inverse IBP loss increases slightly. Similarly, while we (theoretically and empirically) observe a polynomial reduction in tightness with increasing width, the inverse IBP loss stays almost constant. Note the logarithmic scale (and orders of magnitude larger changed) for tightness and the linear scale for the inverse IBP loss. This difference in trend is unsurprising as the custom initialization of Shi et al. (2021) is designed to keep IBP bound width constant over network depth and width. We thus conclude that tightness and (inverse) IBP loss yield fundamentally different results and insights when analyzing networks at initialization.




Figure 15: Tightness and inverse IBP loss at initialization depending on width and depth.

**IBP Loss after Training** We show the inverse IBP loss (left and center) and tightness (right) after IBP training depending on training perturbation size  $\epsilon$ , evaluated at training  $\epsilon$  (left) or constant  $\epsilon = 10^{-3}$  (center) in Figure 16. We observe that inverse IBP loss, in contrast to tightness, is heavily dependent on the perturbation magnitude used for evaluation (compare left and center), making it poorly suited to analyze the effects of changing perturbation magnitude. Further, when using the most natural perturbation magnitude, the  $\epsilon$  used during training and certification (left), we observe completely different trends to tightness. For very small perturbation magnitudes, the inverse IBP

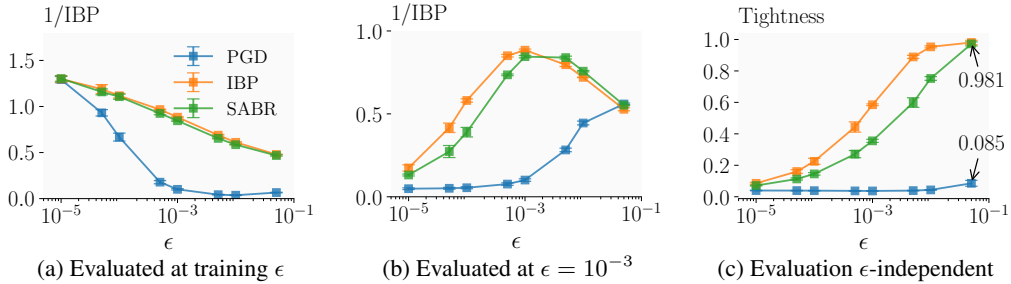


Figure 16: Tightness (right) and inverse IBP loss (left and center) after IBP training depending on training perturbation size  $\epsilon$ , evaluated at training  $\epsilon$  (left) or constant  $\epsilon = 10^{-3}$  (center).

Table 3: Certified and standard accuracy depending on network width.

Dataset	$\epsilon$	Method	Width	Accuracy	Certified	Tightness
MNIST	0.1	IBP	1×	85.70	67.71	0.871
			2×	88.42	73.77	0.857
			4×	<b>90.31</b>	<b>79.89</b>	0.803

loss is very high, suggesting high (perturbation) robustness, but both inverse IBP loss evaluated with a larger  $\epsilon$  and tightness are low, showing that it neither permits precise analysis with IBP nor is necessarily robust, again highlighting the difference between tightness and (inverse) IBP loss.

#### D.5 TIGHTNESS AFTER IBP TRAINING

To confirm that wider models improve certified accuracy while slightly reducing tightness across network architectures, we also consider fully connected networks, which used to be the default in neural network verification (Singh et al., 2019b; 2018). We increase the width of a fully connected ReLU network with 6 hidden layers from 100 to 400 neurons and indeed observe a significant increase in certified accuracy (see Table 3).