
IBP Regularization for Verified Adversarial Robustness via Branch-and-Bound

Alessandro De Palma¹ Rudy Bunel² Krishnamurthy (Dj) Dvijotham³ M. Pawan Kumar² Robert Stanforth²

Abstract

Recent works have tried to increase the verifiability of adversarially trained networks by running the attacks over domains larger than the original perturbations and adding various regularization terms to the objective. However, these algorithms either underperform or require complex and expensive stage-wise training procedures, hindering their practical applicability. We present IBP-R, a novel verified training algorithm that is both simple and effective. IBP-R induces network verifiability by coupling adversarial attacks on enlarged domains with a regularization term, based on inexpensive interval bound propagation, that minimizes the gap between the non-convex verification problem and its approximations. By leveraging recent branch-and-bound frameworks, we show that IBP-R obtains state-of-the-art verified robustness-accuracy trade-offs for small perturbations on CIFAR-10 while training significantly faster than relevant previous work. Additionally, we present UPB, a novel branching strategy that, relying on a simple heuristic based on β -CROWN, reduces the cost of state-of-the-art branching algorithms while yielding splits of comparable quality.

1. Introduction

The existence of adversarial examples (Szegedy et al., 2014; Goodfellow et al., 2015) has raised widespread concerns on the robustness of neural networks. As a consequence, many authors have promptly devised algorithms to formally prove the robustness of trained networks (Katz et al., 2017; Ehlers, 2017; Bunel et al., 2018; Zhang et al., 2018; Raghunathan et al., 2018). At the same time, a number of works have focused on training networks for adversarial robustness: first by defending against specific attacks (adversarial training) (Madry et al., 2018), then providing formal guarantees

about attack-independent robustness (verified training) (Dvijotham et al., 2018; Wong & Kolter, 2018; Mirman et al., 2018). The vast majority of verified training methods operate by backpropagating over over-approximations of the network’s loss under adversarial perturbations, and obtain state-of-the-art results for large perturbations (Zhang et al., 2020; Xu et al., 2020; Lyu et al., 2021). However, these training schemes are typically unable to benefit from tight over-approximations at verification time, hence requiring relatively large networks to perform at their best. A recent line of work has better leveraged network capacity by enhancing the verifiability of adversarially-trained networks. These algorithms can exploit tighter over-approximations but they either underperform (Xiao et al., 2019) or require expensive procedures in order to reach state-of-the-art performance on small perturbations (Balunovic & Vechev, 2020).

The recent VNN-COMP-21, an international competition on neural network verification (Bak et al., 2021) highlighted significant scaling improvements in exact verification algorithms (Henriksen & Lomuscio, 2021; Serre et al., 2021; De Palma et al., 2021c; Wang et al., 2021). We aim to leverage these developments by presenting IBP-R, a novel and inexpensive verified training algorithm that induces network verifiability by: (i) running adversarial attacks over domains that are significantly larger than the target perturbations, (ii) exploiting Interval Bound Propagation (IBP) (Mirman et al., 2018; Gowal et al., 2018) to minimize the area of the convex hull of the activations, a commonly employed relaxation within recent verification frameworks. We show that, in spite of its speed and conceptual simplicity, IBP-R yields state-of-the-art results under small perturbations on CIFAR-10. In particular, under ℓ_∞ perturbations of radius $\epsilon_{\text{ver}} = 2/255$, networks trained via IBP-R attain, on average: a verified accuracy of 61.97, a robust accuracy of 66.39 under MI-FGSM attacks (Dong et al., 2018) and a natural accuracy of 78.19. In this setting, IBP-R trains in less than a third of the runtime of COLT (Balunovic & Vechev, 2020). Furthermore, for $\epsilon_{\text{ver}} = 8/255$, IBP-R performs competitively with COLT while almost halving its runtime.

Finally, motivated by the task of evaluating the verifiability of networks trained via IBP-R, we present a simple and novel branching strategy, named UPB, for complete verification via branch-and-bound (Bunel et al., 2018). UPB leverages dual information from the recent β -CROWN al-

¹University of Oxford ²DeepMind ³Google Brain. Correspondence to: Alessandro De Palma <adepalma@robots.ox.ac.uk>.

1st Workshop on Formal Verification of Machine Learning, Baltimore, Maryland, USA. Colocated with ICML 2022. Copyright 2022 by the author(s).

gorithm (Wang et al., 2021) to heuristically rank the quality of the possible branching decisions. We show that, at a cost equivalent to a single gradient backpropagation through the network, UPB obtains a verification performance comparable to the more expensive and state-of-the-art FSB strategy (De Palma et al., 2021c).

2. Background

In the following, we will use boldface letters to denote vectors (for example, \mathbf{x}), uppercase letters to denote matrices (for example, W_k), and brackets for intervals (for example, $[\hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k]$). Furthermore, we will write \odot for the Hadamard product, $\mathbb{1}_c$ for the indicator vector on condition c , $[\cdot]$ for integer ranges, and employ the following shorthand: $[\mathbf{x}]_+ := \max(\mathbf{x}, \mathbf{0})$.

Let us define the data distribution \mathcal{D} , yielding points $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}^d \times \mathbb{R}^o$ and let us denote by $\boldsymbol{\theta} \in \mathbb{R}^S$ the network parameters. Robust training is concerned with training a neural network $f : \mathbb{R}^S \times \mathbb{R}^d \rightarrow \mathbb{R}^o$ so that a given property $P : \mathbb{R}^o \times \mathbb{R}^o \rightarrow \{0, 1\}$ is satisfied in a region around each input point \mathbf{x} , denoted $\mathcal{C}(\mathbf{x})$. In other words, the parameters $\boldsymbol{\theta}$ must satisfy the following:

$$\mathbf{x}_0 \in \mathcal{C}(\mathbf{x}) \implies P(f(\boldsymbol{\theta}, \mathbf{x}_0), \mathbf{y}) \quad \forall (\mathbf{x}, \mathbf{y}) \in \mathcal{D}. \quad (1)$$

In this work, we will focus on robustness to adversarial perturbations around the input images. Specifically, $\mathcal{C}(\mathbf{x}) := \{\mathbf{x}_0 : \|\mathbf{x}_0 - \mathbf{x}\|_p \leq \epsilon_{\text{ver}}\}$ and P amounts to checking that the predicted and ground truth classification labels match.

2.1. Neural Network Verification

Before delving into the task of training a robust network, we first consider the problem of determining whether a given network is robust or not. This involves the formal verification of condition (1) on the given network, which is generally NP-HARD (Katz et al., 2017). Therefore, its exact verification is often replaced by less expensive approximations, which nevertheless provide formal guarantees for a subset of the properties (incomplete verification). By means of simple transformations, one can represent both f and P from condition (1) via a single network f' of depth n (Bunel et al., 2020b), so that incomplete verification corresponds to the following optimization problem:

$$\begin{aligned} \min_{\mathbf{x}, \hat{\mathbf{x}}} \quad & \hat{x}_n \quad \text{s.t.} \\ & \mathbf{x}_0 \in \mathcal{C}(\mathbf{x}), \\ & \hat{\mathbf{x}}_{k+1} = W_{k+1}\mathbf{x}_k + \mathbf{b}_{k+1} \quad k \in \llbracket 0, n-1 \rrbracket, \\ & (\mathbf{x}_k, \hat{\mathbf{x}}_k) \in \text{Rel}(\sigma, \hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k) \quad k \in \llbracket 1, n-1 \rrbracket, \\ & \hat{\mathbf{x}}_k \in [\hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k] \quad k \in \llbracket 1, n-1 \rrbracket, \end{aligned} \quad (2)$$

where σ denotes the activation function, W_k the weight matrix of the k -th layer of f' , \mathbf{b}_k its bias. The use of

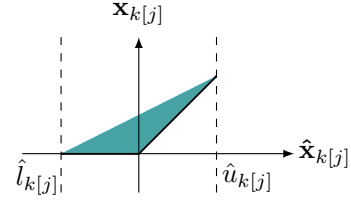


Figure 1: Convex hull for an ambiguous ReLU (Ehlers, 2017).

$\text{Rel}(\sigma, \hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k)$, a convex relaxation of σ , ensures that problem (2) is convex, greatly simplifying its solution. In general, $\text{Rel}(\sigma, \hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k)$ is a function of intermediate bounds $\hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k$, which provide ranges on the network pre-activation variables (for details, see appendix C). In the context of ReLU networks, which are the focus of this work, a popular relaxation choice is the convex hull of the activation, commonly referred to as the Planet relaxation (Ehlers, 2017). If $\hat{\mathbf{l}}_k < 0$ and $\hat{\mathbf{u}}_k > 0$ (ambiguous ReLU), its shape is given by Figure 1. If either $\hat{\mathbf{l}}_k > 0$ or $\hat{\mathbf{u}}_k < 0$, the activation is said to be stable and its convex hull can be represented by a line, greatly improving the tightness of the overall network approximation.

When verifying all properties is a requirement (complete verification), problem (2) is employed as a sub-routine for a global optimization algorithm equivalent to branch-and-bound (Bunel et al., 2018). The goal is to find the sign of the minimum of a non-convex problem (reported in appendix B) whose domain is a subset of the feasible region from problem (2). Complete verifiers hence proceed by recursively splitting the domain (branching) and solving the resulting convex sub-problems (bounding) until a definite answer can be provided. For ReLU activations, the branching is usually performed by splitting the domain of an ambiguous ReLU into its two stable subdomains.

2.2. Training via the Robust Loss

In order to enforce condition (1) during training, one typically defines a surrogate loss $\mathcal{L} : \mathbb{R}^o \times \mathbb{R}^o \rightarrow \mathbb{R}$, and seeks to minimize the worst-case empirical risk within $\mathcal{C}(\mathbf{x})$, referred to as the robust loss:

$$\min_{\boldsymbol{\theta}} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \left[\max_{\mathbf{x}' \in \mathcal{C}(\mathbf{x})} \mathcal{L}(f(\boldsymbol{\theta}, \mathbf{x}'), \mathbf{y}) \right]. \quad (3)$$

The exact computation of $\max_{\mathbf{x}' \in \mathcal{C}(\mathbf{x})} \mathcal{L}(f(\boldsymbol{\theta}, \mathbf{x}'), \mathbf{y})$ entails the use of a complete verification algorithm (§2.1), which is too expensive to be employed during training. Therefore, the robust loss is typically replaced by an approximation: adversarial training algorithms (Madry et al., 2018) rely on lower bounds, while certified training algorithms (Gowal et al., 2018; Zhang et al., 2020) employ upper bounds. Lower bounds are computed by using so-called adversarial attacks: algorithms, such as PGD (Madry et al., 2018), that heuristically search for misclassified (adversarial) examples in the input space.

Upper bounds are instead computed by solving an instance of problem (2) where $\text{Rel}(\sigma, \hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k)$ typically represents the IBP hyper-rectangle (Mirman et al., 2018) or linear bounds on the activation (Wong & Kolter, 2018; Zhang et al., 2018). Adversarial training yield models with strong standard accuracy and empirical robustness. However, such robustness is often hard to demonstrate via a formal verification method, and might potentially break under stronger attacks. Stronger robustness guarantees are instead provided by certified training algorithms, at the expense of the standard network accuracy and with longer training times.

2.3. Hybrid Training Methods

A line of recent work seeks to bridge the gap between adversarial and certified training by modifying the regions over which the attacks are performed and adding specialized regularization terms.

Xiao et al. (2019) demonstrate that the verified robust accuracy of PGD-trained networks (Madry et al., 2018) can be increased by adding ℓ_1 regularization and a term encouraging ReLU stability to problem (3):

$$\min_{\boldsymbol{\theta}} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \left[\max_{\mathbf{x}' \in \mathcal{C}(\mathbf{x})} \mathcal{L}(f(\boldsymbol{\theta}, \mathbf{x}'), \mathbf{y}) + \lambda \|\boldsymbol{\theta}\|_1 + \rho \sum_{j=1}^k \tanh(1 - \hat{\mathbf{l}}_k(\boldsymbol{\theta}) \odot \hat{\mathbf{u}}_k(\boldsymbol{\theta}))^T \mathbf{1}, \quad (4)$$

where $\hat{\mathbf{l}}_k(\boldsymbol{\theta})$ and $\hat{\mathbf{u}}_k(\boldsymbol{\theta})$, which depend on the network parameters, are computed via a tightened version of IBP.

Balunovic & Vechev (2020) propose to employ adversarial training layer-wise, running the attacks over convex outer-approximations of frozen subsets of the network. Let us denote by f^j a subset of network f that starts at the j -th layer, and by $\boldsymbol{\theta}^j$ its parameters. Furthermore, let $\mathcal{C}_j(\mathbf{x})$ represent an outer-approximation of the j -th latent space obtained via the zonotope relaxation (Zhang et al., 2018), relying on zonotope intermediate bounds approximated via Cauchy random projections (Li et al., 2007). Convex Layer-wise Adversarial Training (COLT), operates on the following objective at the j -th stage of the training:

$$\min_{\boldsymbol{\theta}^j(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \mathbb{E} \left\{ \begin{array}{l} \kappa \left[\begin{array}{l} \max_{\mathbf{x}' \in \mathcal{C}_j(\mathbf{x})} \mathcal{L}(f^j(\boldsymbol{\theta}^j, \mathbf{x}'), \mathbf{y}) \\ + \rho_j \left[-\hat{\mathbf{l}}_{j+1}(\boldsymbol{\theta}^j) \right]_+^T \left[\hat{\mathbf{u}}_{j+1}(\boldsymbol{\theta}^j) \right]_+ \end{array} \right] \\ + (1 - \kappa) \max_{\mathbf{x}' \in \mathcal{C}_{j-1}(\mathbf{x})} \mathcal{L}(f^{j-1}(\boldsymbol{\theta}^{j-1}, \mathbf{x}'), \mathbf{y}) \\ + \lambda \|\boldsymbol{\theta}^j\|_1, \end{array} \right. \quad (5)$$

where $\left[-\hat{\mathbf{l}}_{j+1} \right]_+^T \left[\hat{\mathbf{u}}_{j+1} \right]_+$ is a regularizer for the $(j + 1)$ -th latent space, inducing ReLU stability and minimizing the area of the zonotope relaxation for ambiguous ReLUs. As when computing intermediate bounds for $\mathcal{C}_j(\mathbf{x})$, the regularizer is computed via approximate zonotope bounds. In order to gradually transition from one training stage to the other, κ is linearly increased from 0 to 1 in the first phase

of the training. At the first stage ($j = 0$), the loss transitions from the natural loss (without any adversarial component) to the regularized PGD loss. COLT performs particularly well for smaller perturbation radii, for which it yields state-of-the-art results. However, its complexity and stage-wise nature make it relatively hard to deploy in practice. For instance, Balunovic & Vechev (2020) employ a different value for both ρ_j and the train-time perturbation radius, which affects both $\mathcal{C}_j(\mathbf{x})$ and intermediate bounds, at each training stage. Further details are provided in appendix A.

3. Training via IBP Regularization

Certified training algorithms that directly employ upper bounds to the robust loss (3) do not typically benefit from the use of more accurate verification algorithms than those they were trained with (Zhang et al., 2020). On the other hand, the hybrid training methods from §2.3 are designed to be verified with tighter bounds, potentially encoding (part of) the network as a MILP (Tjeng et al., 2019). In light of the recent scaling improvements of complete verifiers (Bak et al., 2021), we present a robust training method designed for recent branch-and-bound frameworks (De Palma et al., 2021c; Wang et al., 2021), capable of preserving COLT’s effectiveness while simplifying its training procedure.

Training objective Intuitively, verification is easier if the network is robust by a large margin, and if the employed relaxation accurately represents the network. Therefore, we propose a simple training scheme revolving around the following two features: (i) adversarial attacks run over significantly larger domains than those employed at test time, (ii) a term minimizing the gap between the relaxations used for verification and the original neural network domain. Given its flexibility, small cost, and widespread usage, we aim to exclusively rely on IBP for bounding computations. Details for IBP can be found in appendix D. We name the resulting method IBP Regularization (IBP-R), which operates on the following objective:

$$\min_{\boldsymbol{\theta}} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \left\{ \begin{array}{l} \kappa \left[\begin{array}{l} \max_{\mathbf{x} \in \mathcal{C}_+(\mathbf{x})} \mathcal{L}(f(\boldsymbol{\theta}, \mathbf{x}), \mathbf{y}) + \\ \frac{\rho}{2} \sum_{j=1}^n \left[-\hat{\mathbf{l}}_j(\boldsymbol{\theta}) \right]_+^T \left[\hat{\mathbf{u}}_j(\boldsymbol{\theta}) \right]_+ \end{array} \right] \\ + (1 - \kappa) \mathcal{L}(f(\boldsymbol{\theta}, \mathbf{x}), \mathbf{y}) + \lambda \|\boldsymbol{\theta}\|_1, \end{array} \right. \quad (6)$$

where $\frac{1}{2} \left[-\hat{\mathbf{l}}_j(\boldsymbol{\theta}) \right]_+^T \left[\hat{\mathbf{u}}_j(\boldsymbol{\theta}) \right]_+$ represents the area of the widely-employed ReLU convex hull represented in Figure 1, and $\mathcal{C}_+(\mathbf{x})$ is a superset of the original input specification from condition (1). For adversarial robustness specifications, $\mathcal{C}_+(\mathbf{x}) := \{\mathbf{x}_0 \mid \|\mathbf{x}_0 - \mathbf{x}\|_p \leq \alpha \epsilon_{\text{ver}}\}$, with $\alpha \geq 1.6$ in our experiments (see §6). Note that we regularize over all the activations of the network at once, with the same regularization coefficient. An in-depth comparison with COLT is available in appendix A.

Regularization masking When a property holds, the tightness of the employed relaxations is fundamental in order to swiftly provide a formal guarantee via branch-and-bounds methods. On the other hand, verifying that the property does not hold typically implies finding counter-examples via adversarial attacks (Bak et al., 2021). As a consequence, there is no need to encourage tightness by minimizing $\frac{1}{2}[-\hat{\mathbf{l}}_j(\boldsymbol{\theta})]_+^T [\hat{\mathbf{u}}_j(\boldsymbol{\theta})]_+$ when $P(f(\boldsymbol{\theta}, \mathbf{x}_0), \mathbf{y}) \forall \mathbf{x}_0 \in \mathcal{C}(\mathbf{x})$ is unlikely to hold after training. In order to take this observation into account, we propose to mask the convex hull regularizer when a counter-example is found for the current sample. Denoting by $\tilde{\mathbf{x}}$ the point found by the train-time adversarial attack, we perform the following substitution in objective (6):

$$\frac{1}{2}[-\hat{\mathbf{l}}_j(\boldsymbol{\theta})]_+^T [\hat{\mathbf{u}}_j(\boldsymbol{\theta})]_+ \rightarrow \frac{\mathbb{1}_{P(f(\boldsymbol{\theta}, \tilde{\mathbf{x}}), \mathbf{y})}}{2} [-\hat{\mathbf{l}}_j(\boldsymbol{\theta})]_+^T [\hat{\mathbf{u}}_j(\boldsymbol{\theta})]_+.$$

Training details As common for recent certified training algorithms (Balunovic & Vechev, 2020; Zhang et al., 2020), κ is linearly increased from 0 to 1 at the beginning of training (mixing). Similarly to IBP (Gowal et al., 2018) and CROWN-IBP (Zhang et al., 2020), we also linearly increase the effective perturbation radius from 0 to $\alpha\epsilon_{\text{ver}}$ while mixing the objectives.

4. Verification Framework

As explained in §3, IBP-R is designed to facilitate the verification of trained networks via recent branch-and-bound frameworks. We will now first present the details of the employed complete verifier (§4.1), then present a novel branching strategy in §4.2.

4.1. Branch-and-Bound Setup

Owing to its modularity and performance on large COLT-trained networks (Bak et al., 2021, cifar2020 benchmark), we base our verifier on the OVAL branch-and-bound framework (Bunel et al., 2020a;b; De Palma et al., 2021c) from VNN-COMP-2021 (Bak et al., 2021). Given that IBP-R explicitly seeks to minimize the area of the ReLU convex hull (see Figure 1), we instantiate the framework so as to use β -CROWN (Wang et al., 2021) for the bounding (see §2.1), a state-of-the-art solver for the employed relaxation, designed for use within branch-and-bound. In line with De Palma et al. (2021c), intermediate bounds are never updated after branching, and they are individually computed via α -CROWN (Xu et al., 2021). Furthermore, the dual variables of each bounding computation are initialized to the values associated to the parent node (that is, the bounding performed before the last split), and the number of dual iterations is dynamically adjusted to reduce the bounding time (De Palma et al., 2021c). Counter-examples are found using the MI-FGSM (Dong

et al., 2018) adversarial attack, which is run repeatedly for each property, using a variety of hyper-parameter settings. Verification is run with a timeout of 1800 seconds, and terminated early when the property is likely to time out. We now present the employed branching strategy.

4.2. UPB Branching

In spite of its strong empirical performance on COLT-trained networks (Bak et al., 2021), the FSB branching strategy (De Palma et al., 2021c), commonly employed for β -CROWN (Wang et al., 2021), requires $O(n)$ CROWN-like bounding computations per split. In order to reduce the branching overhead, we present Upper Planet Bias (UPB), a novel and simpler branching strategy that yields splitting decisions of comparable quality at the cost of a single gradient backpropagation through the network.

The popular SR (Bunel et al., 2020b) and FSB branching strategies partly rely on estimates of the sensitivity of FastLin bounds (Wong & Kolter, 2018) to the splitting of an ambiguous ReLUs (De Palma et al., 2021c). As the employed relaxation for the output bounding is typically much tighter, these estimates are often unreliable. In order to improve branching performance, FSB couples these estimates with an expensive bounding step. We propose to remove the need to compute bounds at branching time by re-using dual information from branch-and-bound’s bounding step, which we perform using β -CROWN (see §4.1). Specifically, we propose to score branching decisions according to a dual term associated to the bias of the upper linear constraint from the Planet relaxation (Figure 1) for each ambiguous neuron:

$$s_{\text{UPB},k} = \frac{[-\hat{\mathbf{l}}_j]_+^T [\hat{\mathbf{u}}_j]_+}{(\hat{\mathbf{u}}_k - \mathbf{l}_k)} \odot [\bar{\boldsymbol{\lambda}}_k]_+ \quad k \in \llbracket 1, n-1 \rrbracket, \quad (7)$$

where $[\bar{\boldsymbol{\lambda}}_k]_+$ is computed by evaluating equation (15) in appendix E on the dual variables computed for the branch-and-bound bounding step. Therefore, the cost of computing scores (7) for all neurons corresponds to that of a single gradient backpropagation. Intuitively, as $s_{\text{UPB},k}$ will disappear from the dual objective after splitting, we employ it as a proxy for a branching decision’s potential bounding improvement. See appendix E for further details.

5. Related Work

Many popular certified training algorithms work by upper bounding the robust loss (3) via some combination of IBP and linear bounds on the activation function. IBP (Gowal et al., 2018; Mirman et al., 2018), CAP (Wong & Kolter, 2018; Wong et al., 2018), and CROWN-IBP (Zhang et al., 2020) all fall in this category. Shi et al. (2021) recently proposed a series of techniques to shorten the usually long training schedules of these algorithms. Xu et al. (2020) pro-

vide minor improvements on CROWN-IBP by changing the way the loss function is incorporated into problem (2). The above family of methods produce state-of-the-art results for larger perturbation radii. Regularization-based techniques, instead, tend to perform better on smaller radii (see §2.3). Further works have focused on achieving robustness via specialized network architectures (Lyu et al., 2021; Zhang et al., 2021), Lipschitz constant estimation for perturbations in the ℓ_2 norm (Huang et al., 2021), or under randomized settings (Cohen et al., 2019; Salman et al., 2019a): these methods are out of the scope of the present work.

As outlined in §2, widely-employed neural network relaxations model the convex hull of the activation function, referred to as the convex barrier due to its popularity (Salman et al., 2019b), or on even looser convex outer-approximations such as FastLin (Wong & Kolter, 2018), CROWN (Zhang et al., 2018), DeepZ (Singh et al., 2018), or DeepPoly (Singh et al., 2019b). These relaxations are relatively inexpensive yet very effective when adapted for complete verification via branch-and-bound, and are hence at the core of the α - β -CROWN (Xu et al., 2021; Wang et al., 2021) and OVAL frameworks (Bunel et al., 2020a; De Palma et al., 2021c). A number of works have recently focused on devising tighter neural network relaxations (Singh et al., 2019a; Anderson et al., 2020; Tjandraatmadja et al., 2020; Müller et al., 2022). These have been integrated into recent branch-and-bound verifiers (De Palma et al., 2021a;b; Ferrari et al., 2022) and yield strong results for harder verification properties on medium-sized networks.

While the employed relaxations are a fundamental component of a complete verifier, the overall speed-accuracy trade-offs are greatly affected by the employed branching strategy. Small networks with few input dimensions can be quickly verified by recursively splitting the input region (Bunel et al., 2018; Royo et al., 2019). On the other hand, activation splitting (Ehlers, 2017; Katz et al., 2017) (see §2.1, §4.2) performs better for larger convolutional networks (Bunel et al., 2020b; De Palma et al., 2021c) and enjoyed recent developments based on graph neural networks (Lu & Kumar, 2020) or for use with tighter relaxations (Ferrari et al., 2022).

6. Experiments

In this section, we present an experimental evaluation of the IBP-R certified training algorithm (§6.1), then evaluate the performance of our UPB branching strategy (§6.2). The implementation of our training algorithm is based on Jax (Bradbury et al., 2018), while verification is performed post-training by using a modified version (see §4) of the OVAL framework, implemented in PyTorch (Paszke et al., 2019).

6.1. Verified Training

We evaluate the efficacy of our IBP-R algorithm (§3) by replicating the CIFAR-10 (Krizhevsky & Hinton, 2009) experiments from Balunovic & Vechev (2020) and comparing against COLT, which we ported to Jax for fairness (resulting in significant speed-ups as shown in table 1, see appendix F.2 for details). We focus our comparison on COLT, as it is the best-performing instance of the hybrid training algorithms detailed in §2.3, family to which IBP-R belongs. Furthermore, to the best of our knowledge, it yields state-of-the-art results for small perturbations and ReLU networks. Details concerning the employed architecture, hyper-parameters, and the computational setup can be found in appendix F. Timing results were executed on a Nvidia Titan V GPU.

Table 1 reports the results of our experiment, as well as relevant results from the literature. Specifically, we report results for: (i) CROWN-IBP (Zhang et al., 2020) and the improved IBP algorithm by Lyu et al. (2021), representing state-of-the-art certified training algorithms based on upper bounding the robust loss (see §2.2), (ii) Xiao et al. (2019), sharing many similarities with IBP-R (see §2.3), and (iii) the original COLT experiments from Balunovic & Vechev (2020). IBP-R excels on the smaller ℓ_∞ perturbation radius ($\epsilon_{\text{ver}} = 2/255$), displaying verified and standard accuracies comparable (given the experimental variability) to COLT, and a larger empirically robust accuracy (for details on the employed attack, see §4.1). Furthermore, IBP-R training is more than three times faster than COLT, making its use particularly convenient on this setup. The masking (see §3) does not appear to be beneficial on smaller perturbations. Relevant results from the literature all underperform compared to IBP-R, which hence achieves state-of-the-art results on this benchmark. On the larger perturbation radius ($\epsilon_{\text{ver}} = 8/255$), masking the regularization has a positive effect. The masked version of IBP-R performs comparably with COLT in $\approx 57\%$ of its runtime, attaining larger standard and empirically robust accuracies, and smaller verified accuracy. However, in this context, all regularization-based methods (including IBP-R) are outperformed by algorithms upper bounding the robust loss (Zhang et al., 2020; Lyu et al., 2021). Nevertheless, these works employ data augmentation and significantly larger networks than the one used in our experiments. We therefore conjecture that IBP-R, which scales better than COLT with the network depth (see runtimes in table 1), will become more competitive when evaluated on comparable settings.

6.2. Branching

In order to test the efficacy of our UPB branching strategy (§4.2), we time verification for the first 500 CIFAR-10 test samples on two IBP-R-trained networks (one per perturbation radius) from the experiment of §6.1. Images that

IBP Regularization for Verified Adversarial Robustness

Perturbation	Method	Standard accuracy	Robust accuracy	Verified accuracy	Runtime [s]	
$\epsilon_{\text{ver}} = 2/255$	COLT	78.37 ± 0.24	65.66 ± 0.13	61.88 ± 0.11	$3.05 \times 10^4 \pm 1.21 \times 10^2$ ‡	
	IBP-R	78.19 ± 0.52	66.39 ± 0.12	61.97 ± 0.18	$9.34 \times 10^3 \pm 2.95 \times 10^1$	
	IBP-R w/ MASKING	78.22 ± 0.26	66.28 ± 0.17	61.69 ± 0.29	$9.63 \times 10^3 \pm 4.42 \times 10^1$	
	LITERATURE RESULTS					
	(Shi et al., 2021)*	66.84	/	52.85		
	(Zhang et al., 2020)*	71.52	59.72	53.97	9.13×10^4 †	
$\epsilon_{\text{ver}} = 8/255$	(Balunovic & Vechev, 2020)	78.4	/	60.50		
	(Xiao et al., 2019)	61.12	49.92	45.93		
	COLT	51.94 ± 0.14	31.68 ± 0.23	28.73 ± 0.23	$1.03 \times 10^4 \pm 1.70 \times 10^1$ ‡	
	IBP-R	51.43 ± 0.21	31.89 ± 0.11	27.87 ± 0.01	$5.92 \times 10^3 \pm 2.95 \times 10^1$	
	IBP-R w/ MASKING	52.74 ± 0.30	32.78 ± 0.33	27.55 ± 0.22	$5.89 \times 10^3 \pm 3.38 \times 10^1$	
	LITERATURE RESULTS					
(Shi et al., 2021)*	48.28 ± 0.40	/	34.42 ± 0.32	9.51×10^3 ◊		
(Zhang et al., 2020)*	54.50	34.26	30.50	9.13×10^4 †		
(Balunovic & Vechev, 2020)	51.70	/	27.50			
(Xiao et al., 2019)	40.45	26.78	20.27			

* the employed 7-layer network has 17.2×10^6 parameters, as opposed to the 2.1×10^6 parameters of the 5-layer and 4-layer networks respectively used for our $\epsilon_{\text{ver}} = 2/255$ and $\epsilon_{\text{ver}} = 8/255$ experiments. Differently from our work, data augmentation (random horizontal flips and croppings) is used.

◊ the training is performed on an Nvidia RTX 2080 Ti GPU, which is significantly faster than that employed in our experiments.

† the training is performed on 4 Nvidia RTX 2080 Ti GPUs.

‡ on the same setup, the original PyTorch implementation runs in 1.74×10^5 seconds for $\epsilon_{\text{ver}} = 2/255$, and 4.49×10^4 seconds for $\epsilon_{\text{ver}} = 8/255$.

Table 1: Performance of different verifiably robust training algorithms under ℓ_∞ norm perturbations, on the CIFAR10 dataset with $\epsilon_{\text{ver}} = 2/255$ and $\epsilon_{\text{ver}} = 8/255$. The table reports mean and sample standard deviation over three repetitions for our experiments, over five repetitions for (Shi et al., 2021) on $\epsilon_{\text{ver}} = 8/255$. The remaining results from the literature were executed with a single seed. The method with the best average performance for each perturbation radius is highlighted in bold.

are incorrectly classified are discarded. We keep the branch-and-bound settings fixed to those of §4.1, and benchmark against the following branching strategies: FSB (De Palma et al., 2021c), and SR (Bunel et al., 2020b). Appendix G replicates the experiment on two COLT-trained networks.

Figure 2 shows that UPB improves average verification times compared to FSB (by roughly 13% and 2.5% for the $\epsilon_{\text{ver}} = 2/255$ and $\epsilon_{\text{ver}} = 8/255$ networks, respectively). On the larger perturbation radius, this holds in spite of a larger average number of visited subproblems, highlighting the cost of FSB’s bounding-based selection step. On the other hand, UPB verifies more properties within the timeout on both networks. Furthermore, it reduces the number of visited subproblems on the smaller perturbation, testifying the efficacy of the selected domain splits. Both UPB and FSB yield significantly faster verification than SR on the considered problems. The results show that UPB is less expensive than the state-of-the-art FSB algorithm, while producing branching decisions of comparable quality.

7. Conclusions

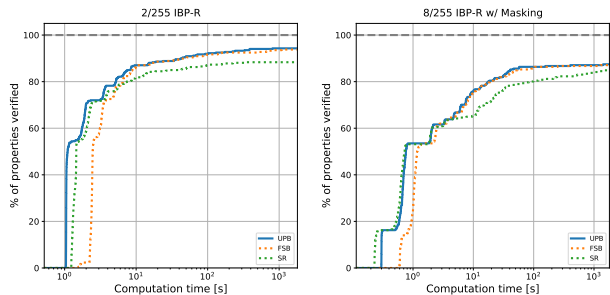
Many state-of-the-art verified training algorithms require very large networks to obtain good robustness-accuracy trade-offs. Methods designed to exploit tight bounds at verification-time better exploit network capacity but they either underperform or involve extremely complex training procedures. We introduced IBP-R, a simple and intuitive robust training algorithm designed to induce verifiability via recent branch-and-bound algorithms. We show that, by minimizing the area of the Planet relaxation via IBP bounds for all network activations, and using PGD attacks over

larger perturbations, one can obtain state-of-the-art certified accuracy results on small perturbations without large compromises in standard accuracy. Finally, in order to ease the task of verifying the trained networks, we presented UPB, a straightforward and inexpensive branching strategy that yields branching decisions as effective as the state-of-the-art. We believe our results could be further improved by leveraging recent improvements on standard IBP training (Shi et al., 2021), as well as larger network architectures: these are interesting avenues for future work.

Method	$\epsilon_{\text{ver}} = 2/255$ IBP-R			$\epsilon_{\text{ver}} = 8/255$ IBP-R w/ Masking		
	time [s]	subproblems*	%Timeout	time [s]	subproblems*	%Timeout
UPB	113.42	930.83	5.70	237.24	935.80	12.55
FSB	127.65	1244.01	6.22	243.29	795.22	12.92
SR	215.89	5185.27	11.66	303.76	6744.72	15.13

* computed on the properties that did not time out for neither UPB nor FSB. The inclusion of timed-out results in the average leads to an overestimation of the number of subproblems for the less expensive branching strategy.

(a) Comparison of average runtime, average number of solved subproblems and the percentage of timed out properties. The best performing method is highlighted in bold.



(b) Cactus plots: percentage of solved properties as a function of runtime. Baselines are represented by dotted lines.

Figure 2: Complete verification performance of different branching strategies on two IBP-R-trained CIFAR-10 networks from §6.1.

References

- Anderson, R., Huchette, J., Ma, W., Tjandraatmadja, C., and Vielma, J. P. Strong mixed-integer programming formulations for trained neural networks. *Mathematical Programming*, 2020.
- Bak, S., Liu, C., and Johnson, T. The second international verification of neural networks competition (VNN-COMP 2021): Summary and results. *arXiv preprint arXiv:2109.00498*, 2021.
- Balunovic, M. and Vechev, M. Adversarial training and provable defenses: Bridging the gap. *International Conference on Learning Representations*, 2020.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. JAX median scaling issues, 2020. URL <https://github.com/google/jax/issues/4379>.
- Bunel, R., Turkaslan, I., Torr, P. H., Kohli, P., and Kumar, M. P. A unified view of piecewise linear neural network verification. *Neural Information Processing Systems*, 2018.
- Bunel, R., De Palma, A., Desmaison, A., Dvijotham, K., Kohli, P., Torr, P. H., and Kumar, M. P. Lagrangian decomposition for neural network verification. *Conference on Uncertainty in Artificial Intelligence*, 2020a.
- Bunel, R., Lu, J., Turkaslan, I., Kohli, P., Torr, P., and Kumar, M. P. Branch and bound for piecewise linear neural network verification. *Journal of Machine Learning Research*, 21(2020), 2020b.
- Cohen, J., Rosenfeld, E., and Kolter, Z. Certified adversarial robustness via randomized smoothing. In *International Conference on Machine Learning*, 2019.
- De Palma, A., Behl, H. S., Bunel, R., Torr, P. H. S., and Kumar, M. P. Scaling the convex barrier with active sets. *International Conference on Learning Representations*, 2021a.
- De Palma, A., Behl, H. S., Bunel, R., Torr, P. H. S., and Kumar, M. P. Scaling the convex barrier with sparse dual algorithms. *arXiv preprint arXiv:2101.05844*, 2021b.
- De Palma, A., Bunel, R., Desmaison, A., Dvijotham, K., Kohli, P., Torr, P. H., and Kumar, M. P. Improved branch and bound for neural network verification via Lagrangian decomposition. *arXiv preprint arXiv:2104.06718*, 2021c.
- Dong, Y., Liao, F., Pang, T., Su, H., Zhu, J., Hu, X., and Li, J. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 9185–9193, 2018.
- Dvijotham, K., Goyal, S., Stanforth, R., Arandjelovic, R., O’Donoghue, B., Uesato, J., and Kohli, P. Training verified learners with learned verifiers. *arxiv:1805.10265*, 2018.
- Ehlers, R. Formal verification of piece-wise linear feed-forward neural networks. *Automated Technology for Verification and Analysis*, 2017.
- Ferrari, C., Mueller, M. N., Jovanović, N., and Vechev, M. Complete verification via multi-neuron relaxation guided branch-and-bound. *International Conference on Learning Representations*, 2022.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *International Conference on Learning Representations*, 2015.
- Goyal, 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. *Workshop on Security in Machine Learning, NeurIPS*, 2018.
- Henriksen, P. and Lomuscio, A. Deepsplit: An efficient splitting method for neural network verification via indirect effect analysis. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI21)*, 2021.
- Hickey, T., Ju, Q., and Van Emden, M. H. Interval arithmetic: From principles to implementation. *Journal of the ACM (JACM)*, 2001.
- Huang, Y., Zhang, H., Shi, Y., Kolter, J. Z., and Anandkumar, A. Training certifiably robust neural networks with efficient local lipschitz bounds. In *Neural Information Processing Systems*, 2021.
- Katz, G., Barrett, C., Dill, D., Julian, K., and Kochenderfer, M. Reluplex: An efficient SMT solver for verifying deep neural networks. *Computer Aided Verification*, 2017.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2015.
- Krizhevsky, A. and Hinton, G. Learning multiple layers of features from tiny images. *Master’s thesis, Department of Computer Science, University of Toronto*, 2009.
- Li, P., Hastie, T. J., and Church, K. W. Nonlinear estimators and tail bounds for dimension reduction in l1 using

- cauchy random projections. In *Conference on Learning Theory*, 2007.
- Lu, J. and Kumar, M. P. Neural network branching for neural network verification. In *International Conference on Learning Representations*, 2020.
- Lyu, Z., Guo, M., Wu, T., Xu, G., Zhang, K., and Lin, D. Towards evaluating and training verifiably robust neural networks. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. *International Conference on Learning Representations*, 2018.
- Mirman, M., Gehr, T., and Vechev, M. Differentiable abstract interpretation for provably robust neural networks. *International Conference on Machine Learning*, 2018.
- Müller, M. N., Makarchuk, G., Singh, G., Püschel, M., and Vechev, M. PRIMA: General and precise neural network certification via scalable convex hull approximations. *Proceedings of the ACM on Programming Languages*, 2022.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 2019.
- Raghunathan, A., Steinhardt, J., and Liang, P. S. Semidefinite relaxations for certifying robustness to adversarial examples. *Neural Information Processing Systems*, 2018.
- Royo, V. R., Calandra, R., Stipanovic, D. M., and Tomlin, C. J. Fast neural network verification via shadow prices. *arXiv preprint arXiv:1902.07247*, 2019.
- Salman, H., Yang, G., Li, J., Zhang, P., Zhang, H., Razenshteyn, I., and Bubeck, S. Provably robust deep learning via adversarially trained smoothed classifiers. In *Neural Information Processing Systems*, 2019a.
- Salman, H., Yang, G., Zhang, H., Hsieh, C.-J., and Zhang, P. A convex relaxation barrier to tight robustness verification of neural networks. *Neural Information Processing Systems*, 2019b.
- Serre, F., Müller, C., Singh, G., Püschel, M., and Vechev, M. Scaling polyhedral neural network verification on GPUs. In *Machine Learning and Systems (MLSys)*, 2021.
- Shi, Z., Wang, Y., Zhang, H., Yi, J., and Hsieh, C.-J. Fast certified robust training with short warmup. In *Neural Information Processing Systems*, 2021.
- Singh, G., Gehr, T., Mirman, M., Püschel, M., and Vechev, M. Fast and effective robustness certification. *Neural Information Processing Systems*, 2018.
- Singh, G., Ganvir, R., Püschel, M., and Vechev, M. Beyond the single neuron convex barrier for neural network certification. *Neural Information Processing Systems*, 2019a.
- Singh, G., Gehr, T., Püschel, M., and Vechev, M. An abstract domain for certifying neural networks. *Proceedings of the ACM on Programming Languages*, 2019b.
- Sunaga, T. Theory of an interval algebra and its application to numerical analysis. *RAAG Memoirs*, 1958.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. Intriguing properties of neural networks. *International Conference on Learning Representations*, 2014.
- Tjandraatmadja, C., Anderson, R., Huchette, J., Ma, W., Patel, K., and Vielma, J. P. The convex relaxation barrier, revisited: Tightened single-neuron relaxations for neural network verification. *Neural Information Processing Systems*, 2020.
- Tjeng, V., Xiao, K., and Tedrake, R. Evaluating robustness of neural networks with mixed integer programming. *International Conference on Learning Representations*, 2019.
- Wang, S., Zhang, H., Xu, K., Lin, X., Jana, S., Hsieh, C.-J., and Kolter, J. Z. Beta-CROWN: Efficient bound propagation with per-neuron split constraints for complete and incomplete neural network verification. *Advances in Neural Information Processing Systems*, 2021.
- Wong, E. and Kolter, Z. Provable defenses against adversarial examples via the convex outer adversarial polytope. *International Conference on Machine Learning*, 2018.
- Wong, E., Schmidt, F., Metzen, J. H., and Kolter, J. Z. Scaling provable adversarial defenses. *Neural Information Processing Systems*, 2018.
- Xiao, K., Tjeng, V., Shafiq, N. M., and Madry, A. Training for faster adversarial robustness verification via inducing relu stability. *International Conference on Learning Representations*, 2019.
- Xu, K., Shi, Z., Zhang, H., Wang, Y., Chang, K.-W., Huang, M., Kaikhura, B., Lin, X., and Hsieh, C.-J. Automatic perturbation analysis for scalable certified robustness and beyond. In *Neural Information Processing Systems*, 2020.
- Xu, K., Zhang, H., Wang, S., Wang, Y., Jana, S., Lin, X., and Hsieh, C.-J. Fast and complete: Enabling complete

neural network verification with rapid and massively parallel incomplete verifiers. In *International Conference on Learning Representations*, 2021.

Zhang, B., Cai, T., Lu, Z., He, D., and Wang, L. Towards certifying l-infinity robustness using neural networks with l-inf-dist neurons. In *International Conference on Machine Learning*, 2021.

Zhang, H., Weng, T.-W., Chen, P.-Y., Hsieh, C.-J., and Daniel, L. Efficient neural network robustness certification with general activation functions. *Neural Information Processing Systems*, 2018.

Zhang, H., Chen, H., Xiao, C., Goyal, S., Stanforth, R., Li, B., Boning, D., and Hsieh, C.-J. Towards stable and efficient training of verifiably robust neural networks. *International Conference on Learning Representations*, 2020.

A. Comparison between IBP-R and COLT

In this section, we provide a detailed comparison between COLT by Balunovic & Vechev (2020) (§2.3) and IBP-R (§3) in the context of training for adversarial robustness. Equation (8) pictorially highlights the differences between the two algorithms, which are summarized in table 2.

Training Detail	COLT	IBP-R
Stages	Proceeds in a stage-wise fashion: at each stage, a subset of the network parameters (all the θ that are not in θ^j) is frozen, and the remaining parameters are optimized over. The number of training stages is $O(n)$.	All network parameters are optimized at once.
PGD domains	The attacks are carried out in the space of the j -th activations (\mathbf{x}_j). The frozen subset of the network is replaced by a zonotope-based convex outer-approximation, employed to define $\mathbb{C}_j(\mathbf{x})$ and $\mathbb{C}_{j-1}(\mathbf{x})$, the domains of the PGD attacks. At each stage, the zonotope domains are computed for train-time perturbations larger than those to be verified: $\epsilon_{\text{train},j} > \epsilon_{\text{ver}}$. In practice, the CIFAR-10 experiments from Balunovic & Vechev (2020) use $1.05 \leq \epsilon_{\text{train},j} \leq 1.512$. For the first stage: $\mathbb{C}_0(\mathbf{x}) := \mathcal{C}_+(\mathbf{x})$, $\mathbb{C}_{-1}(\mathbf{x}) := \mathbf{x}$.	The attacks are performed in the network input space, over a superset of the perturbations employed at verification. Our notation for $\mathcal{C}_+(\mathbf{x})$ corresponds to setting $\epsilon_{\text{train}} := \alpha \epsilon_{\text{ver}}$, with $1.6 \leq \alpha \leq 2.1$ in our experiments (after the mixing phase). Note that the first stage of COLT displays the same attack structure, with smaller employed α values.
Intermediate bounds	Computed via an approximation of the zonotope relaxation based on Cauchy random projections (Li et al., 2007). These bounds, computed for perturbation radius $\epsilon_{\text{train},j}$, are used both for the regularization term and to define the zonotope outer-approximations (which depend on intermediate bounds).	Computed via IBP, for perturbation radius ϵ_{train} . Used for the regularization term.
Bounds regularization	Minimizes the area of the zonotope relaxation of ambiguous ReLUs for a single layer per stage. This term produces a non-null gradient only for the parameters of the $(j+1)$ -th layer.	Minimizes the area of the Planet relaxation (in practice, one half of the area of the zonotope relaxation) of ambiguous ReLUs for all layers at once (including the output space \hat{x}_n).
Hyper-parameters	Each stage is potentially associated to a different regularization coefficient ρ_j , to a different train-time perturbation radius $\epsilon_{\text{train},j}$, and to a different learning rate η_j . In practice, Balunovic & Vechev (2020) tune the values for the first stage ($\rho_0, \epsilon_{\text{train},0}, \eta_0$), and then respectively decay ρ_0 and $\epsilon_{\text{train},0}$, and increase ρ_0 , by a fixed and tunable quantity at each stage. Finally, Balunovic & Vechev (2020) set $\rho_{n-2} = 0$ in all cases.	ρ, α , and η are not altered throughout training.
Mixing phase	At each stage, κ is linearly increased from 0 to 1.	κ is linearly increased from 0 to 1, ϵ_{train} is linearly increased from 0 to $\alpha \epsilon_{\text{ver}}$.

Table 2: Main differences between the COLT (§2.3) and IBP-R (§3) verified training algorithms.

$$\begin{array}{cc}
 \text{COLT:} & \text{IBP-R:} \\
 \left[\min_{\theta^j} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \left\{ \begin{array}{l} \kappa \left[\max_{\mathbf{x}' \in \mathbb{C}_j(\mathbf{x})} \mathcal{L}(f^j(\theta^j, \mathbf{x}'), \mathbf{y}) \right. \right. \\ \left. \left. + \rho_j \left[-\hat{\mathbf{I}}_{j+1}(\theta^j) \right]_+^T \left[\hat{\mathbf{u}}_{j+1}(\theta^j) \right]_+ \right] \right. \\ \left. + (1 - \kappa) \max_{\mathbf{x}' \in \mathbb{C}_{j-1}(\mathbf{x})} \mathcal{L}(f^{j-1}(\theta^{j-1}, \mathbf{x}'), \mathbf{y}) \right. \\ \left. + \lambda \|\theta^j\|_1, \right. \\ \left. \forall j \in \llbracket 0, n-2 \rrbracket \right. \end{array} \right] & \min_{\theta} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \left\{ \begin{array}{l} \kappa \left[\max_{\mathbf{x} \in \mathcal{C}_+(\mathbf{x})} \mathcal{L}(f(\theta, \mathbf{x}), \mathbf{y}) + \right. \\ \left. \frac{\rho}{2} \sum_{j=1}^n \left[-\hat{\mathbf{I}}_j(\theta) \right]_+^T \left[\hat{\mathbf{u}}_j(\theta) \right]_+ \right] \\ \left. + (1 - \kappa) \mathcal{L}(f(\theta, \mathbf{x}), \mathbf{y}) + \lambda \|\theta\|_1. \right. \end{array} \right.
 \end{array} \tag{8}$$

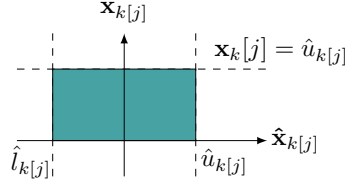


Figure 3: Depiction of the IBP hyper-rectangle.

B. Complete Verification Problem

Provided the network is in canonical form (Bunel et al., 2020b), complete verification amounts to finding sign of the minimum of the following problem, of which problem (2) is a convex outer-approximation:

$$\begin{aligned}
 \min_{\mathbf{x}, \hat{\mathbf{x}}} \quad & \hat{x}_n \quad \text{s.t.} \\
 & \mathbf{x}_0 \in \mathcal{C}(\mathbf{x}), \\
 & \hat{\mathbf{x}}_{k+1} = W_{k+1}\mathbf{x}_k + \mathbf{b}_{k+1} \quad k \in \llbracket 0, n-1 \rrbracket, \\
 & \mathbf{x}_k = \sigma(\hat{\mathbf{x}}_k) \quad k \in \llbracket 1, n-1 \rrbracket.
 \end{aligned} \tag{9}$$

C. Intermediate Bounds

As seen in §2.1, the network relaxations employed for incomplete verification (and, hence, complete verification via branch-and-bound) are defined in terms of bounds on the network pre-activations (intermediate bounds $\hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k$). Intermediate bounds are computed by solving instances of problem (2) over subsets of the network. For instance, for $\hat{\mathbf{l}}_i[j]$, the lower bound on $\hat{\mathbf{x}}_i[j]$:

$$\begin{aligned}
 \min_{\mathbf{x}, \hat{\mathbf{x}}} \quad & \hat{\mathbf{x}}_i[j] \\
 & \mathbf{x}_0 \in \mathcal{C}(\mathbf{x}), \\
 & \hat{\mathbf{x}}_{k+1} = W_{k+1}\mathbf{x}_k + \mathbf{b}_{k+1} \quad k \in \llbracket 0, i-1 \rrbracket, \\
 & (\mathbf{x}_k, \hat{\mathbf{x}}_k) \in \text{Rel}(\sigma, \hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k) \quad k \in \llbracket 1, i-1 \rrbracket, \\
 & \hat{\mathbf{x}}_k \in [\hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k] \quad k \in \llbracket 1, i-1 \rrbracket,
 \end{aligned} \tag{10}$$

where intermediate bounds until the $(i-1)$ -th layer are needed. Overall, problem (10) needs to be solved twice per neuron: once for the lower bound, once for the upper bound, by flipping the sign of the last layer’s weights. Therefore, intermediate bounds are often computed by relying on looser relaxations than for the output bounding (that is, solving for $\min_{\mathbf{x}, \hat{\mathbf{x}}} \hat{x}_n$). See §4.1 for details on how we compute intermediate bounds within branch-and-bound.

D. Interval Bound Propagation

Interval bound propagation (Gowal et al., 2018; Mirman et al., 2018), a simple application of interval arithmetic (Sunaga, 1958; Hickey et al., 2001) to neural networks, implies solving a version of problem (2) where $\text{Rel}(\sigma, \hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k)$ is the hyper-rectangle depicted in Figure 3. We will use \mathcal{B} to denote the corresponding feasible region. Furthermore, let us write $[\mathbf{x}]_- := \min(\mathbf{x}, \mathbf{0})$ and $\hat{\mathbf{l}}_n := \min_{\mathbf{x}, \hat{\mathbf{x}}} \hat{x}_n$ s.t. $(\mathbf{x}, \hat{\mathbf{x}}) \in \mathcal{B}$. Due to the simplicity of the relaxation, the problem enjoys the following closed form solution:

$$\hat{\mathbf{l}}_n := \min_{\mathbf{x}, \hat{\mathbf{x}}} \hat{x}_n = [W_n]_+ \hat{\mathbf{l}}_{n-1} + [W_n]_- \hat{\mathbf{u}}_{n-1} + \mathbf{b}_n.$$

Upper bounds can be computed by replacing $\hat{\mathbf{l}}_{n-1}$ with $\hat{\mathbf{u}}_{n-1}$, and viceversa. Note that the bounds at layer n only depend on the intermediate bounds at layer $(n-1)$. Therefore, output bounds and all intermediate bounds can be computed at once by forward-propagating the bounds from the first layers, at the total cost of four network evaluations. Let d be the input dimensionality: $\mathbf{x}_0 \in \mathbb{R}^d$. IBP is significantly less expensive than relaxations based on linear bounds (Wong & Kolter, 2018; Zhang et al., 2018), which incur a cost equivalent to $O(d)$ network evaluations for the same computation (Xu et al., 2020).

E. UPB Branching: Beta-CROWN Dual and Planet relaxation

Our UPB (see §4.2) branching strategy makes use of variables from the Beta-CROWN (Wang et al., 2021) dual objective. We will now describe the objective as well as its relationship to the Planet relaxation (see §2.1).

E.1. Planet Relaxation

Let us denote by $\text{Conv}(\sigma, \hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k)$ the element-wise convex hull of the activation function, as a function of intermediate bounds. For ReLUs, this corresponds to the Planet relaxation, which is depicted in Figure 1 for the ambiguous case. By replacing $\text{Rel}(\sigma, \hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k)$ with $\text{Conv}(\sigma, \hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k)$ in problem (2), we obtain:

$$\begin{aligned}
 \min_{\mathbf{x}, \hat{\mathbf{x}}} \quad & \hat{x}_n \quad \text{s.t.} \\
 & \mathbf{x}_0 \in \mathcal{C}(\mathbf{x}), \\
 & \hat{\mathbf{x}}_{k+1} = W_{k+1}\mathbf{x}_k + \mathbf{b}_{k+1} \quad k \in \llbracket 0, n-1 \rrbracket, \\
 & (\mathbf{x}_k, \hat{\mathbf{x}}_k) \in \text{Conv}(\sigma, \hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k) \quad k \in \llbracket 1, n-1 \rrbracket, \\
 & \hat{\mathbf{x}}_k \in [\hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k] \quad k \in \llbracket 1, n-1 \rrbracket.
 \end{aligned} \tag{11}$$

Problem (11) can be alternatively written as:

$$\begin{aligned}
 \max_{\alpha \in [0,1]} \min_{\mathbf{x}, \hat{\mathbf{x}}} \quad & \hat{x}_n \quad \text{s.t.} \quad \mathbf{x}_0 \in \mathcal{C}, \\
 & \hat{\mathbf{x}}_{k+1} = W_{k+1}\mathbf{x}_k + \mathbf{b}_{k+1} \quad k \in \llbracket 0, n-1 \rrbracket, \\
 & \mathbf{a}_k(\alpha_k) \odot \hat{\mathbf{x}}_k + \mathbf{b}_k \leq \mathbf{x}_k \leq \bar{\mathbf{a}}_k \odot \hat{\mathbf{x}}_k + \bar{\mathbf{b}}_k \quad k \in \llbracket 1, n-1 \rrbracket, \\
 & \hat{\mathbf{x}}_k \in [\hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k] \quad k \in \llbracket 1, n-1 \rrbracket,
 \end{aligned} \tag{12}$$

where the coefficients are defined as follows:

$$\begin{aligned}
 \mathbf{a}_k(\alpha_k) = \alpha_k, \quad \bar{\mathbf{a}}_k = \frac{\hat{\mathbf{u}}_k}{\hat{\mathbf{u}}_k - \hat{\mathbf{l}}_k}, \quad \bar{\mathbf{b}}_k = \frac{-\hat{\mathbf{l}}_k \odot \hat{\mathbf{u}}_k}{\hat{\mathbf{u}}_k - \hat{\mathbf{l}}_k} \quad & \text{if } \hat{\mathbf{l}}_k \leq 0 \text{ and } \hat{\mathbf{u}}_k \geq 0, \\
 \bar{\mathbf{a}}_k = \mathbf{a}_k(\alpha_k) = 0 \quad & \text{if } \hat{\mathbf{u}}_k \leq 0, \\
 \bar{\mathbf{a}}_k = \mathbf{a}_k = 1 \quad & \text{if } \hat{\mathbf{l}}_k \geq 0, \\
 \bar{\mathbf{b}}_k = 0 \quad & \text{in all cases.}
 \end{aligned} \tag{13}$$

For ambiguous ReLUs, $\bar{\mathbf{a}}_k \odot \hat{\mathbf{x}}_k + \bar{\mathbf{b}}_k$, represents the upper bounding constraint of Figure 1. The bias of such constraint, $\bar{\mathbf{b}}_k = \frac{-\hat{\mathbf{l}}_k \odot \hat{\mathbf{u}}_k}{\hat{\mathbf{u}}_k - \hat{\mathbf{l}}_k}$, gives the name to our branching strategy (Upper Planet Bias).

E.2. Beta-CROWN Dual

Within branch-and-bound for ReLU networks, the constraints of the form $\hat{\mathbf{x}}_k \in [\hat{\mathbf{l}}_k, \hat{\mathbf{u}}_k]$ can be usually omitted from (12), except when they capture the additional constraints associated to the domain splits (split constraints). For simplicity, we will enforce split constraints on all stable neurons, regardless of whether stability comes from actual split constraints or held before splitting. In this context, using (Salman et al., 2019b, equations (8), (9), (38)), the Lagrangian relaxation of problem (12) can be written as follows:

$$\begin{aligned}
 \max_{\alpha \in [0,1], \mu, \lambda, \beta} \min_{\mathbf{x}, \hat{\mathbf{x}}} \quad & \begin{aligned} & W_n \mathbf{x}_{n-1} + b_n + \sum_{k=1}^{n-1} \boldsymbol{\mu}_k^T (\hat{\mathbf{x}}_k - W_k \mathbf{x}_{k-1} - \mathbf{b}_k) \\ & + \sum_{k=1}^{n-1} [\boldsymbol{\lambda}_k]_-^T (\mathbf{x}_k - (\mathbf{a}_k(\alpha_k) \odot \hat{\mathbf{x}}_k + \mathbf{b}_k)) + \sum_{k=1}^{n-1} [\boldsymbol{\lambda}_k]_+^T (\mathbf{x}_k - (\bar{\mathbf{a}}_k \odot \hat{\mathbf{x}}_k + \bar{\mathbf{b}}_k)) \\ & + \sum_{k=1}^{n-1} \boldsymbol{\beta}_k^T \mathbf{x}_k \mathbb{1}_{\hat{\mathbf{u}}_k \leq 0} - \sum_{k=1}^{n-1} \boldsymbol{\beta}_k^T \mathbf{x}_k \mathbb{1}_{\hat{\mathbf{l}}_k \geq 0} \end{aligned} \\
 \text{s.t.} \quad & \mathbf{x}_0 \in \mathcal{C}.
 \end{aligned} \tag{14}$$

By enforcing the coefficient of the unconstrained \mathbf{x} and $\hat{\mathbf{x}}$ terms to be null, we obtain the following:

$$\begin{aligned}
 d_P &= \max_{\alpha \in [0, 1], \beta \geq 0} \left\{ \min_{\mathbf{x}_0 \in \mathcal{C}} \left(-\bar{\boldsymbol{\mu}}_1^T W_1 \mathbf{x}_0 \right) + b_n - \sum_{k=1}^{n-1} \left([\bar{\boldsymbol{\lambda}}_k]_-^T \mathbf{b}_k + [\bar{\boldsymbol{\lambda}}_k]_+^T \bar{\mathbf{b}}_k + \bar{\boldsymbol{\mu}}_k^T \mathbf{b}_k \right) \right\} \\
 \text{s.t.} \quad &\bar{\boldsymbol{\lambda}}_{n-1} = -W_n^T, \\
 &\bar{\boldsymbol{\mu}}_k = \bar{\mathbf{a}}_k \odot [\bar{\boldsymbol{\lambda}}_k]_+ + \boldsymbol{\alpha}_k \odot [\bar{\boldsymbol{\lambda}}_k]_- + \mathbf{s}_k \odot \boldsymbol{\beta}_k \quad k \in \llbracket 1, n-1 \rrbracket, \\
 &\bar{\boldsymbol{\lambda}}_{k-1} = W_k^T \bar{\boldsymbol{\mu}}_k \quad k \in \llbracket 2, n-1 \rrbracket,
 \end{aligned} \tag{15}$$

where: $\mathbf{s}_k = \mathbb{1}_{\hat{\mathbf{u}}_k \geq 0} - \mathbb{1}_{\hat{\mathbf{u}}_k \leq 0}$.

Problem (15) corresponds to the β -CROWN objective, as it can be easily seen by comparing it with (Wang et al., 2021, equation (20)) and pointing out that, in their formulation, the input domain represents ℓ_∞ norm perturbations of radius ϵ . The dual variables β , which give the name to the algorithm, are necessary only for the neurons whose domains have been split within branch-and-bound (Wang et al., 2021).

E.3. UPB Branching

The key observation behind our UPB branching strategy is that the $[\bar{\boldsymbol{\lambda}}_k]_+^T \bar{\mathbf{b}}_k$ term is present only for ambiguous neurons (see the coefficients in equation (13), note that $[\bar{\boldsymbol{\lambda}}_k]_-^T \mathbf{b}_k = 0$). We can hence heuristically employ it as a proxy for the improvement that a split constraint will have on d_P from (15). Replacing \mathbf{b}_k with its definition in $[\bar{\boldsymbol{\lambda}}_k]_+^T \bar{\mathbf{b}}_k$ yields the branching scores $s_{\text{UPB},k}$ in equation (7). Given the values for α and β obtained in the last bounding step within branch-and-bound, $[\bar{\boldsymbol{\lambda}}_k]_+$, and hence the scores, can be computed using their definition in problem (15), which has a cost equivalent to a single gradient backpropagation through the network. The UPB branching strategy then proceeds by enforcing split constraints on the neuron associated to the largest $s_{\text{UPB},k}$ score throughout the network.

F. Experimental Details

We now present experimental details that were omitted from §6. In particular, we describe the computational setup, network architectures, employed hyper-parameters, and details concerning the Jax porting of the COLT algorithm (Balunovic & Vechev, 2020).

F.1. Experimental Setting, Hyper-parameters

All the experiments were run on a single GPU, either an Nvidia Titan Xp, or an Nvidia Titan V. The timing experiments for verified training were all run on an Nvidia Titan V GPU, on a machine with a 20-core Intel i9-7900X CPU. The branching experiments were instead consistently run on an Nvidia Titan XP GPU, on a machine with a 12-core Intel i7-6850K CPU.

Training hyper-parameters COLT was run with the hyper-parameters provided by the authors (Balunovic & Vechev, 2020), whereas the hyper-parameters for IBP-R are listed in Figure 4. Note that the learning rate is annealed, at each epoch, only after the mixing phase of training. Please refer to equation (6) for the meaning of the various hyper-parameters. We did not tune the parameters of the PGD attacks, which were set as for COLT: we report them for convenience.

Verification hyper-parameters We now complement section §4.1 with omitted details concerning the configuration of the OVAL branch-and-bound framework (Bunel et al., 2020a;b; De Palma et al., 2021c). These details apply to both the training and the branching experiments. The entire verification procedure is run on a single GPU. For the UPB and SR branching strategies, the branching and bounding steps are performed in parallel for batches of 600 and 1200 subproblems, respectively. For the FSB branching strategy, these numbers were reduced to 500 and 1000 subproblems, respectively, in order to prevent PyTorch out-of-memory errors. β -CROWN is run with a dynamically adjusted number of iterations (see (De Palma et al., 2021c, section 5.1.2)) of the Adam optimizer (Kingma & Ba, 2015), with a learning rate of 0.1, decayed by 0.98 at each iteration. Similarly, α -CROWN for the intermediate bounds is run with Adam for 5 iterations, a learning rate of 1, decayed by 0.98 at each iteration. Early termination is triggered when an exponential moving average of the expected branching improvement, computed on the subproblem with the smallest lower bound within the current sub-problem batch, suggests that the decision threshold will be crossed after the timeout. The time to deplete the sub-problem queue (estimated via the runtime per bounding batch) is also added to the estimated time to termination.

Hyperparameter	Value	Hyperparameter	Value
IBP-R with and without masking		IBP-R with and without masking	
Optimizer	SGD	Optimizer	SGD
Learning rate	10^{-2}	Learning rate	10^{-2}
Learning rate exponential decay	0.95	Learning rate exponential decay	0.95
Batch size	100	Batch size	150
Total training steps	800	Total training steps	800
Mixing steps	600	Mixing steps	600
PGD attack steps	8	PGD attack steps	8
PGD attack step size	0.25	PGD attack step size	0.25
α	2.1	λ	10^5
λ	2×10^5	$\frac{\rho}{2}$	10^4
$\frac{\rho}{2}$	10^4	IBP-R	
		α	1.7
		$\frac{\rho}{2}$	5×10^3
		IBP-R with Masking	
		α	1.6
		$\frac{\rho}{2}$	10^2

(a) $\epsilon_{\text{ver}} = 2/255$.

(b) $\epsilon_{\text{ver}} = 8/255$.

Figure 4: IBP-R hyper-parameters for the experiment of table 1.

Network Specifications	Network Architecture
Perturbation radius: $\epsilon_{\text{ver}} = 2/255$ Activation: ReLU Total activations: 49402 Total parameters: 2133736	Conv2d(3, 32, 3, stride=1, padding=1) Conv2d(32,32,4, stride=2, padding=1) Conv2d(32,128,4, stride=2, padding=1) linear layer of 250 hidden units linear layer of 10 hidden units
Perturbation radius: $\epsilon_{\text{ver}} = 8/255$ Activation: ReLU Total activations: 16643 Total parameters: 2118856	Conv2d(3, 32, 5, stride=2, padding=2) Conv2d(32,128,4, stride=2, padding=1) linear layer of 250 hidden units linear layer of 10 hidden units

Table 3: Specifications of the employed network architectures for the experiments of §6.

Network architectures Table 3 reports the details of the employed network architectures for both the training and the branching experiments.

F.2. Jax Porting of COLT

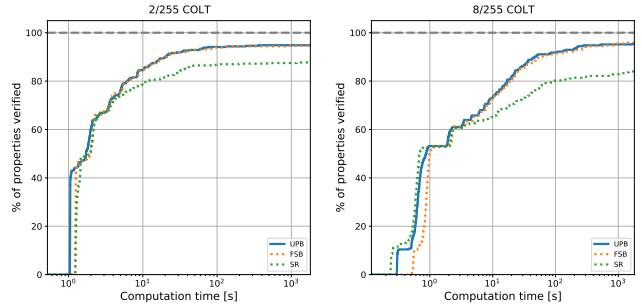
In order to ensure a fair timing comparison with our IBP-R, we ported COLT, whose original implementation is in PyTorch (Paszke et al., 2019), to Jax (Bradbury et al., 2018). The porting resulted in speed-up factors of around 5.7 and 4.4 for the $\epsilon_{\text{ver}} = 2/255$ and $\epsilon_{\text{ver}} = 8/255$ experiments, respectively: see table 1. The standard accuracy results were similar to the original implementation (see table 1), testifying the validity of the porting. The verified accuracy of our experiments is larger than the one reported in the literature: this is likely due to the different, and arguably more effective, verification procedure that we employed.

We conclude this subsection by reporting the main functional difference of our Jax porting with respect to the original implementation. COLT’s Cauchy random projections-approximated zonotope intermediate bounds require extensive use of median values computed over arrays. The median values are employed as estimators for the ℓ_1 norm (Li et al., 2007) of the zonotope propagation matrices. Unfortunately, these operations do not scale very well in Jax (Bradbury et al., 2020). Therefore, we rely on a different ℓ_1 estimator, based on the geometric mean (Li et al., 2007), and we clip the employed Cauchy samples to ensure numerical stability during training.

Method	$\epsilon_{\text{ver}} = 2/255$ COLT			$\epsilon_{\text{ver}} = 8/255$ COLT		
	time [s]	subproblems*	%Timeout	time [s]	subproblems*	%Timeout
UPB	98.87	257.17	5.16	100.19	2924.80	4.46
FSB	104.41	379.77	5.41	101.00	2840.52	4.09
SR	229.77	5894.53	12.29	318.21	20916.78	15.99

* computed on the properties that did not time out for neither UPB nor FSB. The inclusion of timed-out results in the average leads to an overestimation of the number of subproblems for the less expensive branching strategy.

(a) Comparison of average runtime, average number of solved subproblems and the percentage of timed out properties. The best performing method is highlighted in bold.



(b) Cactus plots: percentage of solved properties as a function of runtime. Baselines are represented by dotted lines.

Figure 5: Complete verification performance of different branching strategies, on two COLT-trained CIFAR-10 networks from §6.1.

G. Supplementary Branching Experiments

We now complement the branching results in §6.2 by repeating the same complete verification experiment on two COLT-trained networks from table 1.

Figure 5(a) confirms the results from §6.2, albeit with reduced margins between UPB and FSB. UPB yields small ($< 6\%$) improvements on the average verification times with respect to FSB. In addition, UPB increases the average number of visited subproblems on the $\epsilon_{\text{ver}} = 8/255$ network, and reduces it for $\epsilon_{\text{ver}} = 2/255$, where it also decreases the number of visited subproblems and timed-out properties. As for the IBP-R-trained networks, UPB yields branching decisions competitive with those of FSB while incurring smaller overheads. Finally, as seen in §6.2, SR is significantly slower than both UPB and FSB.