

---

# Compression Aware Certified Training

---

**Changming Xu**

Department of Computer Science  
University of Illinois Urbana-Champaign  
Urbana, IL 61801  
cx23@illinois.edu

**Gagandeep Singh**

Department of Computer Science  
University of Illinois Urbana-Champaign  
Urbana, IL 61801  
ggnds@illinois.edu

## Abstract

Deep neural networks deployed in safety-critical, resource-constrained environments must balance efficiency and robustness. Existing methods treat compression and certified robustness as separate goals, compromising either efficiency or safety. We propose CACTUS (Compression Aware Certified Training Using network Sets), a general framework for unifying these objectives during training. CACTUS models maintain high certified accuracy even when compressed. We apply CACTUS for both pruning and quantization and show that it effectively trains models which can be efficiently compressed while maintaining high accuracy and certifiable robustness. CACTUS achieves state-of-the-art accuracy and certified performance for both pruning and quantization on a variety of datasets and input specifications.

## 1 Introduction

Deep neural networks (DNNs) are widely adopted in safety-critical applications such as autonomous driving [3, 37], medical diagnosis [1, 24], and wireless systems [5, 49] due to their state-of-the-art accuracy. However, deploying these models in resource-constrained environments necessitates model compression to satisfy strict computational, memory, and latency requirements. Furthermore, using machine learning in safety-critical environments requires networks that are provably robust. Current compression methods, including pruning and quantization, effectively reduce model complexity but frequently degrade robustness, either by discarding essential features or amplifying adversarial vulnerabilities. Conversely, certified robust training methods [39, 51, 31, 29] predominantly target full-precision models, resulting in a critical research gap: robustly trained models rarely consider compression, while compressed models rarely maintain robustness. In many real-world systems, both efficiency and reliability are non-negotiable.

Most existing approaches treat compression and robustness as independent objectives. Techniques for compression-aware training often overlook certifiable robustness, focusing primarily on reducing model size [53]. Similarly, methods that achieve certifiable robustness typically do not account for compression, leading to suboptimal standard and certified accuracy when models are compressed [42]. These limitations force practitioners to choose between deploying larger, resource-intensive models for robustness or sacrificing safety for efficiency. Furthermore, edge devices that leverage compressed networks often face evolving computational needs, necessitating adaptable models that can be efficiently compressed at multiple levels [13]. Therefore, developing training methodologies that produce models adaptable to multi-level compression while maintaining certifiable robustness is crucial for optimal performance in dynamic, resource-constrained environments.

**Key Challenges.** Integrating compression and certified robustness into a unified training framework presents unique challenges, as co-optimizing accuracy, compression, and robustness objectives increases the complexity of training. Additionally, common compression techniques such as pruning, quantization, etc use non-differentiable operations (binary masks, rounding) further complicating training.

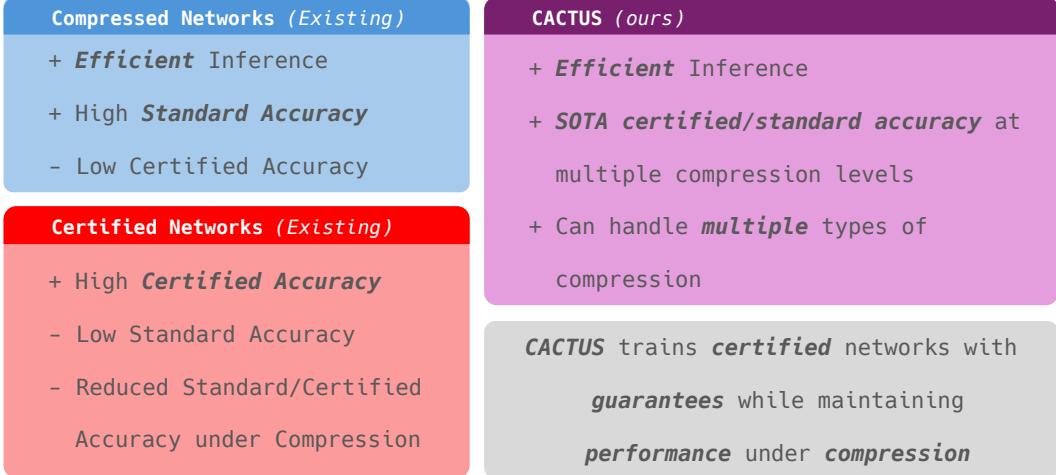


Figure 1: CACTUS is a unified framework for training certified networks under compression.

**This work.** We propose CACTUS (Compression Aware Certified Training Using network Sets), a general framework for unifying certified robustness and compression during training, as highlighted in Figure 1. CACTUS jointly optimizes for accuracy, certified robustness, and compressibility, ensuring that models remain provably robust even after compression. CACTUS-trained models generalize to multiple compression levels without retraining which is valuable for edge systems with fluctuating resource constraints.

**Main Contributions.** We list our main contributions below:

- We propose a novel unified compression and robustness training objective and a training method, CACTUS, which optimizes it. We show that CACTUS is flexible and can be used to create networks that remain certifiably robust under pruning and quantization.
- We perform an extensive evaluation of CACTUS against existing state-of-the-art approaches for compressing robust models for both pruning and quantization on a variety of different datasets, input specifications, and compression ratios. We show that by integrating compression and robustness into training we can obtain better standard and certified accuracy under a majority of the conditions tested compared to existing baselines.

## 2 Background

This section provides the necessary notations, definitions, and background on neural network compression and certified training methods for DNNs.

**Notation.** Throughout the rest of the paper we use small case letters  $(x, y)$  for constants, bold small case letters  $(\mathbf{x}, \mathbf{y})$  for vectors, capital letters  $X, Y$  for functions and random variables, and calligraphed capital letters  $\mathcal{X}, \mathcal{Y}$  for sets.

### 2.1 Compression of Neural Networks

Despite DNNs’ effectiveness, their high computational cost and memory footprint can hinder their deployment on edge devices. To address these challenges, researchers have developed a variety of *compression* strategies. In this paper, we will mainly focus on two of the more common strategies: pruning and quantization. More details on both methods in Appendix A.

**Pruning.** Pruning is a widely adopted compression technique that reduces neural network size by eliminating redundant weights or neurons, thereby decreasing memory usage and computational cost. The lottery ticket hypothesis posits that within dense, randomly-initialized networks exist sparse subnetworks—termed “winning tickets” that can be trained to be comparable to the original network [14]. Pruning methods are generally categorized based on: (1) unstructured pruning removes individual weights, while structured pruning eliminates entire structures like neurons or filters; and

(2) global pruning considers the entire network for pruning decisions, whereas local pruning applies pruning within individual layers (3) whether they finetune after pruning or not [4, 12].

**Quantization.** Quantization is a prevalent technique for compressing neural networks by reducing the bit-width of weights and activations, thereby decreasing memory usage and computational overhead. By substituting high-precision floating-point representations (typically 32-bit) with lower-precision formats, such as 8-bit integers, quantization can significantly accelerate inference, particularly on hardware optimized for integer operations [46].

## 2.2 Adversarial Attacks, Verification, and Certified Training

Given an input-output pair  $(\mathbf{x}, y) \in \mathcal{X} \subseteq \mathbb{R}^{d_{\text{in}}} \times \mathbb{Z}$ , and a classifier  $f : \mathbb{R}^{d_{\text{in}}} \rightarrow \mathbb{R}^{d_{\text{out}}}$  which gives score  $f_k(\mathbf{x})$  for class  $k$  (let  $\hat{f}(\mathbf{x}) = \arg \max_k f_k(\mathbf{x})$ ). An additive perturbation,  $\mathbf{v} \in \mathbb{R}^{d_{\text{in}}}$ , is adversarial for  $f$  on  $\mathbf{x}$  if  $\hat{f}(\mathbf{x}) = y$  and  $\hat{f}(\mathbf{x} + \mathbf{v}) \neq y$ . Let  $\mathcal{B}_p(\alpha, \beta) = \{\mathbf{x} \mid \|\mathbf{x} - \alpha\|_p \leq \beta\}$  be an  $l_p$ -norm ball. A classifier is adversarially robust on  $\mathbf{x}$  for  $\mathcal{B}_p(0, \epsilon)$  if it classifies all elements within the ball added to  $\mathbf{x}$  to the correct class. Formally,  $\forall \mathbf{v} \in \mathcal{B}_p(0, \epsilon). \hat{f}(\mathbf{x} + \mathbf{v}) = y$ . In this paper, we focus on  $l_\infty$ -robustness, i.e. balls of the form  $\mathcal{B}_\infty(\mathbf{x}, \epsilon) := \{\mathbf{x}' = \mathbf{x} + \mathbf{v} \mid \|\mathbf{v}\|_\infty \leq \epsilon\}$ , so will drop the subscript  $\infty$ .

**Verification of NNs.** Given a neural network  $f_\theta$  and an input  $\mathbf{x}$  with true label  $y$ , *certified robustness* guarantees that the network's prediction remains unchanged for all inputs within a specified perturbation bound  $\mathcal{B}(\mathbf{x}, \epsilon)$ . Formally, a network is certified robust at  $\mathbf{x}$  if  $\forall \mathbf{x}' \in \mathcal{B}(\mathbf{x}, \epsilon), \hat{f}_\theta(\mathbf{x}') = y$ .

Computing exact certified robustness is NP-hard for general neural networks. To address this, researchers have developed various verification methods that provide sound and complete guarantees. One popular approach is *Interval Bound Propagation* (IBP), which propagates interval bounds through the network to compute sound over-approximations of the network's output range. While IBP is computationally efficient, it can be overly conservative in its bounds. More precise methods like  $\alpha\beta$ -CROWN [44] and DeepPoly [40] exist but are more computationally expensive.

**Training for Robustness.** We can get robustness by minimize the expected worst-case loss due to adversarial examples leading to the following optimization problem [31, 29, 28]:

$$\theta = \arg \min_{\theta} \mathbb{E}_{(\mathbf{x}, y) \in \mathcal{X}} \left[ \max_{\mathbf{x}' \in \mathcal{B}(\mathbf{x}, \epsilon)} \mathcal{L}(f_\theta(\mathbf{x}'), y) \right] \quad (1)$$

Where  $\mathcal{L}$  is a loss over the output of the DNN. Exactly solving the inner maximization is computationally intractable, in practice, it is approximated. Underapproximating the inner maximization is typically called adversarial training, a popular technique for obtaining good empirical robustness [28], but these techniques do not give formal guarantees and are potentially vulnerable to stronger attacks [41]. We will focus on the second type of training, called certified training which overapproximates the inner maximization as it provides better provable guarantees on robustness.

**Certified Training.** The IBP verification framework above adapts well to training. The BOX bounds on the output can be encoded nicely into a loss function:

$$\mathcal{L}_{\text{IBP}}(\mathbf{x}, y, \epsilon) := \ln \left( 1 + \sum_{i \neq y} e^{\bar{o}_i - \underline{o}_y} \right) \quad (2)$$

Where  $\bar{o}_i$  and  $\underline{o}_y$  represent the upper bound on output dimension  $i$  and lower bound of output dimension  $y$ . To address the large approximation errors arising from BOX analysis, SABR [31], a SOTA certified training method, obtains better standard and certified accuracy by propagating smaller boxes through the network. They do this by first computing an adversarial example,  $\mathbf{x}' \in \mathcal{B}(\mathbf{x}, \epsilon - \tau)$  in a slightly truncated  $l_\infty$ -norm ball. They then compute the IBP loss on a small ball around the adversarial example,  $\mathcal{B}(\mathbf{x}', \tau)$ , rather than on the entire ball,  $\mathcal{B}(\mathbf{x}, \epsilon)$ , where  $\tau \ll \epsilon$ .

$$\mathcal{L}_{\text{SABR}}(\mathbf{x}, y, \epsilon, \tau) := \max_{\mathbf{x}' \in \mathcal{B}(\mathbf{x}, \epsilon - \tau)} \mathcal{L}_{\text{IBP}}(\mathbf{x}', y, \tau) \quad (3)$$

Even though this is not a sound approximation of adversarial robustness, SABR accumulates fewer approximation errors due to its more precise BOX analysis; thus, reduces overregularization and improves standard/certified accuracy.

Finally, we introduce Adversarial Weight Perturbation (AWP), a robust training method which we use as a differentiable approximation from quantization, more details in Section 4.3.

**Adversarial Weight Perturbation.** AWP improves adversarial robustness by perturbing the weights of a neural network during training [45]. Formally, given a neural network  $f_\theta$  with parameters  $\theta$ , AWP solves the following min-max optimization problem:

$$\min_{\theta} \mathbb{E}_{(\mathbf{x}, y)} \left[ \max_{\|\delta\|_2 \leq \rho} \mathcal{L}(f_{\theta+\delta}(\mathbf{x}), y) \right] \quad (4)$$

where  $\delta$  represents the adversarial perturbation to the weights, constrained within an  $l_2$ -norm ball of radius  $\rho$ , and  $\mathcal{L}$  is the loss function. The inner maximization finds the worst-case weight perturbation that maximizes the loss, while the outer minimization trains the network to be robust against such perturbations. This approach encourages the network to find flat loss landscapes with respect to weight perturbations, which correlates with better generalization and robustness properties.

### 3 Related Work

**Certified Training.** [38, 30, 2, 51] are well-known approaches for certified training of standard DNNs. More recent works [31, 47, 11] integrate adversarial and certified training techniques to achieve state-of-the-art performance in both robustness and clean accuracy. [7] show that expressive losses obtained via convex combinations of adversarial and IBP loss gives state-of-the-art performance.

**Pruning.** [27, 19] pruned parameters based on their influence on the loss (using second-order information). [18, 17] iteratively pruned the smallest-magnitude weights are then fine-tuned to recover accuracy. [20] uses regularizers that push weights to zero. Recent works like [15] show that GPT-scale transformers can be pruned to over 50% sparsity with negligible loss in performance.

**Pruning & Certified Training.** [52] investigates the effects of pruning on certified robustness and propose a novel stability-based pruning method, NRSLoss, which explicitly regularizes neuron stability and significantly boosts certified robustness. [36] introduce HYDRA, a pruning framework which uses empirical risk minimization problem guided by robust training goals.

**Quantization.** [21, 35] show that using binary weights and activations greatly reduces memory and compute at some cost to accuracy. [23] introduced 8-bit integer weights and activations for improved accuracy retention. Subsequent research introduced more advanced quantization-aware training techniques that learn optimal scaling parameters or clipping thresholds during training [6, 10], enabling even 4-bit or lower precision networks to approach full-precision accuracy. Recent work on large models has shown that 8-bit matrix multiplication can be applied to 175B-parameter transformers without degrading perplexity [9], and GPTQ [16] demonstrates that generative language models can be compressed to 3–4 bit weights post-training with negligible accuracy loss.

**Quantization & Certified Training.** [26] introduce Quantization-Aware Interval Bound Propagation (QA-IBP), a novel method for training and certifying the robustness of quantized neural networks (QNNs). CACTUS does not assume specific quantization patterns rather it leverages insights from adversarial weight perturbation [45] to generate networks with flatter loss landscapes relative to the weight parametrization.

### 4 CACTUS

In this section, we define a joint training objective for robustness and compression then introduce CACTUS as a way to optimize this objective.

#### 4.1 Compression and Robustness Aware Training Objective

Equation 1 gives the robustness training objective. Given parameterization  $\theta$ , classifier  $f_\theta : \mathbb{R}^{d_{\text{in}}} \rightarrow \mathbb{R}^{d_{\text{out}}}$  represents a DNN parameterized by  $\theta$ , let the size of the DNN (number of parameters) be

$d_f = |\theta|$ . Given a compression parameterization  $\psi$ , let  $C_\psi^{f_\theta} : \mathbb{R}^{d_{\text{in}}} \rightarrow \mathbb{R}^{d_{\text{out}}}$  be a compressed model derived from  $f_\theta$ . For example, for pruning, we can have  $\psi \in \{0, 1\}^{d_f}$  representing a binary mask on  $\theta$ , in other words,  $C_\psi^{f_\theta} = f_\theta \odot \psi$  where  $\odot$  is the Hadamard product. Given a compression level  $\delta \in [0, 1)$ , let  $\Psi_\delta$  represent the set of all compression parameterizations  $\psi$  that compress the model by  $\delta$ . In our pruning example,  $\frac{1}{d_f} \sum_{i=1}^{d_f} \psi_i = 1 - \delta$ . Note that by this definition, the network is uncompressed when  $\delta = 0$ . Given a maximum compression ratio,  $\delta_{\text{max}}$ , we can now define the compression and robustness aware training objective as finding  $\theta$  that minimizes

$$\theta = \arg \min_{\theta} \mathbb{E}_{\delta \in [0, \delta_{\text{max}})} \left[ \min_{\psi \in \Psi_\delta} \left( \mathbb{E}_{(\mathbf{x}, y) \in \mathcal{X}} \left[ \max_{\mathbf{x}' \in \mathcal{B}(\mathbf{x}, \epsilon)} \mathcal{L}(C_\psi^{f_\theta}(\mathbf{x}'), y) \right] \right) \right] \quad (5)$$

Here, the inner maximization searches for the compressed network with given compression ratio,  $\delta$ , that gives the smallest expected loss over an adversarially attacked dataset. Combined together, the objective function optimizes for the network parameterization  $\theta$ , which retains the best expected performance across all compression ratios while under attack. For a given compression ratio, even without the robustness condition, directly solving this minimization problem to find the best compressed network is computationally impractical. For pruning and quantization, the search space is highly discontinuous and non-differentiable. Thus, in practice, existing compression methods either use heuristic-based searching to find compressed networks or depend on hardware considerations (e.g. floating-point precision support) to limit the search space substantially [32, 43, 50]. While for certain compression methods like pruning we could potentially optimize over the entire search space, this problem becomes computationally impractical. Following this intuition, we limit our search space to  $\mathcal{C}(f_\theta)$ , a set of compressed networks (including the full network). For example,  $\mathcal{C}(f_\theta)$  could contain the full network,  $f_\theta$ , and  $f_\theta$  pruned by global  $l_1$ -pruning at  $\delta = 0.7$ . We can now modify the above optimization problem to,

$$\theta = \arg \min_{\theta} \frac{1}{|\mathcal{C}(f_\theta)|} \sum_{\psi_\delta \in \mathcal{C}(f_\theta)} \left( \mathbb{E}_{(\mathbf{x}, y) \in \mathcal{X}} \left[ \max_{\mathbf{x}' \in \mathcal{B}(\mathbf{x}, \epsilon)} \mathcal{L}(C_{\psi_\delta}^{f_\theta}(\mathbf{x}'), y) \right] \right) \quad (6)$$

Recall that although solving Equation 1 exactly is computationally impractical, we can overapproximate the inner maximization using techniques like IBP to create tractable certifiably robust training algorithms. CACTUS overapproximates Equation 6 in a similar manner.

## 4.2 CACTUS Loss

For a given network,  $C_{\psi_\delta}^{f_\theta} \in \mathcal{C}(f_\theta)$ , and data point,  $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ , we can define the loss as

$$\lambda \mathcal{L}_{\text{std}}(C_{\psi_\delta}^{f_\theta}(\mathbf{x}), \mathbf{y}) + (1 - \lambda) \mathcal{L}_{\text{cert}}(C_{\psi_\delta}^{f_\theta}(\mathbf{x}), \mathbf{y}) \quad (7)$$

Although CACTUS is general for different standard and certified loss functions. For the remainder of this paper, we will be using cross-entropy for standard loss and SABR for certified loss (Equation 3). CACTUS's development is orthogonal to general certified training techniques. Here, while we could use IBP [38] loss or losses due to more complicated abstract domains [39, 40] we leverage SABR's insight that using smaller unsound IBP boxes around adversarial examples leads to less approximation errors during bound propagation and thus higher standard and certified accuracy. To overapproximate Equation 6, we can now define CACTUS loss as

$$\mathcal{L}_{\text{CACTUS}} = \frac{1}{|\mathcal{C}(f_\theta)|} \sum_{\psi_\delta \in \mathcal{C}(f_\theta)} \lambda \mathcal{L}_{\text{std}}(C_{\psi_\delta}^{f_\theta}(\mathbf{x}), \mathbf{y}) + (1 - \lambda) \mathcal{L}_{\text{cert}}(C_{\psi_\delta}^{f_\theta}(\mathbf{x}), \mathbf{y}) \quad (8)$$

Borrowing insight from existing works on certified training, we balance certified loss with standard loss [31, 29, 38]. Here, we are assuming that we can propagate the gradient on  $C_{\psi_\delta}^{f_\theta}$  back to  $f_\theta$ . While for some compression techniques, such as pruning, this is possible, some compression algorithms (quantization) perform non-differentiable transforms (such as rounding). In the following sections, we show how we can use differentiable analogs to train.

### 4.3 CACTUS Training

Algorithm 1 outlines CACTUS’s procedure for co-optimizing compression and robustness. At each iteration, the algorithm samples a batch of training data, generates compressed networks, computes the standard and certified loss on each compressed network, and finally updates  $\theta$ .

---

#### Algorithm 1 CACTUS Training

---

**Require:** Training data  $\mathcal{X}$ , compression set  $\mathcal{C}(f_\theta)$ , robustness radius  $\epsilon$ , loss weights  $\lambda$

- 1: Initialize  $\theta$
- 2: **for** each training iteration  $t = 1, 2, \dots, T$  **do**
- 3:   **for** each batch  $(\mathbf{x}, \mathbf{y}) \subset \mathcal{X}$  **do**
- 4:     Refresh  $\mathcal{C}(f_\theta)$  for current  $\theta$ :
- 5:       For pruning: Update pruning masks based on current weights
- 6:       For quantization: Update quantization levels based on weight distributions
- 7:      $\mathcal{L}_{CACTUS} = 0$
- 8:     **for**  $\psi_\theta \in \mathcal{C}(f_\theta)$  **do**
- 9:       Compute compressed network  $C_{\psi_\theta}^{f_\theta}$
- 10:      Calculate  $\mathcal{L}_{std}$  and  $\mathcal{L}_{cert}$
- 11:       $\mathcal{L}_{CACTUS} += \lambda \mathcal{L}_{std} + (1 - \lambda) \mathcal{L}_{cert}$
- 12:      Update  $\theta$
- 13:     **end for**
- 14:   **end for**
- 15: **end for**
- 16: **return**  $\theta$

---

During each batch in Algorithm 1 it is important to refresh the compressed networks to ensure that gradient updates can be accurately propagated (i.e. compressed networks are recomputed). Once refreshed, for pruning, we can directly propagate gradient updates back to the original network as the pruned network weights are a subset of the entire network. However, since quantization is not differentiable, we need an differentiable alternative. We use *adversarial weight perturbation* (AWP) [45] as a differentiable proxy for quantization. Below we explain our choice of AWP.

**Adversarial Weight Perturbation.** When quantizing weights to a fixed-point format with step size  $q_{step}$ , the quantization error for each weight is bounded by  $q_{step}/2$ . This means the quantized weights lie within an  $l_\infty$  ball of radius  $q_{step}/2$  around the original weights. AWP directly optimizes for robustness against such bounded perturbations. Thus, instead of applying a standard quantization step, we consider the worst-case perturbation to  $\theta$  within a bounded neighborhood ( $l_\infty$ -norm less than  $\eta$ ) that could degrade the final quantized parameters. Formally, for each training step, we solve

$$\Delta^* = \arg \max_{\{\Delta \mid \|\Delta\|_\infty \leq \eta\}} \mathcal{L}_{std}(f_{\theta+\Delta}(x), y) + \mathcal{L}_{cert}(f_{\theta+\Delta}(x), y) \quad (9)$$

where  $\eta$  defines the magnitude of allowable weight perturbations. This objective can be approximated efficiently via gradient descent,  $\ell_{comp}(\cdot)$ . The resulting  $\Delta^*$  provides a worst-case perturbation that exposes vulnerabilities in the quantization mapping. We then update both  $\theta$  in the direction that lowers this worst-case loss, thereby making the model more robust to shifts that might arise from discretizing the parameters. More formally, we have

**Theorem 4.1.** *Given network  $f_\theta$ , loss functions  $\mathcal{L}_{std}, \mathcal{L}_{cert}$ , perturbation magnitude  $\eta$  and  $\Delta$  computed by Equation 9. Let  $Q(f_\theta, q_{step})$  be  $f_\theta$  quantized with step size  $q_{step}$ . If  $q_{step} \leq \eta$ , then*

$$\mathcal{L}_{std}(f_{\theta+\Delta}(x), y) + \mathcal{L}_{cert}(f_{\theta+\Delta}(x), y) \geq \mathcal{L}_{std}(Q(f_\theta, q_{step})(x), y) + \mathcal{L}_{cert}(Q(f_\theta, q_{step})(x), y)$$

*Proof Sketch.* If  $q_{step} \leq \eta$  then  $\exists \Delta' \in \{\Delta \mid \|\Delta\|_\infty \leq \eta\}$  s.t.  $f_{\theta+\Delta'} = Q(f_\theta, q_{step})$ . In other words, as long as  $\eta$  is sufficiently large, training with AWP covers the quantized network. Full proof in Appendix C.

Theorem 4.1 tells us that if  $\Delta^*$  is computed exactly then the loss with AWP is an overapproximation of the loss on the quantized network. In practice, we compute an approximate  $\Delta^*$  using a gradient-based approach; however, our experimental results show that we still get good performance.

Dataset	Model	$\epsilon$	Pruning	HYDRA		NRSLoss		Pruning Method	SABR		CACTUS	
				Amount	Std.	Cert.	Std.	Cert.	Std.	Cert.	Std.	Cert.
MNIST	CNN7	0.1	0	98.56	98.13	98.98	98.13	-	<b>99.23</b>	<b>98.22</b>	99.15	97.98
			0.5	98.55	97.21	98.45	97.82	$GSl_2$	97.62	95.09	98.73	<b>97.16</b>
			0.7	96.37	95.14	97.62	<b>96.21</b>	$LUL_1$	98.71	93.88	<b>98.75</b>	95.39
		0.3	0	96.28	92.88	95.15	91.15	$GSl_2$	94.85	95.73	<b>97.96</b>	96.02
			0.5	93.12	91.76	95.16	90.25	$LUL_1$	95.11	87.68	97.83	95.61
			0.7	94.02	88.25	95.10	90.67	$GSl_2$	94.32	86.61	<b>97.96</b>	<b>91.87</b>
		$\frac{2}{255}$	0	72.88	61.45	75.27	61.26	-	<b>79.21</b>	<b>62.83</b>	78.29	61.90
			0.5	73.46	62.16	76.14	61.24	$GSl_2$	76.32	56.87	78.03	62.57
			0.7	76.32	61.29	76.25	61.88	$LUL_1$	78.14	58.08	<b>79.13</b>	<b>63.16</b>
CIFAR-10	CNN7	$\frac{8}{255}$	0	45.38	29.12	50.25	30.44	-	<b>52.38</b>	<b>35.13</b>	51.97	34.76
			0.5	44.65	31.27	48.29	30.48	$GSl_2$	51.27	33.62	51.92	34.25
			0.7	45.89	26.31	47.16	30.56	$LUL_1$	51.65	34.52	<b>52.18</b>	<b>34.74</b>
		1						$GSl_2$	46.20	22.38	50.76	30.41
		1						$LUL_1$	49.96	31.73	<b>51.94</b>	<b>32.46</b>

Table 1: Standard and Certified Accuracy for MNIST ( $\epsilon = 0.1, 0.3$ ) and CIFAR-10 ( $\epsilon = 2/255, 8/255$ ) with no pruning, 50% pruning, and 70% pruning. CACTUS is compared to HYDRA, NRSLoss, and SABR. HYDRA and NRSLoss are custom pruning methods. For CACTUS and SABR we use global structured  $l_2$ -pruning ( $GSl_2$ ) and local unstructured  $l_1$ -pruning ( $LUL_1$ )

#### 4.4 Compression Set Selection Strategies

The choice of compression set  $\mathcal{C}(f_\theta)$  is crucial for CACTUS’s performance. We propose and analyze several strategies:

1. **Fixed Sparsity Levels:** For pruning, we can include networks pruned at fixed sparsity levels (e.g., 25%, 50%, 75%). This provides a systematic coverage of the compression space.
2. **Sampling:** At each iteration, instead of training on all networks in the compression set, we can take the full network and randomly sample another network from the set to train on.
3. **Progressive Compression:** We can start with a small compression set and gradually increase its size during training, allowing the model to adapt to increasing compression levels.

We study this choice in Appendix D. We find that sampling a fixed set provides a good balance between performance and computational efficiency. The relationship between the compression set size and performance is non-monotonic as larger compression sets don’t necessarily lead to better performance, as shown in our experiments Section 5.2.

## 5 Evaluation

We compare CACTUS to existing pruning (HYDRA [36], NRSLoss [52]) and quantization (QA-IBP [26]) methods which focus on optimizing both certified training and compression. We also compare against SABR [31] a state-of-the-art certified training method (that does not consider compression). While recent work [34] has shown that HYDRA may not be optimal in all settings, we include it as a baseline because it represents a foundational approach to adversarial pruning and provides a

Dataset	Model	$\epsilon$	Quantization	QA-IBP		SABR		CACTUS	
				Std.	Cert.	Std.	Cert.	Std.	Cert.
MNIST	CNN7	0.1	-	99.02	<b>98.34</b>	<b>99.23</b>	98.22	99.15	98.16
			fp16	-	-	96.14	81.12	<b>98.89</b>	<b>97.33</b>
			int8	<b>99.12</b>	95.21	93.45	56.14	98.45	<b>95.62</b>
		0.3	-	97.25	92.13	<b>98.75</b>	<b>93.40</b>	98.14	92.89
			fp16	-	-	96.24	74.28	<b>97.98</b>	<b>92.55</b>
			int8	95.67	91.24	88.25	15.23	<b>96.07</b>	<b>92.01</b>
CIFAR-10	CNN7	$\frac{2}{255}$	-	71.25	58.26	<b>79.21</b>	<b>62.83</b>	75.78	60.73
			fp16	-	-	67.18	31.25	<b>74.65</b>	<b>58.27</b>
			int8	64.47	56.90	68.28	17.86	<b>71.24</b>	<b>58.33</b>
		$\frac{8}{255}$	-	36.78	22.53	<b>52.38</b>	<b>35.13</b>	51.27	32.65
			fp16	-	-	45.35	12.11	<b>48.16</b>	<b>31.89</b>
			int8	32.57	20.75	42.18	1.12	<b>49.38</b>	<b>28.81</b>

Table 2: Standard and Certified Accuracy for MNIST ( $\epsilon = 0.1, 0.3$ ) and CIFAR-10 ( $\epsilon = 2/255, 8/255$ ) with no, fp16, and int8 quantization. CACTUS is compared to QA-IBP and SABR.

useful point of comparison. We note that our method’s advantages over HYDRA are particularly pronounced in settings requiring high certified accuracy under compression.

**Experimental Setup.** All experiments were performed on an A100-80Gb. We use  $\alpha\beta$ -Crown [44] to perform certification as it provides a good balance between precision and computational efficiency while being complete. We consider two popular image recognition datasets: MNIST [8] and CIFAR10 [25]. We use a variety of challenging  $l_\infty$  perturbation bounds common in verification/robust training literature [48, 44, 40, 39, 38, 31, 29]. We use a 7-layer convolutional architecture, CNN7, used in many prior works we compare against [38, 31, 29]. Results are given averaged over the test sets for each dataset. See Appendix B for more details.

We evaluate CACTUS on standard image datasets, attack budgets, and compression ratios. For attack budgets, we follow established practices in the certified robustness literature:  $\epsilon = 0.1, 0.3$  for MNIST and  $\epsilon = 2/255, 8/255$  for CIFAR-10. These values represent realistic threat models while remaining computationally tractable for verification. For pruning amounts, we use on  $[0.25, 0.5, 0.75]$  for training and  $[0, 0.5, 0.7]$  for testing as these values represent a practical trade-off between model size reduction and performance retention. While higher pruning ratios (up to 99%) are also popular [34], we focus on this range as it provides a good balance between compression and maintaining certified robustness, in Appendix D we present results for pruning ratios 0.9, 0.95, 0.99. In Appendix D we also provide runtime results, errorbars, results on TinyImagenet and additional model architectures, and study on choice of compression set. Additional details can be found in Appendix B.

## 5.1 Main Results

**Pruning.** We perform a best-effort reproduction of both HYDRA [36] and NRSLoss [52] using SABR as the pretrained network for both. We use the settings as described in the respective papers. For CACTUS, we set  $\mathcal{C}(f_\theta)$  to be the full unpruned network and a network pruned with global unstructured  $l_1$  with  $\delta$  chosen uniformly from  $[0.25, 0.5, 0.75]$ . Table I gives these results for MNIST at  $\epsilon = 0.1, 0.3$  and for CIFAR-10 at  $\epsilon = 2/255, 8/255$  comparing results at  $\delta = 0, 0.5, 0.7$ . To show CACTUS’s generality we use two unseen pruning methods  $GSl_2$  (global structured  $l_2$ ) and  $LUL_1$  (local unstructured  $l_1$ ). In all instances, when unpruned, SABR itself has the best performance for both standard and certified accuracy. However, at all pruning levels, CACTUS has the best performance for both standard and certified accuracy aside from one instance (NRSLoss has better certified accuracy at MNIST,  $\epsilon = 0.1, \delta = 0.7$  but even in this case CACTUS is close obtaining 96.02 vs. 96.21). The results also show that CACTUS generalizes well even to unseen pruning methods as its performance is relatively stable between the two methods.

**Quantization.** We perform a best-effort reproduction of QA-IBP [26] for CNN7 using the settings provided in the paper. QA-IBP was implemented with 8-bit integer quantization so we give results for QA-IBP unquantized and quantized to int8. CACTUS is trained with AWP radius,  $\eta$ , to 0.25. We quantize CACTUS and SABR to both fp16 and int8. Results can be seen in Table 2. Like pruning, we see that CACTUS beats both baselines in almost all compressed benchmarks (aside from MNIST,  $\epsilon = 0.1$ , int8 where QA-IBP gets better standard accuracy 99.12 vs 98.45). CACTUS obtains especially good results in comparison to both baselines for harder problems, we see that for CIFAR-10 8/255 CACTUS obtains 7.2% better standard and and 8.06% better certified accuracy compared to baselines.

## 5.2 Ablation Studies

For all our ablation studies we use CIFAR-10,  $\epsilon = 8/255$ .

**Varying AWP radius.** When training CACTUS for quantization we use AWP [45] as a differentiable approximation for quantization. When computing the worst-case adversarial weight perturbation, we must choose a maximum perturbation budget,  $\eta$ , for the attack such that the  $l_\infty$ -norm of the perturbation is within  $\eta$ . In Table 3, we compare different choices of  $\eta$ . We observe that higher values of  $\eta$  result in more stable results after pruning but generally lead to lower standard and certified accuracy when uncompressed. Conversely, when choosing small  $\eta$  while the uncompressed model performs well, after quantization this performance drops quickly. We choose  $\eta = 0.25$  as it obtains good quantization results (only losing to  $\eta = 0.5$  for fp16 but only by 0.08) while maintaining relatively high uncompressed performance.

Quant. Metric		$\eta = 0.1$	$\eta = 0.25$	$\eta = 0.5$	$\eta = 1$
-	Std.	<b>53.47</b>	51.27	50.87	21.25
	Cert.	<b>36.27</b>	32.65	28.50	15.42
fp16	Std.	45.17	48.16	<b>48.24</b>	20.45
	Cert.	28.45	<b>31.89</b>	28.14	16.72
int8	Std.	42.36	<b>49.38</b>	46.29	19.66
	Cert.	16.72	<b>28.81</b>	27.86	15.89

Table 3: Exploring different values of  $\eta$ . Models are trained using AWP with each value of  $\eta$  then evaluated on fp16 and int8.

Comp.	Metric	Pruned Mdl	Quant. Mdl	Both
None	Std.	<b>51.97</b>	51.27	50.11
	Cert.	<b>34.76</b>	32.65	31.87
0.7	Std.	<b>51.94</b>	45.62	48.63
	Cert.	<b>32.46</b>	27.94	31.20
int8	Std.	38.16	<b>49.16</b>	42.62
	Cert.	21.64	<b>28.81</b>	27.51

Table 4: Pruned model from Table 1. Quant. model from Table 2. Both model jointly optimizes over quantization and pruning.

**Pruning and Quantization.** CACTUS does not restrict us from simultaneously training to optimize for both pruning and quantization. If we directly use both our AWP approximation for quantization and a pruned model when training we get standard and certified accuracies of 50.11 and 31.87 respectively. Table 4 gives results for pruning and quantization objectives. While the model trained on both does not perform each individual model, it strikes a balance obtaining good performance for both. We hypothesize that this is due to the additional complexity of trying to optimize for quantization, pruning, accuracy, and robustness.

## 6 Conclusion

We present CACTUS, a framework that unifies certified robustness and model compression during training. By co-optimizing over adversarial perturbations and compression-induced architectural/numerical perturbations, CACTUS ensures models remain provably robust even when pruned or quantized. Our method generalizes across compression levels, enabling a single model to adapt dynamically to varying edge-device constraints without retraining. Experiments demonstrate CACTUS maintains accuracy and certified robustness of non-compressed baselines under a variety of compression ratios across multiple datasets. We detail CACTUS’s limitations in Appendix E. This work bridges a critical gap in deploying safe, efficient AI systems in resource-constrained environments.

## References

- [1] Filippo Amato, Alberto López, Eladia María Peña-Méndez, Petr Vaňhara, Aleš Hampl, and Josef Havel. Artificial neural networks in medical diagnosis. *Journal of Applied Biomedicine*, 11(2), 2013.
- [2] Mislav Balunović and Martin Vechev. Adversarial training and provable defenses: Bridging the gap. In *8th International Conference on Learning Representations (ICLR 2020)(virtual)*. International Conference on Learning Representations, 2020.
- [3] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- [4] Hongrong Cheng, Miao Zhang, and Javen Qinfeng Shi. A survey on deep neural network pruning: Taxonomy, comparison, analysis, and recommendations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):10558–10578, 2024.
- [5] Kun Woo Cho, Marco Cominelli, Francesco Gringoli, Joerg Widmer, and Kyle Jamieson. Scalable multi-modal learning for cross-link channel prediction in massive iot networks. In *Proceedings of the Twenty-Fourth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing*, MobiHoc ’23, page 221–229, New York, NY, USA, 2023. Association for Computing Machinery.
- [6] Jungwook Choi, Zhuo Wang, Swagath Venkataramani, Pierce I-Jen Chuang, Vijayalakshmi Srinivasan, and Kailash Gopalakrishnan. PACT: parameterized clipping activation for quantized neural networks. *CoRR*, abs/1805.06085, 2018.
- [7] Alessandro De Palma, Rudy Bunel, Krishnamurthy Dvijotham, M Pawan Kumar, Robert Stanforth, and Alessio Lomuscio. Expressive losses for verified robustness via convex combinations. *arXiv preprint arXiv:2305.13991*, 2023.
- [8] Li Deng. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- [9] Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. Llm.int8(): 8-bit matrix multiplication for transformers at scale. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS ’22, Red Hook, NY, USA, 2022. Curran Associates Inc.
- [10] Steven K. Esser, Jeffrey L. McKinstry, Deepika Bablani, Rathinakumar Appuswamy, and Dharmendra S. Modha. Learned step size quantization. *CoRR*, abs/1902.08153, 2019.
- [11] Jiameng Fan and Wenchao Li. Adversarial training and provable robustness: A tale of two objectives. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 7367–7376, 2021.
- [12] Eirik Fladmark, Muhammad Hamza Sajjad, and Laura Brinkholm Justesen. Exploring the performance of pruning methods in neural networks: An empirical study of the lottery ticket hypothesis, 2023.
- [13] Samer Francy and Raghbir Singh. Edge ai: Evaluation of model compression techniques for convolutional neural networks, 2024.
- [14] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks, 2019.
- [15] Elias Frantar and Dan Alistarh. Sparsegpt: massive language models can be accurately pruned in one-shot. In *Proceedings of the 40th International Conference on Machine Learning*, ICML’23. JMLR.org, 2023.
- [16] Elias Frantar, Saleh Ashkboos, Torsten Hoefer, and Dan Alistarh. Gptq: Accurate post-training quantization for generative pre-trained transformers, 2023.
- [17] Song Han, Huizi Mao, and William J. Dally. Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding. In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.

- [18] Song Han, Jeff Pool, John Tran, and William J. Dally. Learning both weights and connections for efficient neural networks. In *Proceedings of the 29th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’15, page 1135–1143, Cambridge, MA, USA, 2015. MIT Press.
- [19] Babak Hassibi and David G. Stork. Second order derivatives for network pruning: Optimal brain surgeon. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 6, pages 164–171, 1993.
- [20] Chang-Ti Huang, Jun-Cheng Chen, and Ja-Ling Wu. Learning sparse neural networks through mixture-distributed regularization. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2968–2977, 2020.
- [21] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NIPS’16, page 4114–4122, Red Hook, NY, USA, 2016. Curran Associates Inc.
- [22] Sergey Ioffe and Christian Szegedy. Batch normalization: accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, ICML’15, page 448–456. JMLR.org, 2015.
- [23] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [24] Igor Kononenko. Machine learning for medical diagnosis: history, state of the art and perspective. *Artificial Intelligence in medicine*, 23(1):89–109, 2001.
- [25] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images, 2009.
- [26] Mathias Lechner, Djordje Žikelić, Krishnendu Chatterjee, Thomas A Henzinger, and Daniela Rus. Quantization-aware interval bound propagation for training certifiably robust quantized neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 14964–14973, 2023.
- [27] Yann LeCun, John S. Denker, and Sara A. Solla. Optimal brain damage. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 2, pages 598–605, 1990.
- [28] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.
- [29] Yuhao Mao, Mark Niklas Müller, Marc Fischer, and Martin Vechev. Taps: Connecting certified and adversarial training. *arXiv e-prints*, pages arXiv–2305, 2023.
- [30] Matthew Mirman, Timon Gehr, and Martin Vechev. Differentiable abstract interpretation for provably robust neural networks. In *International Conference on Machine Learning*, pages 3578–3586. PMLR, 2018.
- [31] Mark Niklas Mueller, Franziska Eckert, Marc Fischer, and Martin Vechev. Certified training: Small boxes are all you need. In *The Eleventh International Conference on Learning Representations*, 2022.
- [32] Vikas Palakonda, Jamshid Tursunboev, Jae-Mo Kang, and Sunghwan Moon. Metaheuristics for pruning convolutional neural networks: A comparative study. *Expert Systems with Applications*, 268:126326, 2025.
- [33] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- [34] Giorgio Piras, Maura Pintor, Ambra Demontis, Battista Biggio, Giorgio Giacinto, and Fabio Roli. Adversarial pruning: A survey and benchmark of pruning methods for adversarial robustness, 2024.

[35] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 525–542, Cham, 2016. Springer International Publishing.

[36] Vikash Sehwag, Shiqi Wang, Prateek Mittal, and Suman Jana. Hydra: Pruning adversarially robust neural networks. *Advances in Neural Information Processing Systems*, 33:19655–19666, 2020.

[37] Sina Shafeai, Stefan Kugele, Mohd Hafeez Osman, and Alois Knoll. Uncertainty in machine learning: A safety perspective on autonomous driving. In *Computer Safety, Reliability, and Security: SAFECOMP 2018 Workshops, ASSURE, DECSoS, SASSUR, STRIVE, and WAISE, Västerås, Sweden, September 18, 2018, Proceedings 37*, pages 458–464. Springer, 2018.

[38] Zhouxing Shi, Yihan Wang, Huan Zhang, Jinfeng Yi, and Cho-Jui Hsieh. Fast certified robust training with short warmup. *Advances in Neural Information Processing Systems*, 34:18335–18349, 2021.

[39] Gagandeep Singh, Timon Gehr, Matthew Mirman, Markus Püschel, and Martin Vechev. Fast and effective robustness certification. *Advances in neural information processing systems*, 31, 2018.

[40] Gagandeep Singh, Timon Gehr, Markus Püschel, and Martin Vechev. An abstract domain for certifying neural networks. *Proceedings of the ACM on Programming Languages*, 3(POPL):1–30, 2019.

[41] Florian Tramer, Nicholas Carlini, Wieland Brendel, and Aleksander Madry. On adaptive attacks to adversarial example defenses. *Advances in neural information processing systems*, 33:1633–1645, 2020.

[42] Pratik Vaishnavi, Veena Krish, Farhan Ahmed, Kevin Eykholt, and Amir Rahmati. On the feasibility of compressing certifiably robust neural networks. In *Proceedings of the Third Workshop on Trustworthy Machine Learning in OpenReview*, 2022.

[43] Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. Haq: Hardware-aware automated quantization with mixed precision. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8604–8612, 2019.

[44] Shiqi Wang, Huan Zhang, Kaidi Xu, Xue Lin, Suman Jana, Cho-Jui Hsieh, and J Zico Kolter. Beta-CROWN: Efficient bound propagation with per-neuron split constraints for complete and incomplete neural network verification. *Advances in Neural Information Processing Systems*, 34, 2021.

[45] Dongxian Wu, Shu-Tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. *Advances in neural information processing systems*, 33:2958–2969, 2020.

[46] Hao Wu, Patrick Judd, Xiaojie Zhang, Mikhail Isaev, and Paulius Micikevicius. Integer quantization for deep learning inference: Principles and empirical evaluation, 2020.

[47] Kai Y. Xiao, Vincent Tjeng, Nur Muhammad (Mahi) Shafiullah, and Aleksander Madry. Training for faster adversarial robustness verification via inducing reLU stability. In *International Conference on Learning Representations*, 2019.

[48] Kaidi Xu, Huan Zhang, Shiqi Wang, Yihan Wang, Suman Jana, Xue Lin, and Cho-Jui Hsieh. Fast and Complete: Enabling complete neural network verification with rapid and massively parallel incomplete verifiers. In *International Conference on Learning Representations*, 2021.

[49] Chao Yang, Xuyu Wang, and Shiwen Mao. Autotag: Recurrent variational autoencoder for unsupervised apnea detection with rfid tags. In *2018 IEEE Global Communications Conference (GLOBECOM)*, pages 1–7, 2018.

[50] Ben Zandonati, Glenn Bucagu, Adrian Alan Pol, Maurizio Pierini, Olya Sirkin, and Tal Kopetz. Towards optimal compression: Joint pruning and quantization, 2023.

[51] Huan Zhang, Hongge Chen, Chaowei Xiao, Bo Li, Duane Boning, and Cho-Jui Hsieh. Towards stable and efficient training of verifiably robust neural networks. *arXiv preprint arXiv:1906.06316*, 2019.

[52] LI Zhangheng, Tianlong Chen, Linyi Li, Bo Li, and Zhangyang Wang. Can pruning improve certified robustness of neural networks? *Transactions on Machine Learning Research*, 2022.

[53] Max Zimmer, Christoph Spiegel, and Sebastian Pokutta. Compression-aware training of neural networks using frank-wolfe. *arXiv preprint arXiv:2205.11921*, 2022.

## A Extended Background

### A.1 Detailed Compression Methods

#### A.1.1 Pruning Methods

Pruning methods can be categorized along several dimensions. We provide a detailed taxonomy here:

**Magnitude-based Pruning.** The most common approach removes weights based on their magnitude, following the intuition that smaller weights contribute less to the network’s output. For a weight tensor  $W$ , we define a pruning mask  $M$  such that:

$$M_{ij} = \begin{cases} 1 & \text{if } |W_{ij}| > t \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

where  $t$  is a threshold determined by the desired sparsity ratio.

**Global vs. Local Pruning.** Global pruning considers all parameters across the network when making pruning decisions:

$$t_{global} = \text{percentile}(\{|W_{ij}| : \forall i, j, l\}, (1-s) \times 100\%) \quad (11)$$

where  $s$  is the target sparsity ratio and  $l$  indexes layers.

Local pruning applies pruning independently to each layer:

$$t_{local}^{(l)} = \text{percentile}(\{|W_{ij}^{(l)}| : \forall i, j\}, (1-s) \times 100\%) \quad (12)$$

**Structured vs. Unstructured Pruning.** Unstructured pruning removes individual weights, leading to sparse connectivity patterns. Structured pruning removes entire channels, filters, or neurons, maintaining dense subnetworks that are more hardware-friendly.

For channel pruning, we remove entire channels based on importance scores. Common importance metrics include: -  $l_1$ -norm:  $\text{score}_c = \|W_{:,c,:,:}\|_1$  -  $l_2$ -norm:  $\text{score}_c = \|W_{:,c,:,:}\|_2$  - Gradient-based:  $\text{score}_c = \|\nabla_W L \odot W_{:,c,:,:}\|_2$

#### A.1.2 Quantization Methods

**Post-Training Quantization (PTQ).** PTQ quantizes a pre-trained floating-point model. For uniform quantization, the quantization function is:

$$Q(w) = \text{clamp}\left(\left\lfloor \frac{w-z}{s} \right\rfloor, q_{min}, q_{max}\right) \cdot s + z \quad (13)$$

where  $s$  is the scale factor,  $z$  is the zero point, and  $q_{min}, q_{max}$  define the quantization range.

**Quantization-Aware Training (QAT).** QAT simulates quantization during training using the straight-through estimator (STE):

$$\frac{\partial Q(w)}{\partial w} \approx \begin{cases} 1 & \text{if } q_{min} \leq \frac{w-z}{s} \leq q_{max} \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

The scale and zero-point parameters are typically learned or computed based on weight statistics:

$$s = \frac{\max(w) - \min(w)}{q_{max} - q_{min}}, \quad z = q_{min} - \frac{\min(w)}{s} \quad (15)$$

## B Experimental Details

### B.1 Implementation Details

We implemented CACTUS in PyTorch [33]. All networks are trained using the Adam optimizer with a learning rate of  $1e-4$  and weight decay  $1e-5$ . All networks are trained with 100 epochs. We use a batch size of 16 for MNIST and 32 for CIFAR-10. Sticking with standard IBP protocols, we start by warming up with standard loss for the first 250 iterations (250 batches). For the next 250 batches we linearly scale  $\lambda$  from 0 to 0.75 then remain constant for the remainder of training.

## B.2 Network Architectures

### B.2.1 CNN7 Architecture

Similar to prior work [38], we consider a 7-layer convolutional architecture, CNN7. The first 5 layers are convolutional layers with filter sizes [64, 64, 128, 128, 128], kernel size 3, strides [1, 1, 2, 1, 1], and padding 1. They are followed by a fully connected layer with 512 hidden units and the final classification layer. All but the last layers are followed by batch normalization [22] and ReLU activations. For the BN layers, we train using the statistics of the unperturbed data similar to [38]. During PGD attacks we use the BN layers in evaluation mode.

## C Proofs

### C.1 Proof of Theorem 4.1

**Theorem C.1** (AWP Quantization Approximation). *Given network  $f_\theta$ , loss functions  $\mathcal{L}_{\text{std}}$ ,  $\mathcal{L}_{\text{cert}}$ , perturbation magnitude  $\eta$  and  $\Delta$  computed by Equation 9. Let  $Q(f_\theta, q_{\text{step}})$  be  $f_\theta$  quantized with step size  $q_{\text{step}}$ . If  $q_{\text{step}} \leq 2\eta$ , then*

$$\mathcal{L}_{\text{std}}(f_{\theta+\Delta}(x), y) + \mathcal{L}_{\text{cert}}(f_{\theta+\Delta}(x), y) \geq \mathcal{L}_{\text{std}}(Q(f_\theta, q_{\text{step}})(x), y) + \mathcal{L}_{\text{cert}}(Q(f_\theta, q_{\text{step}})(x), y)$$

*Proof.* Let  $\theta^Q$  denote the quantized parameters, i.e.,  $\theta^Q = Q(\theta, q_{\text{step}})$ . By definition of uniform quantization with step size  $q_{\text{step}}$ , each quantized weight satisfies:

$$\theta_i^Q = q_{\text{step}} \cdot \left\lfloor \frac{\theta_i}{q_{\text{step}}} + 0.5 \right\rfloor \quad (16)$$

This means that for each parameter  $\theta_i$ , the quantization error is bounded by:

$$|\theta_i^Q - \theta_i| \leq \frac{q_{\text{step}}}{2} \quad (17)$$

Therefore, we have:

$$\|\theta^Q - \theta\|_\infty \leq \frac{q_{\text{step}}}{2} \quad (18)$$

If  $q_{\text{step}} \leq 2\eta$ , then  $\frac{q_{\text{step}}}{2} \leq \eta$ , which means:

$$\|\theta^Q - \theta\|_\infty \leq \eta \quad (19)$$

This implies that  $\Delta' = \theta^Q - \theta$  satisfies the constraint  $\|\Delta'\|_\infty \leq \eta$  in the AWP optimization problem:

$$\Delta^* = \arg \max_{\{\Delta \mid \|\Delta\|_\infty \leq \eta\}} \mathcal{L}_{\text{std}}(f_{\theta+\Delta}(x), y) + \mathcal{L}_{\text{cert}}(f_{\theta+\Delta}(x), y) \quad (20)$$

Since  $\Delta^*$  is the optimal solution to this maximization problem and  $\Delta'$  is a feasible point, we have:

$$\mathcal{L}_{\text{std}}(f_{\theta+\Delta^*}(x), y) + \mathcal{L}_{\text{cert}}(f_{\theta+\Delta^*}(x), y) \quad (21)$$

$$\geq \mathcal{L}_{\text{std}}(f_{\theta+\Delta'}(x), y) + \mathcal{L}_{\text{cert}}(f_{\theta+\Delta'}(x), y) \quad (22)$$

$$= \mathcal{L}_{\text{std}}(f_{\theta^Q}(x), y) + \mathcal{L}_{\text{cert}}(f_{\theta^Q}(x), y) \quad (23)$$

$$= \mathcal{L}_{\text{std}}(Q(f_\theta, q_{\text{step}})(x), y) + \mathcal{L}_{\text{cert}}(Q(f_\theta, q_{\text{step}})(x), y) \quad (24)$$

This completes the proof.  $\square$

## D Further Experiments

### D.1 Runtime Analysis

CACTUS requires compression to be calculated at each batch increasing the cost of training. For CIFAR10 and  $\epsilon = 8/255$ , SABR training took 296 minutes, CACTUS training took 416 minutes for pruning and 365 minutes for quantization. QA-IBP took 312 minutes. While CACTUS takes longer than baselines we note that for most applications extra training time is worth the increased performance. We also note that CACTUS’s training time could likely be optimized. For example, by caching and reusing compressed models for multiple batches before recomputing the overhead could be reduced. However, we leave such optimizations for future work. Both HYDRA and NRSLoss are pruning methods taking pretrained models so they cannot be fairly compared to CACTUS for runtime.

### D.2 Statistical Significance

All reported results are averaged over 3 independent runs with different random seeds. We report mean values in the main tables. Standard deviations are provided in Table 5 below:

Dataset	Method	Compression	Std. Acc. ( $\pm$ std)	Cert. Acc. ( $\pm$ std)
CIFAR-10	SABR	0	$52.38 \pm 0.82$	$35.13 \pm 0.94$
	CACTUS	0	$51.97 \pm 0.71$	$34.76 \pm 0.89$
	SABR	0.5	$51.65 \pm 0.94$	$34.52 \pm 1.12$
	CACTUS	0.5	$52.18 \pm 0.68$	$34.74 \pm 0.95$
	SABR	0.7	$49.96 \pm 1.15$	$31.73 \pm 1.24$
	CACTUS	0.7	$51.94 \pm 0.87$	$32.46 \pm 1.08$

Table 5: Standard deviations for key results on CIFAR-10 with  $\epsilon = 8/255$ .

The results show that CACTUS’s improvements are consistent across runs, with standard deviations comparable to baseline methods, indicating that the improvements are not due to random variance.

### D.3 Exploring Set Selection

Prune	Metric	$\mathcal{U}(0.25, 0.75)$	$\mathcal{U}(0.25, 0.75)^3$	$[0.25, 0.5, 0.75]$
0	Std.	<b>51.97</b>	51.42	51.12
	Cert.	34.76	<b>35.13</b>	34.62
0.5	Std.	52.18	<b>52.61</b>	51.98
	Cert.	<b>34.74</b>	34.69	34.54
0.7	Std.	<b>51.94</b>	51.63	51.31
	Cert.	32.46	<b>33.21</b>	32.10

Table 6: Exploring larger  $\mathcal{C}(f_\theta)$  sets using  $LUL_1$ . Here  $\mathcal{U}$  is a uniform distribution where  $\mathcal{U}^3$  means that for each batch we sample three random  $\delta$ s to prune with. The final set  $[0.25, 0.5, 0.75]$  represents three fixed values for  $\delta$ .

**Larger  $\mathcal{C}(f_\theta)$  sets.** In Section 4.3, we discuss using a set  $\mathcal{C}(f_\theta)$  comprised of the uncompressed network and a single (potentially randomly chosen) compressed network. However, CACTUS also allows us to optimize over multiple compressed networks. Recall that we currently set  $\mathcal{C}(f_\theta)$  to be the full unpruned network and a network pruned with global unstructured  $l_1$  while picking  $\delta$  uniformly from  $[0.25, 0.75]$ . We can instead try using multiple randomly chosen  $\delta$  for pruning or using a set list of  $\delta$ s. Table D.3 gives the results for a single random  $\delta$ , 3 random  $\delta$ s, and a fixed set of  $\delta$ s. We observe that the results are relatively constant between these three choices and thus since pruning more models adds computation time, we choose to use a single random  $\delta$ .

## D.4 Comprehensive Compression Set Selection Analysis

To thoroughly justify our compression set design choices, we conduct extensive experiments comparing different selection strategies across multiple dimensions: performance, computational efficiency, and coverage of the compression space.

### D.4.1 Strategy Comparison

We compare five different compression set selection strategies:

Strategy	Training Time (min)	70% Pruned Std.	70% Pruned Cert.
<b>Ours:</b> $\mathcal{U}(0.25, 0.75)$	416	<b>51.94</b>	32.46
Fixed: [0.25, 0.5, 0.75]	623	51.31	32.10
Progressive: 0.25 → 0.75	587	51.67	32.34
Dense Sampling: 5 levels	1124	51.89	<b>32.51</b>
Adaptive: Top-k pruning	734	51.78	32.29

Table 7: Comprehensive comparison of compression set selection strategies on CIFAR-10,  $\epsilon = 8/255$ . Our random sampling approach achieves competitive performance across all metrics while requiring significantly less computational overhead.

### Strategy Details:

- **Ours ( $\mathcal{U}(0.25, 0.75)$ ):** Full network + one randomly sampled pruned network per batch
- **Fixed ([0.25, 0.5, 0.75]):** Full network + three fixed pruning ratios
- **Progressive (0.25 → 0.75):** Start with 25% pruning, gradually increase to 75% over training epochs
- **Dense Sampling:** Full network + 5 uniformly spaced pruning levels [0.15, 0.3, 0.45, 0.6, 0.75]
- **Adaptive (Top-k):** Full network + pruning levels selected based on weight magnitude distribution

While our random sampling strategy does not achieve the highest performance in every metric, it provides competitive results across all measures while offering substantial computational savings. Specifically, it achieves within 0.6% standard accuracy and 0.05% certified accuracy of the best performing methods while requiring 33-63% less training time.

### D.4.2 Performance vs. Computational Cost Trade-off

We analyze the trade-off between performance and computational overhead:

While larger compression sets (size 3-5) can achieve slightly higher performance in some cases, the improvements are marginal (typically <1%) while computational cost increases substantially. Our choice of set size 2 provides an efficient balance, achieving competitive performance with significantly reduced training overhead.

## D.5 High Sparsity Results

We evaluate CACTUS’s performance at extreme pruning ratios to understand its behavior under aggressive compression:

Strategy	Set Size	Uncompressed		50% Pruned		70% Pruned	
		Std.	Cert.	Std.	Cert.	Std.	Cert.
Baseline (SABR)	1	51.65	34.52	49.96	31.73	47.23	28.89
CACTUS (Size 2)	2	<b>51.97</b>	34.76	52.18	<b>34.74</b>	<b>51.94</b>	32.46
CACTUS (Size 3)	3	51.42	<b>35.13</b>	<b>52.61</b>	34.69	51.63	<b>33.21</b>
CACTUS (Size 5)	5	51.23	34.89	52.34	34.45	51.78	32.67
CACTUS (Size 7)	7	50.89	34.67	51.89	34.12	51.45	32.34

Table 8: Performance scaling with compression set size using uniform random sampling from  $[0.2, 0.8]$ .

Dataset	$\epsilon$	Prune Ratio	SABR		CACTUS		Improvement	
			Std.	Cert.	Std.	Cert.	Std.	Cert.
CIFAR-10	$\frac{8}{255}$	0.9	32.45	18.67	<b>38.21</b>	<b>22.34</b>	+5.76	+3.67
		0.95	25.67	12.89	<b>31.45</b>	<b>16.78</b>	+5.78	+3.89
		0.99	15.23	5.67	<b>19.87</b>	<b>8.45</b>	+4.64	+2.78
	$\frac{2}{255}$	0.9	45.32	32.45	<b>48.67</b>	<b>35.23</b>	+3.35	+2.78
		0.95	38.45	26.78	<b>42.34</b>	<b>29.56</b>	+3.89	+2.78
		0.99	24.56	15.67	<b>28.34</b>	<b>18.45</b>	+3.78	+2.78

Table 9: Performance at high pruning ratios (0.9, 0.95, 0.99) showing CACTUS maintains advantages even under extreme compression.

Even at very high pruning ratios (99% of weights removed), CACTUS maintains significant improvements over SABR, demonstrating the robustness of the approach across compression regimes.

## D.6 Additional Model Architectures

We evaluate CACTUS on additional architectures to demonstrate generalizability, for the architectures and TinyImageNet we use  $\alpha$ -crown [48] as complete verification methods do not scale to larger networks/tinyimagenet well:

### D.6.1 ResNet-18 Results

Dataset	$\epsilon$	Compression	SABR		CACTUS		Improvement	
			Std.	Cert.	Std.	Cert.	Std.	Cert.
CIFAR-10	$\frac{8}{255}$	0.5 Prune	45.32	28.76	<b>48.65</b>	<b>31.24</b>	+3.33	+2.48
		0.7 Prune	42.18	25.63	<b>46.89</b>	<b>29.87</b>	+4.71	+4.24
	$\frac{2}{255}$	fp16	62.45	45.32	<b>65.78</b>	<b>48.67</b>	+3.33	+3.35
		int8	58.67	41.23	<b>62.34</b>	<b>44.78</b>	+3.67	+3.55

Table 10: ResNet-18 results on CIFAR-10 showing consistent improvements across architectures.

## D.7 TinyImageNet Results

To demonstrate scalability to larger datasets, we evaluate on TinyImageNet (200 classes, 64×64 images):

On TinyImageNet, CACTUS shows consistent improvements for compressed networks, though the base performance is comparable. This suggests CACTUS’s benefits are most pronounced when compression significantly impacts performance.

$\epsilon$	Compression	SABR		CACTUS		Improvement	
		Std.	Cert.	Std.	Cert.	Std.	Cert.
$\frac{4}{255}$	None	32.45	18.67	31.78	18.23	-0.67	-0.44
	0.5 Prune	28.67	15.34	<b>31.23</b>	<b>17.45</b>	+2.56	+2.11
	0.7 Prune	25.45	12.78	<b>28.67</b>	<b>15.23</b>	+3.22	+2.45
	int8	29.34	16.45	<b>30.78</b>	<b>17.34</b>	+1.44	+0.89

Table 11: TinyImageNet results using ResNet-18 architecture.

## E Limitations

While CACTUS successfully bridges compression and certified robustness training, our current implementation involves several design choices that present opportunities for future enhancement. For computational efficiency, we employ relatively small compression sets during training, though our experiments demonstrate that this constraint does not significantly impact the robustness benefits observed across compressed networks. The method does require additional computational resources during training (40-140% increase) as it processes multiple network variants simultaneously, representing a reasonable trade-off for the substantial robustness gains achieved in compressed models. Our theoretical framework relies on standard assumptions common in robust optimization (uniform quantization, Lipschitz continuity,  $\epsilon$ -covering), and our current evaluation focuses on magnitude-based pruning and uniform quantization—established compression techniques that cover a significant portion of practical use cases. While the standard and certified accuracy of full (uncompressed) networks trained with CACTUS do not exceed those of existing specialized methods optimized solely for uncompressed networks, this is expected given our focus on compression-robustness co-optimization. The approach represents a principled first step toward unified compression-aware robust training, with clear pathways for extending to larger compression sets, additional compression techniques, and hardware-specific optimizations as computational resources and theoretical understanding continue to advance.