

Integer-arithmetic-only Certified Robustness for Quantized Neural Networks

Haowen Lin¹, Jian Lou^{2,3,*}, Li Xiong², Cyrus Shahabi¹

¹University of Southern California ²Emory University ³Xidian University

haowenli@usc.edu jlou@xidian.edu.cn lxiong@emory.edu shahabi@usc.edu

Abstract

Adversarial data examples have drawn significant attention from the machine learning and security communities. A line of work on tackling adversarial examples is certified robustness via randomized smoothing that can provide a theoretical robustness guarantee. However, such a mechanism usually uses floating-point arithmetic for calculations in inference and requires large memory footprints and daunting computational costs. These defensive models cannot run efficiently on edge devices nor be deployed on integer-only logical units such as Turing Tensor Cores or integer-only ARM processors. To overcome these challenges, we propose an integer randomized smoothing approach with quantization to convert any classifier into a new smoothed classifier, which uses integer-only arithmetic for certified robustness against adversarial perturbations. We prove a tight robustness guarantee under ℓ_2 -norm for the proposed approach. We show our approach can obtain a comparable accuracy and $4\times \sim 5\times$ speedup over floating-point arithmetic certified robust methods on general-purpose CPUs and mobile devices on two distinct datasets (CIFAR-10 and Caltech-101).

1. Introduction

Recent works in deep learning have demonstrated that well-trained deep neural networks can easily make wrong predictions with high confidence when a sample is perturbed with a small but adversarially-chosen noise [37, 13, 32]. To defend against these attacks, several works have proposed to develop defensive techniques and improve the robustness of deep neural networks [1, 28, 26, 41]. A recent promising line of work focuses on developing certifiably robust classifiers that promise no adversarially perturbed examples within a certified region can alter the classification result [6, 42, 23]. Such certified defenses provide a rigorous guarantee against norm-bounded perturbation attacks and, more importantly, ensure effectiveness under future stronger at-

tacks [6, 25, 23, 38]. One primary theoretical tool for providing the robustness guarantee is randomized smoothing, which derives a smoothed classifier from the base classifier via injecting designated noises, e.g., Gaussian noise. Multiple repeated inferences through the base classifier (i.e., Monte-Carlo estimation) are required to approximate the smoothed classification result for robustly predicting or certifying a single example.

Despite the promising results achieved by many certified robustness algorithms, existing methods almost exclusively focus on floating-point (FP) represented neural networks. However, the vastly adopted compressed neural network models are considered indispensable when one wishes to deploy the networks on storage-, computing resources- and power consumption-limited platforms such as edge devices, mobile devices, and embedded systems. In practice, one of the most successful and mainstream compression methods is quantization [20, 8, 44, 47]. Quantization is a simple yet effective technique that compresses deep neural networks into smaller sizes by replacing model weights and activations from 32-bit floating-point (FP32) with low-bit precision, e.g., 8-bit integer (int8) [47]. Both storage and computational complexity can be reduced using low-bit quantized neural networks [20]. Moreover, Jacob et al. [20] have proposed an integer-arithmetic-only quantization framework that further accelerates inference by using integer multiplication and accumulation for calculation. Performing inference using integer-arithmetic-only operations has several advantages in real application scenarios. For example, it resolves the limitation that floating-point networks cannot be deployed onto digital computing devices such as the recent Turing Tensor Cores or traditional integer-only ARM processors. Moreover, computing with integer arithmetic significantly reduces computing power, making them attractive for energy-constrained edge deployment and some cost-sensitive cloud data centers [10].

Given the under-studied situation of the certified robustness for quantized neural networks, the following research questions naturally arise: Q1. Are adversarial perturbations still effective on quantized neural networks? Q2. Can we reuse the current certified robustness defenses on quantized

*Corresponding Author.

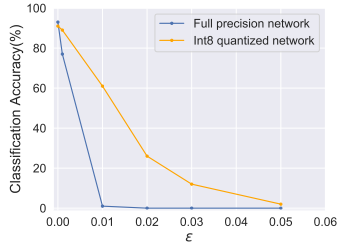


Figure 1. A demonstration of the adversarial perturbation attacks on undefended full precision network (blue) and 8-bit integer network (yellow). The x-axis represents the radius of the noise projected into clean images under ℓ_∞ ball (parameterized by ϵ). Given 100 clean images, the full precision and quantized network achieved 93% and 91% accuracy, respectively. When images are attacked by Projected Gradient Descent attack targeting the full precision model, the accuracy of both classifiers begins to drop. Details are deferred to Supplement.

neural networks? Q3. If not, how can we design a certifiably robust defense making full considerations of the characteristics of quantized neural networks? Q1 and Q2 can be readily answered as follows,

For Q1, we consider the following demonstrating example. We generate adversarial perturbations using Projected Gradient Descent attack [29] and inject the perturbations to 100 randomly selected clean images from CIFAR-10. The result is presented in Figure 1. Although the quantized network manifests slightly stronger robustness than the full precision model, adversarial perturbations can still sabotage the classification performance of the quantized neural network even when the perturbation noise is small. More severely, since the adversarial perturbation used in the example does not consider the characteristics of the quantized neural network as a priori, stronger attacks can be devised once such information is exploited. Thus, it is pressing for us to study certified robustness for quantized neural networks.

For Q2, a new certified robustness mechanism tailored to quantized neural networks is indeed necessary. The reason is that existing certified robustness methods rely on floating-point operations, which are incompatible with integer-arithmetic-only devices. Sometimes even when deployed on platforms that do support floating-point operations, a new integer-arithmetic-only certified robustness method can still be desirable due to the efficiency reason, especially considering that most certified robustness methods invoke repeated inferences to certify on a single example and incur large inference time.

As a result, our main effort in this paper is to answer Q3 with a novel and first integer-arithmetic-only certified robustness mechanism for quantized neural networks against adversarial perturbation attacks. An illustration of our framework and comparison with the existing certified robustness defenses is in Figure 2. In summary, we make the following

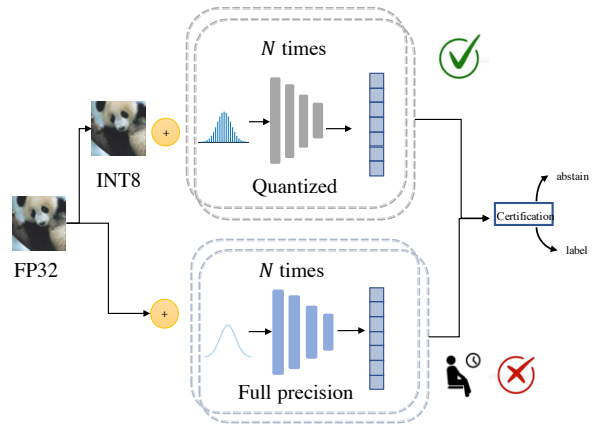


Figure 2. A comparison of certification procedure between IntRS and the floating-point network. Certification requires repeating inference N times. Compared with original floating-point network, IntRS achieves comparable certified accuracy with less inference time.

contributions:

- We devise a new integer-arithmetic-only randomized smoothing mechanism (abbreviated as IntRS) by incorporating the discrete Gaussian noise. More importantly, we rigorously calibrate the certified robustness behavior of IntRS in theory. In practice, we introduce the quantization- and discrete data augmentation- aware as well as the common Monte-Carlo-based estimation.
- We perform experiments with 1) two different base neural network architectures with medium and large scales; 2) two different datasets; 3) two different types of computing devices (general-purpose computer and mobile device), which verify that IntRS achieves similar robustness and accuracy performance, and $4\times$ to $5\times$ efficiency improvement in the inference stage over existing floating-point randomized smoothing method for the original full precision neural networks.

2. Related Work

Certified Defenses. A certifiably robust classifier guarantees a stable classification for any input \mathbf{x} within a certain range (e.g., within an L_p -norm ball) [43]. Certification methods usually fall into two categories: exact and conservative (relaxed). Given a base classifier f and an input \mathbf{x} , exact certification techniques answer the question of whether or not there exists an adversarial perturbation for f at \mathbf{x} . On the other hand, if the conservative algorithms make a certification, it guarantees that there is no perturbation exist, but it could refuse to make certification even for a clean data point. Exact methods, usually based on Mixed Integer Linear Programming [38, 36, 12], and Satisfiability Modulo Theories [22, 2, 21] can be intractable to compute and difficult

to scale up to large and even moderate size neural networks. Conservative certification can be more scalable than exact methods, but it is still challenging to obtain a robust network in the large-scale setting and apply it to high-dimensional data.

Randomized Smoothing. To further scale up certified defenses on larger networks encountered in practice, randomized smoothing methods have been proposed, which provide certified robust guarantee in theory and Monte-Carlo estimation-based certification and prediction in practice. Lecuyer et al. (PixelDP) first applied differential privacy to prove robustness guarantees for randomized smoothing classifier in the case of ℓ_0 - norm attack [23]. Cohen et al. [6] used Neyman-Pearson theorem to prove tight certified robustness bound through the smoothed classifier in ℓ_2 -norm certified defense. More recently, Dvijotham et al. [9] extended randomized smoothing to handle arbitrary ℓ_p settings and prove the robustness using f-divergence. However, all these certified robustness analyses have relied on the continuous probability distribution, which incurs compatibility issues with the integer-arithmetic-only platforms. In addition, Monte-Carlo estimation requires a large number of repeated inferences during certification, which leads to huge computation time, especially for large neural networks. Finally, we note that there are existing works that propose randomized smoothing variants for discrete input domain [24, 3, 45], which are somewhat closer to our problem. However, they merely consider the discrete nature of the input space while still pay little attention to neither applicability to the quantized neural networks nor to the integer-arithmetic-only restriction of many applications with quantized neural networks.

Model Compression: Quantization. Quantization aims to compress a network and save computation by changing the weights and activation of a neural network from 32-bit floating-point representation into lower bit-depth representations. The weights can be quantized to 16-bit [16, 7], 8-bit [20, 40], 4-bit [39, 49] or even with 1-bit (also known as binary quantization) [19, 48]. Existing techniques can be roughly divided into two categories: quantization during training and quantization for inference. Early quantization during training is applied to reduce the network size and make the training process more computational efficient [5]. Guo et al. [15] applied vector quantization to compress the weighing matrix. On the contrary, quantization for inference mainly focuses on accelerating inference speed. For example, Han et al. [17] proposed a three-stage pipeline: pruning, quantization, and fine-tuning to determine the bit widths for convolution layers [17]. A more recent quantization scheme has been proposed in [20], which inserts “fake” quantization operations during training using floating-point operations and builds quantized parameters during the inference stage. It simulates and decreases quantization error while perform-

ing efficient inference on integer-arithmetic-only hardware. We will further extend this quantization scheme and incorporate it into our certified robust networks.

Quantization with Robust Training. Several works try to address the challenge of training deep neural networks on security-critical and resource-limited applications. Gui et al. [14] construct constrained optimization formulation to integrate various existing compression techniques (e.g., pruning, factorization, quantization) into adversarially robust models. However, their approach is only applied to defend specific adversarial attacks and does not provide certified robustness guarantees. Another related line of research aims to connect model compression with the robustness goal through pruning [35, 46, 34]. However, most of the works focus on empirical adversarial defenses, and none of them support deployment with integer-arithmetic-only platforms.

3. Proposed Method

In this section, we present a new certified robustness defense against adversarial perturbation attack under the integer-arithmetic-only context.

3.1. Preliminaries

First, we formally define the adversarial perturbation attack and certified robustness under the integer-arithmetic-only context, as follows.

Definition 3.1. (ℓ_2 -Norm Bounded Integer Adversarial Perturbation Attack) For any input $\mathbf{x} \in \mathbb{H}^d$, where \mathbb{H} is a discrete additive subgroup, the ℓ_2 -norm bounded integer adversarial perturbation attack with magnitude L perturbs \mathbf{x} to $\mathbf{x} + \delta$ in order to alter the classification result, where $\delta \in \mathbb{H}^d$ and $\|\delta\|_2^2 \leq L^2$. Denote all possible $\mathbf{x} + \delta$ within L distance by $\mathbb{B}_{\mathbb{H}}(\mathbf{x}, L)$.

Definition 3.2. (Certified Robustness to ℓ_2 -Norm Bounded Integer Adversarial Example Attack with magnitude L) Denote a multiclass classification model by $f(\mathbf{x}) : \mathcal{X} \mapsto c \in \mathcal{C}$, where c is a label in the labels set $\mathcal{C} = \{1, \dots, C\}$. In general, $f(\mathbf{x})$ outputs a vector of scores $f^{\mathbf{y}}(\mathbf{x}) = (f^{y_1}, \dots, f^{y_C}) \in \mathcal{Y}$, where $\mathcal{Y} = \{\mathbf{y} : \sum_{i=1}^C f^{y_i} = 1, f^{y_i} \in [0, 1]\}$, and $c = \operatorname{argmax}_{i \in \mathcal{C}} f^{y_i}$. A predictive model $f(\mathbf{x})$ is robust to ℓ_2 -norm integer adversarial perturbation attack with magnitude L on input \mathbf{x} , if for all $\mathbf{x}' \in \mathbb{B}_{\mathbb{H}}(\mathbf{x}, L)$, it has $f(\mathbf{x}) = f(\mathbf{x}')$, which is equivalent to

$$f^{y_c}(\mathbf{x}') > \max_{i \in \mathcal{C}: i \neq c} f^{y_i}(\mathbf{x}'). \quad (1)$$

The following lemma is utilized in the proof of the certified robustness guarantee.

Lemma 3.1. (Neyman-Pearson Lemma [30]) Let X, Y be random variables in \mathbb{H}^d with probability mass function

p_X, p_Y . Let $h : \mathbb{H}^d \mapsto \{0, 1\}$ be a random or deterministic function. It has:

1) If $\mathcal{L}(\mathbf{z}) = \{\mathbf{z} \in \mathbb{H}^d : \frac{p_Y(\mathbf{z})}{p_X(\mathbf{z})} \leq \alpha\}$ for some $\alpha > 0$ and $\mathbb{P}(h(X) = 1) \geq \mathbb{P}(X \in \mathcal{L}(\mathbf{z}))$, then $\mathbb{P}(h(Y) = 1) \geq \mathbb{P}(Y \in \mathcal{L}(\mathbf{z}))$;

2) If $\mathcal{L}(\mathbf{z}) = \{\mathbf{z} \in \mathbb{H}^d : \frac{p_Y(\mathbf{z})}{p_X(\mathbf{z})} \geq \alpha\}$ for some $\alpha > 0$ and $\mathbb{P}(h(X) = 1) \leq \mathbb{P}(X \in \mathcal{L}(\mathbf{z}))$, then $\mathbb{P}(h(Y) = 1) \leq \mathbb{P}(Y \in \mathcal{L}(\mathbf{z}))$.

3.2. Integer Randomized Smoothing for Quantized Neural Networks

We propose the integer randomized smoothing (abbreviated as IntRS) for quantized neural networks to ensure certified robustness against L - ℓ_2 -norm bounded integer adversarial perturbation, which involves only integer arithmetic operations. For this purpose, we utilize the following discrete Gaussian random noise.

Definition 3.3. (Discrete Gaussian Distribution) The discrete Gaussian distribution $\mathcal{N}_{\mathbb{H}}(\mu, \sigma^2)$ is a probability distribution supported on \mathbb{H} with location μ and scale σ . Its probability mass function is defined as follows,

$$\mathbb{P}_{X \sim \mathcal{N}_{\mathbb{H}}(\mu, \sigma^2)}[X = x] = \frac{e^{-(x-\mu)^2/2\sigma^2}}{\sum_{h \in \mathbb{H}} e^{-(h-\mu)^2/2\sigma^2}}. \quad (2)$$

The discrete Gaussian noise-based integer randomized smoothing (IntRS) is defined as follows.

Definition 3.4. (Integer Randomized Smoothing with Discrete Gaussian Noise) For an arbitrary base quantized classifier defined on \mathbb{H}^d , for any input $\mathbf{x} \in \mathbb{H}^d$, the smoothed classifier $g(\mathbf{x})$ is defined as

$$g(\mathbf{x}) = \operatorname{argmax}_{c \in \mathcal{C}} \mathbb{P}(f(\mathbf{x} + \boldsymbol{\delta}) = c), \quad \boldsymbol{\delta} \sim \mathcal{N}_{\mathbb{H}}(0, \sigma^2 \mathbf{I}_d). \quad (3)$$

Remark 1. One may wonder why we cannot re-utilize the existing randomized smoothing that relied on continuous Gaussian noise (e.g., [6]) by rounding the noise for quantized neural networks. There are two reasons: 1) Sampling and injecting continuous Gaussian noise still requires float-number operations, which raises the compatibility issue; 2) Although 1) can be addressed via pre-storing rounded continuous Gaussian noises, the rounding error introduces randomness that is difficult to precisely calibrate, which contradicts our goal to guarantee certified robustness rigorously. To form a theoretical proof, we utilize the discrete Gaussian noise, whose distribution is precisely described by eq.(3).

Although the IntRS mechanism parallels the randomized smoothing in the full precision context, it is nontrivial to establish its certified robustness guarantee. A simpler yet indirect way is to first regard IntRS as a differential privacy mechanism [4] and then establish the certified robustness in

a similar manner as PixelDP [23]. We take another direct way, which provides a tighter certified robustness guarantee. The key to our certified robustness guarantee of the IntRS mechanism is by extending the Neyman-Pearson lemma to the discrete Gaussian distribution, as summarized by the following Proposition.

Proposition 3.1. (Neyman-Pearson for Discrete Gaussian with Different Means) For $\mathbf{X} \sim \mathcal{N}_{\mathbb{H}}(\mathbf{x}, \sigma^2 \mathbf{I}_d)$, $\mathbf{Y} \sim \mathcal{N}_{\mathbb{H}}(\mathbf{x} + \boldsymbol{\delta}, \sigma^2 \mathbf{I}_d)$, and $h : \mathbb{H}^d \mapsto \{0, 1\}$, it has:

1) If $\mathcal{L}(\mathbf{z}) = \{\mathbf{z} \in \mathbb{H}^d : \langle \mathbf{z}, \boldsymbol{\delta} \rangle \leq \sigma^2 \ln \alpha + \frac{1}{2}(\|\boldsymbol{\delta}\|_2^2 + 2\langle \mathbf{x}, \boldsymbol{\delta} \rangle)\}$ for some α and $\mathbb{P}(h(\mathbf{x}) = 1) \geq \mathbb{P}(\mathbf{x} \in \mathcal{L}(\mathbf{z}))$, then $\mathbb{P}(h(\mathbf{y}) = 1) \geq \mathbb{P}(\mathbf{y} \in \mathcal{L}(\mathbf{z}))$;

2) If $\mathcal{L}(\mathbf{z}) = \{\mathbf{z} \in \mathbb{H}^d : \langle \mathbf{z}, \boldsymbol{\delta} \rangle \geq \sigma^2 \ln \alpha + \frac{1}{2}(\|\boldsymbol{\delta}\|_2^2 + 2\langle \mathbf{x}, \boldsymbol{\delta} \rangle)\}$ for some α and $\mathbb{P}(h(\mathbf{x}) = 1) \leq \mathbb{P}(\mathbf{x} \in \mathcal{L}(\mathbf{z}))$, then $\mathbb{P}(h(\mathbf{y}) = 1) \leq \mathbb{P}(\mathbf{y} \in \mathcal{L}(\mathbf{z}))$.

Proof. Proof can be found in Appendix H.1. \square

The following Theorem provides the certified robustness to the ℓ_2 -norm bounded integer adversarial perturbation achieved by IntRS.

Theorem 3.1. (Certified Robustness via Integer Randomized Smoothing) Let $f : \mathbb{R}^d \mapsto \mathcal{Y}$ be the base classifier. Let $\mathbf{n} \sim \mathcal{N}_{\mathbb{H}}(0, \sigma^2 \mathbf{I}_d)$. Denote the randomized smoothed classifier by $g(\mathbf{x}) = \operatorname{argmax}_{c \in \mathcal{Y}} \mathbb{P}[f(\mathbf{x} + \mathbf{n}) = c]$. If there exists $c_A \in \mathcal{Y}$ such that the following relation stands

$$\mathbb{P}[f(\mathbf{x} + \mathbf{n}) = c_A] \geq p_{c_A}^{lb} \geq p_{c_B}^{ub} \geq \max_{c_B: c \neq c_A} \mathbb{P}[f(\mathbf{x} + \mathbf{n}) = c], \quad (4)$$

then the randomized smoothed classifier $g(\mathbf{x})$ is certified robust to perturbation $\|\boldsymbol{\delta}\|_2^2 \leq R^2$, i.e., $g(\mathbf{x} + \boldsymbol{\delta}) = c_A$, where R is the certified radius.

Proof. Let $\mathbf{X} = \mathbf{x} + \mathbf{n} \sim \mathcal{N}_{\mathbb{H}}(\mathbf{x}, \sigma^2 \mathbb{I})$, $\mathbf{Y} = \mathbf{x} + \mathbf{n} + \boldsymbol{\delta} \sim \mathcal{N}_{\mathbb{H}}(\mathbf{x} + \boldsymbol{\delta}, \sigma^2 \mathbf{I}_d)$. To show that $g(\mathbf{x} + \boldsymbol{\delta}) = c_A$, we first prove

$$\begin{aligned} \mathbb{P}[f(\mathbf{x} + \mathbf{n} + \boldsymbol{\delta}) = c_A] &> \mathbb{P}[f(\mathbf{x} + \mathbf{n} + \boldsymbol{\delta}) = c_B] \\ \iff \mathbb{P}[f(\mathbf{Y}) = c_A] &> \mathbb{P}[f(\mathbf{Y}) = c_B], \end{aligned} \quad (5)$$

given that $\mathbb{P}[f(\mathbf{X}) = c_A] \geq p_{c_A}^{lb} \geq p_{c_B}^{ub} \geq \mathbb{P}[f(\mathbf{X}) = c_B]$. Notice that that $p_{c_A}^{lb} = \mathbb{P}[X \in \mathcal{S}_A]$, where $\mathcal{S}_A = \{\mathbf{z} : \langle \mathbf{z} - \mathbf{x}, \boldsymbol{\delta} \rangle \leq \sigma \|\boldsymbol{\delta}\|_2 \Phi_{\mathcal{N}_{\mathbb{H}}}^{-1}(p_{c_A}^{lb})\}$ and $p_{c_B}^{ub} = \mathbb{P}[X \in \mathcal{S}_B]$, where $\mathcal{S}_B = \{\mathbf{z} : \langle \mathbf{z} - \mathbf{x}, \boldsymbol{\delta} \rangle \geq \sigma \|\boldsymbol{\delta}\|_2 \Phi_{\mathcal{N}_{\mathbb{H}}}^{-1}(1 - p_{c_B}^{ub})\}$. Then, by Lemma 3.1, we have

$$\mathbb{P}[f(\mathbf{Y}) = c_A] \geq \mathbb{P}[\mathbf{Y} \in \mathcal{S}_A] = p_{c_A}^{lb} \quad (6)$$

$$\mathbb{P}[f(\mathbf{Y}) = c_B] \leq \mathbb{P}[\mathbf{Y} \in \mathcal{S}_B] = p_{c_B}^{ub}, \quad (7)$$

if $\boldsymbol{\delta}$ satisfies the constraint:

$$\|\boldsymbol{\delta}\|_2^2 < \left(\frac{\sigma}{2}(\Phi_{\mathcal{N}_{\mathbb{H}}}^{-1}(p_{c_A}^{lb}) - \Phi_{\mathcal{N}_{\mathbb{H}}}^{-1}(p_{c_B}^{ub}))\right)^2. \quad (8)$$

Thus, we have proved $\mathbb{P}[f(\mathbf{Y}) = c_A] \geq p_{c_A}^{lb} \geq p_{c_B}^{ub} \geq \mathbb{P}[f(\mathbf{Y}) = c_B]$. \square

Practical IntRS. As with all randomized smoothing methods, we resort to Monte-Carlo estimation to approximate the randomized smoothing function $g(\mathbf{x})$, because it is challenging to obtain the precise $g(\mathbf{x})$. Essentially, for certification on example \mathbf{x} , we repeatedly evaluate $f(\mathbf{x} + \mathbf{n})$, $\mathbf{n} \sim \mathcal{N}_{\mathbb{H}}(0, \sigma^2 \mathbf{I}_d)$, and obtain robust prediction or certification result based on the aggregated evaluations. Algorithm 1 summarizes the robust certification, which follows [6].

Algorithm 1 Monte-Carlo estimation and aggregated evaluation for certified robust certification

Input: Base function $f(\cdot)$, inference sample \mathbf{x} , Gaussian noise std σ , repeated number N_1, N_2 , and confidence level α .

Certification:

- 1: Repeat N_1 inferences on $f(\mathbf{x} + \mathbf{n})$, where $\mathbf{n} \sim \mathcal{N}_{\mathbb{H}}(0, \sigma^2 \mathbf{I}_d)$.
 - 2: Collect prediction results: (n_1, \hat{c}_A) : highest prediction count and its label;
 - 3: Repeat N_2 inferences on $f(\mathbf{x} + \mathbf{n})$, where $\mathbf{n} \sim \mathcal{N}_{\mathbb{H}}(0, \sigma^2 \mathbf{I}_d)$.
 - 4: Collect prediction results: n_2 : count of \hat{c}_A ;
 - 5: Let p_A be the one-sided $(1 - \alpha)$ lower confidence interval for $\text{Binomial}(N_2, n_2)$;
 - 6: **if** $p_A > \frac{1}{2}$ **then**
 - 7: **Return** \hat{c}_A and radius $\sigma \Phi^{-1}(p_A)$;
 - 8: **else**
 - 9: ABSTAIN
 - 10: **end if**
-

The following Theorem and its proof summarize that the practical prediction and certification of IntRS can be carried out via integer-arithmetic-only operations.

Theorem 3.2. *The prediction and certification procedure of the practical IntRS can be executed with integer-arithmetic-only operations.*

Proof. The overall prediction and certification procedure consist of the following integer-arithmetic only steps:

1. Repeat the inference N, N_1, N_2 times to estimate \hat{c}_A, \hat{c}_B and counts, which evaluates $\mathbf{x} + \delta$ with quantized neural network.
2. Compute $\Phi_{\mathcal{N}_{\mathbb{H}}}^{-1}(p_A)$. Since the distribution is discrete, the cumulative density function has exactly $|\mathbb{H}|$ elements, which means the cumulative density function can be precomputed and its inverse can be obtained by looking up during inference.
3. Return the squared ℓ_2 norm. Since we do not take least square norm, the squared ℓ_2 norm is closed within \mathbb{H} .

Thus, the robust prediction and certification of the IntRS mechanism can be implemented with integer-arithmetic-only operations. \square

3.3. Practical Training: Discrete Noise and Quantization Aware Training

In this section, we present how to incorporate Quantization Aware Training (QAT) with our randomized mechanism.

Quantization replaces floating-point number representations with low precision fixed-point representations. Since the quantization scheme replaces floating-point weights/activations by N ($N=256$ for 8-bits) fixed-point numbers, if we directly quantize the pre-trained network for inference, rounding errors accumulate, leading to a significant drop in the classification performance. A solution to rectify these drifting errors is to quantize the network during training. However, representing parameters by integers becomes a challenge in backward pass since direct optimization over a discrete domain generates null gradients. To address this issue, QAT keeps carrying out computation using 32-bits floating-point arithmetic while simulating quantization error. In particular, during the forward pass, weight is quantized into integers and then convoluted with the input, which is represented with floating numbers. On the other hand, back-propagation is the same as that of the full precision model.

In order for the base classifier f to classify the labeled examples correctly and robustly, f needs to classify them with their true labels consistently. Thus, we train f with noisy images, which is also adopted by almost all certified robustness approaches. To combine with the QAT scheme, in practice, we train the model with discrete Gaussian data augmentation at variance σ , which is the same as used for prediction and certification. For each training data point, we randomly draw noise from discrete Gaussian distribution. We follow the sampling strategy in [4]. Note that although the algorithm for sampling from discrete Gaussian distribution runs $O(1)$ on average, since we need to add noise onto every channel of every pixel in the image, sampling noise for a batch of 32×32 images takes more than 3 mins. To accelerate training speed, we first sample noises for a batch and store these noises in advance. During training, we randomly shuffle and add the noise to the input.

4. Experiment

4.1. Experiment Setup

Datasets. We evaluate our framework on two standard image classification datasets: CIFAR-10 [31] and Caltech-101 [11]. **CIFAR-10** consists of 50,000 training images and 10,000 test images, where each image is of 32×32 resolution.

Caltech-101 contains 9,144 images of size 300×200 pixels in 102 categories (one of which is background). We use 101 categories for classification (without the background category). We randomly split 80% for training and the remaining images for testing.

Evaluation Metrics. For evaluation metrics, we use certi-

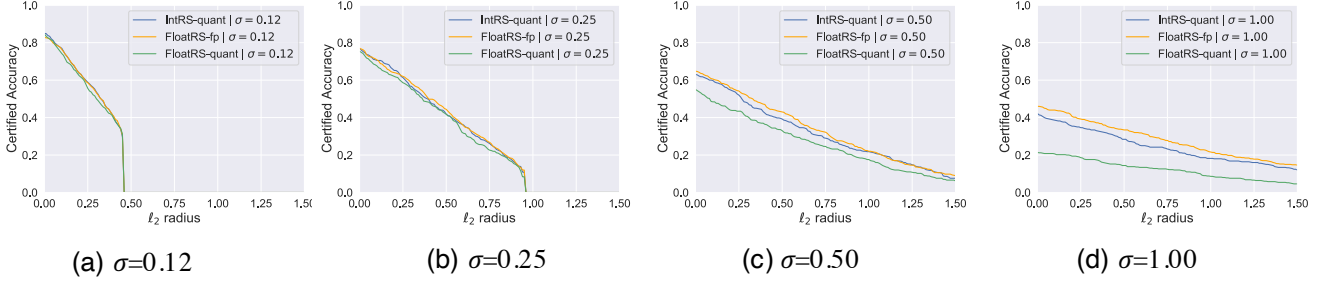


Figure 3. Certified accuracy of IntRS-quant trained CIFAR-10 classifiers vs Float-fp and FloatRS-quant. (Blue,Orange,Green) per σ .

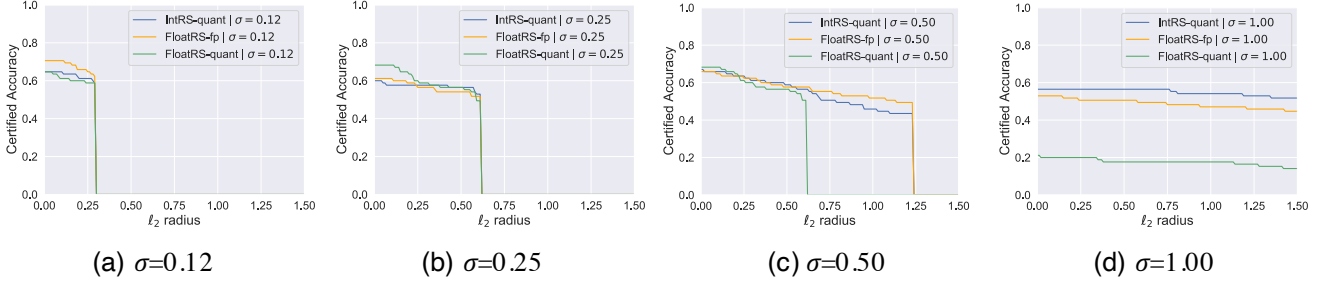


Figure 4. Certified accuracy of IntRS-quant trained Caltech-101 classifiers vs Float-fp and FloatRS-quant. (Blue,Orange,Green) per σ .

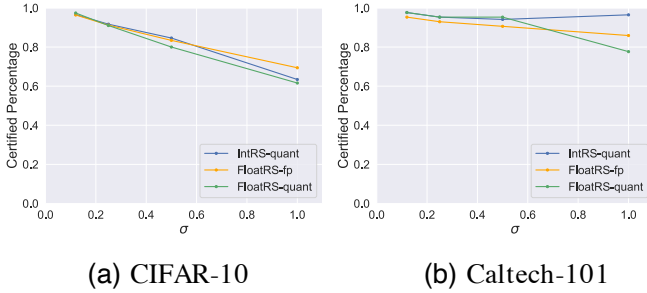


Figure 5. Certified Percentage of IntRS-quant trained Caltech-101 classifiers vs Float-fp and FloatRS-quant per σ over CIFAR-10 (Left) and Caltech-101 (Right).

ified percentage and certified accuracy[6], as follows:

1) Certified Percentage (CP) is defined as the percentage of clean inputs that satisfy the certification criteria in Theorem 3.1 under the corresponding ℓ_2 -norm ball, which can be formulated as

$$\frac{\sum_{i=1}^L |\text{Certification}(N_1, \mathbf{x}_i, \sigma)|}{L}, \quad (9)$$

where $|\text{Certification}(N_1, \mathbf{x}_i, \sigma)| = 1$ if it does not abstain and 0, otherwise. L represents the size of dataset.

2) Certified Accuracy (CA) at radius r is defined as the fraction of the test set in which the randomized smoothing function g makes a correct prediction without abstaining

within an ℓ_2 ball of radius r , which can be formulated as

$$\frac{\sum_{i=1}^L [|\text{Certification}(N_1, \mathbf{x}_i, \sigma)| \& \hat{c}_i == c_i \& R_i \geq r]}{\sum_{i=1}^L |\text{Certification}(N_1, \mathbf{x}_i, \sigma)|}, \quad (10)$$

where \hat{c}_i, R_i are the returned results for certification for \mathbf{x}_i . c_i is x_i 's ground-truth label.

Comparison Methods. We compare our proposed IntRS method with

1) FloatRS-fp: Cohen et al. [6] as it was the state-of-the-art provable defence to ℓ_2 -norm bounded adversarial perturbation attacks. We denote it as FloatRS-fp (short for floating-point randomized smoothing for full precision neural networks).

2) FloatRS-quant: We implement a vanilla randomized smoothing approach for quantized neural networks, which first adds continuous Gaussian noises to the images and then quantizes the perturbed images. We denote it as FloatRS-quant (short for floating-point randomized smoothing for quantized neural networks). We stress that FloatRS-quant is incompatible with a real integer-arithmetic-only device since the continuous noise injection step requires floating-point operations.

To conduct a fair comparison, we run experiments with the same training settings when possible (i.e., use the same quantization-aware and noise-aware strategy for the corresponding neural network) and report the comparison results in metrics of CP and CA.

Implementation Details We implement our algorithm in

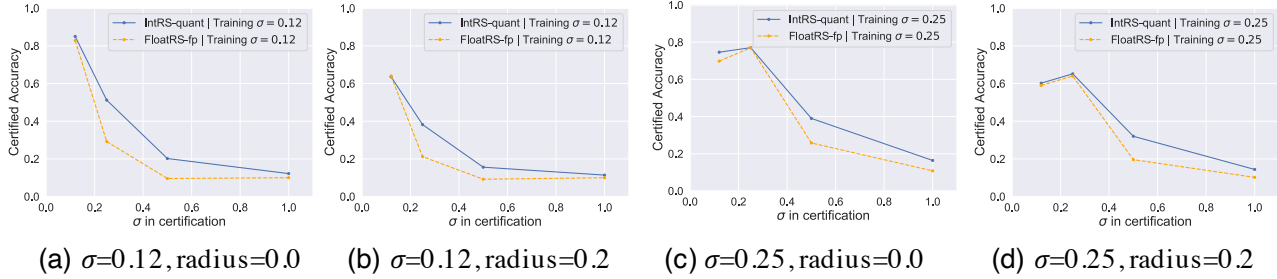


Figure 6. Comparison between IntRS-quant and FloatRS-fp on CIFAR-10. Each model is trained with a fixed σ while certifying using different σ . We perform experiments under train $\sigma = 0.12, 0.25$ and report certified accuracy for radius = 0.0, 0.2.

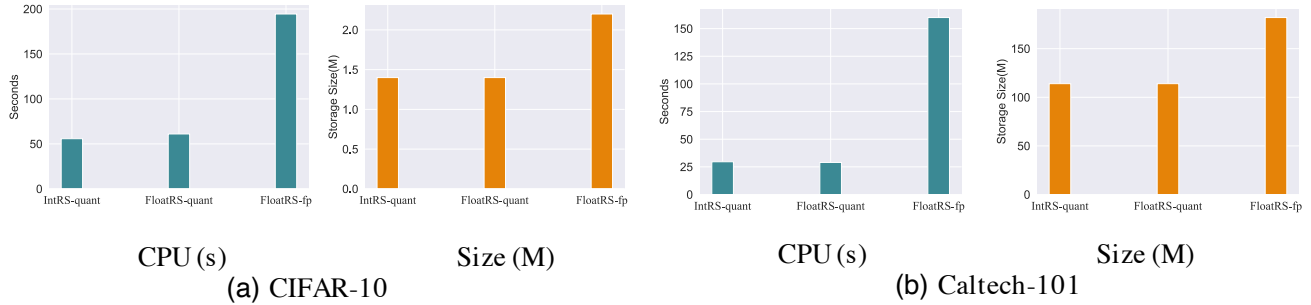


Figure 7. Computational efficiency between IntRS-quant and FloatRS-fp on CIFAR-10 and Caltech-101.

Pytorch [33]. We consider ResNet-20 [18] as the base model architecture for CIFAR-10 and ResNet-50 for Caltech-101. In addition, we adopt a model pretrained from ImageNet for Caltech-101 as the images from these two datasets are similar. On each dataset, we train multiple smoothed classifiers with different σ 's. We compress the model from FP32 representation to int8 representation for the quantized neural networks. We set batch size to 256 for CIFAR-10 and 64 for Caltech-101. The learning rate is set to 0.1 using the step scheduler and SGD optimizer. In all experiments, unless otherwise stated, we set $\alpha = 0.001$ for **Prediction** in Algorithm 1, which corresponds to 99.9% confidence interval, i.e. there is at most a 0.1% chance that **Certification** falsely certifies a non-robust input. In **Certification** algorithm, we use $N_1 = 100$ and $N_2 = 100,000$ in CIFAR-10. We set $N_2 = 1000$ for Caltech-101 during efficiency comparison. We certify the entire CIFAR-10 testset and 100 images for Caltech-101.

4.2. Certified Accuracy and Certified Percentage on Adversarially Perturbed Examples

Figure 3 and Figure 4 plot CA obtained by smoothing with different σ over a range of ℓ_2 at radius r . The blue, orange, green lines are the CA on adversarial examples for IntRS-quant, Float-fp, and FloatRS-quant, respectively. Note that CA drops to zero beyond a certain maximum point for each σ as there is an upper bound to the radius we can

certify. This maximum radius is achieved when p_{α} reaches its maximum and all N_2 samples are classified as the same class), where p_{α} and N_2 are defined in Algorithm 1. Figure 5 reports CP for each dataset. It is observed that IntRS-quant achieves comparable CA and CP on both datasets with mild degradation in CA for a few cases compared with the Float-fp model. This is because quantization reduces the bits per weight and leads to errors by these approximations during computations. Moreover, compared with FloatRS-quant, our algorithm can achieve better performance, especially when σ is large, which means IntRS is not only capable of being deployed on a wider selection of integer-arithmetics-only devices but also better in prediction accuracy.

Trade-off between Accuracy and Robustness. As mentioned by Cohen et al. [6] and [43], we also observe from Figure 3 and Figure 4 that hyperparameter σ controls a trade-off between robustness and accuracy. That is, when σ is low, small radii can be certified with high accuracy, but large radii cannot be certified at all; when σ is high, larger radii can be certified, but smaller radii are certified at lower accuracy. This phenomenon can also be verified in Figure 5 with a trade-off between CP and σ where we observe higher σ leads to lower CP.

Effects of Mismatching Noise Magnitudes for Data Augmentation in Training and Smoothing in Inference. According to [6], they train the base classifier with Gaussian

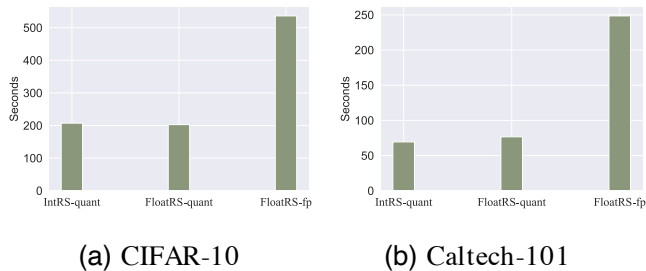


Figure 8. Inference speed between IntRS-quant and FloatRS-fp on mobile device.

noise data augmentation at variance σ^2 in their experiments. If the base classifier f does not see the same noise during training, it will not necessarily learn to classify x under Gaussian noise with its ground truth label. Specifically, they report CA using the same σ as the standard deviation of Gaussian noise data augmentation for training and robust inference. However, in reality, the attacker may use different σ 's during inference time. In Figure 6, we present CA of models trained with fixed σ under different $\sigma = \{0.12, 0.25, 0.5, 1.0\}$ used for certification. Figure 6 (a) and (b) show the result for training with data augmentation using $\sigma = 0.12$ on CIFAR-10 at radius $r = 0.0$ and 2.0 , respectively. Figure 6 (c) and (d) and show the result for training with $\sigma = 0.25$. As shown in the figure, certifying using σ different from that applied in training can result in lower CA. The classifier achieves the best CP when the training σ equals the certification σ , and the classifier performs worse when it is certifying with smaller σ , implying that less randomness is introduced in the inference (in Figure 6 (c) and (d)). We can also observe that our IntRS approach outperforms Float-fp under all different certification σ 's. This indicates that IntRS is less sensitive to the mismatch level compared with the full precision model and thus provides more flexibility during certified robust inference.

4.3. Efficiency Evaluation

We compare the efficiency performance on two different types of devices: general-purpose CPU and mobile device. For both types, we report the average robust inference time per image on CIFAR-10 and Caltech-101 datasets. The counts of Monte-Carlo repetitions are set to 100,000 and 10,000 respectively for CIFAR-10 and Caltech-101.

Efficiency Results on General-purpose CPU. We test running time on an Intel-i7 CPU-powered desktop computer as an instance of the general-purpose device, which in general has more computational capabilities and power consumption than most edge devices. The results on CIFAR-10 and Caltech-101 are in Figure 7. Since Caltech-101 has fewer Monte-Carlo repetitions, certifying an image on it is faster than that on CIFAR-10. The results show that IntRS-quant

for the quantized model requires only 30% of the original Float-fp for the full precision model in inference time (55.8 s v.s. 194.4 s on CIFAR-10 and 29.5 s vs. 159.9s on Caltech-101 on average). The memory cost can be roughly compared by the storage size of the saved model, which shows ours has approximately 40% less storage than that of a full precision model.

Efficiency Results on Mobile Device. Next, we test running time on an iPhone 11 as an instance of the mobile device. The results are reported in Figure 8. On the mobile device, it shows clear efficiency improvement of our IntRS for the quantized model over the FP32 model, where the latter costs over 8 mins to robustly predict a single CIFAR-10 image, which makes it impractical to maintain user engagement. Finally, we note that many edge devices are subject to stricter resource and power restrictions, on which the efficiency improvement of randomized smoothing-type techniques made possible by IntRS can be more significant.

5. Conclusion

In this paper, we proposed an integer-arithmetic-only randomized smoothing mechanism called IntRS, which has made the certified robustness inference for quantized neural networks ever possible for the first time. We rigorously analyzed its certified robustness property based on the discrete Neyman-Pearson lemma when specified to the discrete Gaussian noise. In practice, we incorporated quantization- and discrete data augmentation- aware training, as well as Monte-Carlo-based practical prediction and certification. We evaluated its effectiveness with modern CNN architectures (ResNet-20 and ResNet-50) on two distinct datasets: CIFAR-10 and Caltech-101. We demonstrated through extensive experiments that IntRS can obtain comparable certified accuracy and certified percentage. More importantly, IntRS makes inference efficient on edge devices. Compared to inferences using models with floating point operations, IntRS requires 40% times less storage size and gains $4\times$ to $5\times$ acceleration for inference on general-purpose CPUs and mobile devices.

6. Acknowledgement

This research has been funded in part by NSF CNS-1952192, NIH CTSA Award UL1TR002378, and Cisco Research Award #2738379, NSF grants IIS-1910950 and CNS-2027794 and unrestricted cash gifts from Google and Microsoft. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the sponsors.

References

- [1] Chirag Agarwal, Anh Nguyen, and Dan Schonfeld. Improving robustness to adversarial examples by encouraging discriminative features. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 3801–3505. IEEE, 2019.
- [2] Anish Athalye and Nicholas Carlini. On the robustness of the cvpr 2018 white-box adversarial example defenses. *arXiv preprint arXiv:1804.03286*, 2018.
- [3] Aleksandar Bojchevski, Johannes Klicpera, and Stephan Günnemann. Efficient robustness certificates for discrete data: Sparsity-aware randomized smoothing for graphs, images and more. In *International Conference on Machine Learning*, pages 1003–1013. PMLR, 2020.
- [4] Clément Canonne, Gautam Kamath, and Thomas Steinke. The discrete gaussian for differential privacy. *arXiv preprint arXiv:2004.00010*, 2020.
- [5] Tejalal Choudhary, Vipul Mishra, Anurag Goswami, and Jagannathan Sarangapani. A comprehensive survey on model compression and acceleration. *Artificial Intelligence Review*, pages 1–43, 2020.
- [6] Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In *International Conference on Machine Learning*, pages 1310–1320. PMLR, 2019.
- [7] Dipankar Das, Naveen Mellempudi, Dheevatsa Mudigere, Dhiraj Kalamkar, Sasikanth Avancha, Kunal Banerjee, Srinivas Sridharan, Karthik Vaidyanathan, Bharat Kaul, Evangelos Georganas, et al. Mixed precision training of convolutional neural networks using integer operations. *arXiv preprint arXiv:1802.00930*, 2018.
- [8] Zhen Dong, Zhewei Yao, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. Hawq: Hessian aware quantization of neural networks with mixed-precision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 293–302, 2019.
- [9] Krishnamurthy Dj Dvijotham, Jamie Hayes, Borja Balle, Zico Kolter, Chongli Qin, András György, Kai Xiao, Sven Gowal, and Pushmeet Kohli. A framework for robustness certification of smoothed classifiers using f-divergences. In *ICLR*, 2020.
- [10] Julian Faraone, Martin Kumm, Martin Hardieck, Peter Zipf, Xueyuan Liu, David Boland, and Philip HW Leong. Addnet: Deep neural networks using fpga-optimized multipliers. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 28(1):115–128, 2019.
- [11] Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *2004 conference on computer vision and pattern recognition workshop*, pages 178–178. IEEE, 2004.
- [12] Matteo Fischetti and Jason Jo. Deep neural networks as 0-1 mixed integer linear programs: A feasibility study. *arXiv preprint arXiv:1712.06174*, 2017.
- [13] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [14] Shupeng Gui, Haotao Wang, Chen Yu, Haichuan Yang, Zhangyang Wang, and Ji Liu. Model compression with adversarial robustness: A unified optimization framework. *arXiv preprint arXiv:1902.03538*, 2019.
- [15] Yiwen Guo, Anbang Yao, and Yurong Chen. Dynamic network surgery for efficient dnns. In *Advances in neural information processing systems*, 2016.
- [16] Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. Deep learning with limited numerical precision. In *International conference on machine learning*, pages 1737–1746. PMLR, 2015.
- [17] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [19] 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*, pages 4114–4122. Citeseer, 2016.
- [20] 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*, pages 2704–2713, 2018.
- [21] Guy Katz, Clark Barrett, David L Dill, Kyle Julian, and Mykel J Kochenderfer. Reluplex: An efficient smt solver for verifying deep neural networks. In *International Conference on Computer Aided Verification*, pages 97–117. Springer, 2017.
- [22] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*, 2016.
- [23] Mathias Lecuyer, Vaggelis Atlidakis, Roxana Geambasu, Daniel Hsu, and Suman Jana. Certified robustness to adversarial examples with differential privacy. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 656–672. IEEE, 2019.
- [24] Guang-He Lee, Yang Yuan, Shiyu Chang, and Tommi S. Jaakkola. Tight certificates of adversarial robustness for randomly smoothed classifiers. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 4911–4922, 2019.
- [25] Bai Li, Changyou Chen, Wenlin Wang, and Lawrence Carin. Certified adversarial robustness with additive noise. *arXiv preprint arXiv:1809.03113*, 2018.
- [26] Yuezun Li, Daniel Tian, Ming-Ching Chang, Xiao Bian, and Siwei Lyu. Robust adversarial perturbation on deep proposal-based models. *arXiv preprint arXiv:1809.05962*, 2018.
- [27] Qun Liu and Supratik Mukhopadhyay. Unsupervised learning using pretrained cnn and associative memory bank. In *2018*

- International Joint Conference on Neural Networks (IJCNN)*, pages 01–08. IEEE, 2018.
- [28] Bo Luo, Yannan Liu, Lingxiao Wei, and Qiang Xu. Towards imperceptible and robust adversarial example attacks against neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [29] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- [30] Jerzy Neyman and Egon Sharpe Pearson. Ix. on the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231(694-706):289–337, 1933.
- [31] University of Toronto. Learning multiple layers of features from tiny images. 2012.
- [32] Nicolas Papernot, Patrick McDaniel, Somesh Jha, Matt Fredrikson, Z Berkay Celik, and Ananthram Swami. The limitations of deep learning in adversarial settings. In *2016 IEEE European symposium on security and privacy (EuroS&P)*, pages 372–387. IEEE, 2016.
- [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. In *NIPS*, pages 8026–8037, 2019.
- [34] Vikash Sehwal, Shiqi Wang, Prateek Mittal, and Suman Jana. Towards compact and robust deep neural networks. *arXiv preprint arXiv:1906.06110*, 2019.
- [35] Vikash Sehwal, Shiqi Wang, Prateek Mittal, and Suman Jana. Hydra: Pruning adversarially robust neural networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 7, 2020.
- [36] Gagandeep Singh, Timon Gehr, Markus Püschel, and Martin T Vechev. Boosting robustness certification of neural networks. In *ICLR (Poster)*, 2019.
- [37] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [38] Vincent Tjeng, Kai Xiao, and Russ Tedrake. Evaluating robustness of neural networks with mixed integer programming. *arXiv preprint arXiv:1711.07356*, 2017.
- [39] Peisong Wang, Qiang Chen, Xiangyu He, and Jian Cheng. Towards accurate post-training network quantization via bit-split and stitching. In *International Conference on Machine Learning*, pages 9847–9856. PMLR, 2020.
- [40] Tianzhe Wang, Kuan Wang, Han Cai, Ji Lin, Zhijian Liu, Hanrui Wang, Yujun Lin, and Song Han. Apq: Joint search for network architecture, pruning and quantization policy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2078–2087, 2020.
- [41] Wenjie Wang, Pengfei Tang, Jian Lou, and Li Xiong. Certified robustness to word substitution attack with differential privacy. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1102–1112, 2021.
- [42] Lily Weng, Huan Zhang, Hongge Chen, Zhao Song, Chou-Jui Hsieh, Luca Daniel, Duane Boning, and Inderjit Dhillon. Towards fast computation of certified robustness for relu networks. In *International Conference on Machine Learning*, pages 5276–5285. PMLR, 2018.
- [43] Eric Wong and Zico Kolter. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *International Conference on Machine Learning*, pages 5286–5295. PMLR, 2018.
- [44] Bichen Wu, Yanghan Wang, Peizhao Zhang, Yuandong Tian, Peter Vajda, and Kurt Keutzer. Mixed precision quantization of convnets via differentiable neural architecture search. *arXiv preprint arXiv:1812.00090*, 2018.
- [45] Mao Ye, Chengyue Gong, and Qiang Liu. Safer: A structure-free approach for certified robustness to adversarial word substitutions. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3465–3475, 2020.
- [46] Shaokai Ye, Kaidi Xu, Sijia Liu, Hao Cheng, Jan-Henrik Lambrechts, Huan Zhang, Aojun Zhou, Kaisheng Ma, Yanzhi Wang, and Xue Lin. Adversarial robustness vs. model compression, or both? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 111–120, 2019.
- [47] Dongqing Zhang, Jiaolong Yang, Dongqiangzi Ye, and Gang Hua. Lq-nets: Learned quantization for highly accurate and compact deep neural networks. In *Proceedings of the European conference on computer vision (ECCV)*, pages 365–382, 2018.
- [48] Shuchang Zhou, Yuxin Wu, Zekun Ni, Xinyu Zhou, He Wen, and Yuheng Zou. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*, 2016.
- [49] Shu-Chang Zhou, Yu-Zhi Wang, He Wen, Qin-Yao He, and Yu-Heng Zou. Balanced quantization: An effective and efficient approach to quantized neural networks. *Journal of Computer Science and Technology*, 32(4):667–682, 2017.