

Sparsity Meets Robustness: Channel Pruning for the Feynman-Kac Formalism Principled Robust Deep Neural Nets

Thu Dinh¹, Bao Wang², Andrea Bertozzi², Stanley Osher², and Jack Xin¹

¹ University of California, Irvine, USA
thud2@uci.edu, jxin@math.uci.edu

² University of California, Los Angeles, USA
wangbaonj@gmail.com, {bertozzi,sjo}@math.ucla.edu

Abstract. Deep neural nets (DNNs) compression is crucial for adaptation to mobile devices. Though many successful algorithms exist to compress naturally trained DNNs, developing efficient and stable compression algorithms for robustly trained DNNs remains widely open. In this paper, we focus on a co-design of efficient DNN compression algorithms and sparse neural architectures for robust and accurate deep learning. Such a co-design enables us to advance the goal of accommodating both sparsity and robustness. With this objective in mind, we leverage the relaxed augmented Lagrangian based algorithms to prune the weights of adversarially trained DNNs, at both structured and unstructured levels. Using a Feynman-Kac formalism principled robust and sparse DNNs, we can at least double the channel sparsity of the adversarially trained ResNet20 for CIFAR10 classification, meanwhile, improve the natural accuracy by 8.69% and the robust accuracy under the benchmark 20 iterations of IFGSM attack by 5.42%.

1 Introduction

Robust deep neural nets (DNNs) compression is a fundamental problem for secure AI applications in resource-constrained environments such as biometric verification and facial login on mobile devices, and computer vision tasks for the internet of things (IoT) [6,37,24]. Though compression and robustness have been separately addressed in recent years, much less is studied when both players are present and must be satisfied.

To date, many successful techniques have been developed to compress naturally trained DNNs, including neural architecture re-design or searching [18,43], pruning including structured (weights sparsification) [15,30] and unstructured (channel-, filter-, layer-wise sparsification) [36,17], quantization [44,39,7], low-rank approximation [8], knowledge distillation [27], and many more [1,22].

The adversarially trained (AT) DNN is more robust than the naturally trained (NT) DNN to adversarial attacks [23,2]. However, adversarial training (denoted as AT if no ambiguity arises, and the same for NT) also dramatically

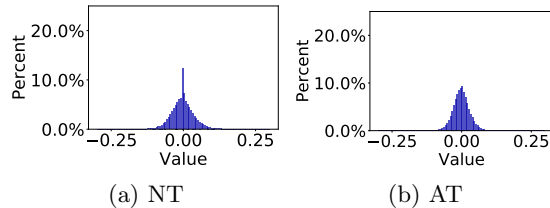


Fig. 1. Histograms of the ResNet20’s weights.

reduces the sparsity of the trained DNN’s weights. As shown in Fig. 1, start from the same default initialization in PyTorch, the NT ResNet20’s weights are much sparser than that of the AT counterpart, for instance, the percent of weights that have magnitude less than 10^{-3} for NT and AT ResNet20 are 8.66% and 3.64% (averaged over 10 trials), resp. This observation motivates us to consider the following two questions:

- 1. *Can we re-design the neural architecture with minimal change on top of the existing one such that the new DNN has sparser weights and better robustness and accuracy than the existing one?*
- 2. *Can we develop efficient compression algorithms to compress the AT DNNs with minimal robustness and accuracy degradations?*

We note that under the AT, the recently proposed Feynman-Kac formalism principled ResNet ensemble [32] has much sparser weights than the standard ResNet, which gives a natural answer to the first question above. To answer the second question, we leverage state-of-the-art relaxed augmented Lagrangian based sparsification algorithms [9,36] to perform both structured and unstructured pruning for the AT DNNs. We focus on unstructured and channel pruning in this work.

1.1 Notation

Throughout this paper we use bold upper-case letters \mathbf{A} , \mathbf{B} to denote matrices, bold lower-case letters \mathbf{x} , \mathbf{y} to denote vectors, and lower case letters x , y and α , β to denote scalars. For vector $\mathbf{x} = (x_1, \dots, x_d)^\top$, we use $\|\mathbf{x}\| = \|\mathbf{x}\|_2 = \sqrt{x_1^2 + \dots + x_d^2}$ to represent its ℓ_2 -norm; $\|\mathbf{x}\|_1 = \sum_{i=1}^d |x_i|$ to represent its ℓ_1 -norm; and $\|\mathbf{x}\|_0 = \sum_{i=1}^d \chi_{\{x_i \neq 0\}}$ to represent its ℓ_0 -norm. For a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, we use $\nabla f(\cdot)$ to denote its gradient. Generally, \mathbf{w}^t represents the set of all parameters of the network being discussed at iteration t , e.g. $\mathbf{w}^t = (\mathbf{w}_1^t, \mathbf{w}_2^t, \dots, \mathbf{w}_M^t)$, where \mathbf{w}_j^t is the weight on the j^{th} layer of the network at the t -th iteration. Similarly, $\mathbf{u}^t = (\mathbf{u}_1^t, \mathbf{u}_2^t, \dots, \mathbf{u}_M^t)$ is a set of weights with the same dimension as \mathbf{w}^t , whose value depends on \mathbf{w}^t and will be defined below. We use $\mathcal{N}(\mathbf{0}, \mathbf{I}_{d \times d})$ to represent the d -dimensional Gaussian, and use notation $O(\cdot)$ to hide only absolute constants which do not depend on any problem parameter.

1.2 Organization

This paper is organized in the following way: In section 2, we list the most related work to this paper. In section 3, we show that the weights of the recently proposed Feynman-Kac formalism principled ResNet ensemble are much sparser than that of the baseline ResNet, providing greater efficiency for compression. In section 4, we present relaxed augmented Lagrangian-based algorithms along with theoretical analysis for both unstructured and channeling pruning of AT DNNs. The numerical results are presented in section 5, followed by concluding remarks. Technical proofs and more related results are provided in the appendix.

2 Related Work

Compression of AT DNNs: [13] considered a low-rank form of the DNN weight matrix with ℓ_0 constraints on the matrix factors in the AT setting. Their training algorithm is a projected gradient descent (PGD) [23] based on the worst adversary. In their paper, the sparsity in matrix factors are unstructured and require additional memory.

Sparsity and Robustness: [14] shows that there is a relationship between the sparsity of weights in the DNN and its adversarial robustness. They showed that under certain conditions, sparsity can improve the DNN’s adversarial robustness. The connection between sparsity and robustness has also been studied recently by [38], [28], and et al. In our paper, we focus on designing efficient pruning algorithms integrated with sparse neural architectures to advance DNNs’ sparsity, accuracy, and robustness.

Feynman-Kac formalism principled Robust DNNs: Neural ordinary differential equations (ODEs) [5] are a class of DNNs that use an ODE to describe the data flow of each input data. Instead of focusing on modeling the data flow of each individual input data, [32,31,21] use a transport equation (TE) to model the flow for the whole input distribution. In particular, from the TE viewpoint, [32] modeled training ResNet [16] as finding the optimal control of the following TE

$$\begin{cases} \frac{\partial u}{\partial t}(\mathbf{x}, t) + G(\mathbf{x}, \mathbf{w}(t)) \cdot \nabla u(\mathbf{x}, t) = 0, & \mathbf{x} \in \mathbb{R}^d, \\ u(\mathbf{x}, 1) = g(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^d, \\ u(\mathbf{x}_i, 0) = y_i, & \mathbf{x}_i \in T, \text{ with } T \text{ being the training set.} \end{cases} \quad (1)$$

where $G(\mathbf{x}, \mathbf{w}(t))$ encodes the architecture and weights of the underlying ResNet, $u(\mathbf{x}, 0)$ serves as the classifier, $g(\mathbf{x})$ is the output activation of ResNet, and y_i is the label of \mathbf{x}_i .


[32] interpreted adversarial vulnerability of ResNet as arising from the irregularity of $u(\mathbf{x}, 0)$ of the above TE. To enhance $u(\mathbf{x}, 0)$ ’s regularity, they added a diffusion term, $\frac{1}{2}\sigma^2 \Delta u(\mathbf{x}, t)$, to the governing equation of (1) which resulting in the convection-diffusion equation (CDE). By the Feynman-Kac formula, $u(\mathbf{x}, 0)$ of the CDE can be approximated by the following two steps:

- Modify ResNet by injecting Gaussian noise to each residual mapping.
- Average the output of n jointly trained modified ResNets, and denote it as En_nResNet .

[32] have noticed that EnResNet can improve both natural and robust accuracies of the AT DNNs. In this work, we leverage the sparsity advantage of EnResNet to push the sparsity limit of the AT DNNs.

3 Regularity and Sparsity of the Feynman-Kac Formalism Principled Robust DNNs’ Weights

From a partial differential equation (PDE) viewpoint, a diffusion term to the governing equation (1) not only smooths $u(\mathbf{x}, 0)$, but can also enhance regularity of the velocity field $G(\mathbf{x}, \mathbf{w}(t))$ [19]. As a DNN counterpart, we expect that when we plot the weights of EnResNet and ResNet at a randomly select layer, the pattern of the former one will look smoother than the latter one. To validate this, we follow the same AT with the same parameters as that used in [32] to train $\text{En}_5\text{ResNet20}$ and ResNet20 , resp. After the above two robust models are trained, we randomly select and plot the weights of a convolutional layer of ResNet20 whose shape is $64 \times 64 \times 3 \times 3$ and plot the weights at the same layer of the first ResNet20 in $\text{En}_5\text{ResNet20}$. As shown in Fig. 2 (a) and (b), most of $\text{En}_5\text{ResNet20}$ ’s weights are close to 0 and they are more regularly distributed in the sense that the neighboring weights are closer to each other than ResNet20 ’s weights. The complete visualization of this randomly selected layer’s weights is shown in the appendix. As shown in Fig. 2 (c) and (d), the weights of $\text{En}_5\text{ResNet20}$ are more concentrated at zero than that of ResNet20 , and most of the $\text{En}_5\text{ResNet20}$ ’s weights are close to zero.



(a) ResNet20 (AT) (b) $\text{En}_5\text{ResNet20}$ (AT) (c) ResNet20 (AT) (d) $\text{En}_5\text{ResNet20}$ (AT)

Fig. 2. (a) and (b): weights visualization; (c) and (d): histogram of weights.

4 Unstructured and Channel Pruning with AT

4.1 Algorithms

In this subsection, we introduce relaxed, augmented Lagrangian-based, pruning algorithms to sparsify the AT DNNs. The algorithms of interest are the Relaxed Variable-Splitting Method (RVSM) [9] for weight pruning (Algorithm 1), and its

variation, the Relaxed Group-wise Splitting Method (RGSM) [36] for channel pruning (Algorithm 2).

Our approach is to apply the RVSM/RGSM algorithm together with robust PGD training to train and sparsify the model from scratch. Namely, at each iteration, we apply a PGD attack to generate adversarial images \mathbf{x}' , which are then used in the forward-propagation process to generate predictions y' . The back-propagation process will then compute the appropriate loss function and apply RVSM/RGSM to update the model. Previous work on RVSM mainly focused on a one-hidden layer setting; In this paper, we extend this result to the general setting. To the best of our knowledge, this is the first result that uses RVSM/RGSM in an adversarial training scenario.

To explain our choice of algorithm, we discuss a classical algorithm to promote sparsity of the target weights, the alternating direction of multiplier method (ADMM) [3,10]. In ADMM, instead of minimizing the original loss function $f(\mathbf{w})$, we seek to minimize the ℓ_1 regularized loss function, $f(\mathbf{w}) + \lambda\|\mathbf{u}\|_1$, by considering the following augmented Lagrangian

$$\mathcal{L}(\mathbf{w}, \mathbf{u}, \mathbf{z}) = f(\mathbf{w}) + \lambda\|\mathbf{u}\|_1 + \langle \mathbf{z}, \mathbf{w} - \mathbf{u} \rangle + \frac{\beta}{2}\|\mathbf{w} - \mathbf{u}\|^2, \quad \lambda, \beta \geq 0. \quad (2)$$

which can be easily solved by applying the following iterations

$$\begin{cases} \mathbf{w}^{t+1} \leftarrow \arg \min_{\mathbf{w}} \mathcal{L}_{\beta}(\mathbf{w}, \mathbf{u}^t, \mathbf{z}^t) \\ \mathbf{u}^{t+1} \leftarrow \arg \min_{\mathbf{u}} \mathcal{L}_{\beta}(\mathbf{w}^{t+1}, \mathbf{u}, \mathbf{z}^t) \\ \mathbf{z}^{t+1} \leftarrow \mathbf{z}^t + \beta(\mathbf{w}^{t+1} - \mathbf{u}^{t+1}) \end{cases} \quad (3)$$

Although widely used in practice, ADMM has several drawbacks when it is used to regularize DNN's weights. First, one can improve the sparsity of the final learned weights by replacing $\|\mathbf{u}\|_1$ with $\|\mathbf{u}\|_0$; but $\|\cdot\|_0$ is not differentiable, thus current theory of optimization does not apply [34]. Second, the update $\mathbf{w}^{t+1} \leftarrow \arg \min_{\mathbf{w}} \mathcal{L}_{\beta}(\mathbf{w}, \mathbf{u}^t, \mathbf{z}^t)$ is not a reasonable step in practice, as one has to fully know how $f(\mathbf{w})$ behaves. In most ADMM adaptation on DNN, this step is replaced by a simple gradient descent. Third, the Lagrange multiplier term, $\langle \mathbf{z}, \mathbf{w} - \mathbf{u} \rangle$, seeks to close the gap between \mathbf{w}^t and \mathbf{u}^t , and this in turn reduces sparsity of \mathbf{u}^t .

The RVSM we will implement is a relaxation of ADMM. RVSM drops the Lagrangian multiplier, and replaces $\lambda\|\mathbf{u}\|_1$ with $\lambda\|\mathbf{u}\|_0$, and resulting in the following relaxed augmented Lagrangian

$$\mathcal{L}_{\beta}(\mathbf{w}, \mathbf{u}) = f(\mathbf{w}) + \lambda\|\mathbf{u}\|_0 + \frac{\beta}{2}\|\mathbf{w} - \mathbf{u}\|^2. \quad (4)$$

The above relaxed augmented Lagrangian can be solved efficiently by the iteration in Algorithm 1. RVSM can resolve all the three issues associated with ADMM listed above in training robust DNNs with sparse weights: First, by removing the linear term $\langle \mathbf{z}, \mathbf{w} - \mathbf{u} \rangle$, one has a closed form formula for the update of \mathbf{u}^t without requiring $\|\mathbf{u}\|_0$ to be differentiable. Explicitly, $\mathbf{u}^t = H_{\sqrt{2\lambda/\beta}}(\mathbf{w}^t) =$

$(w_1^t \chi_{\{|w_1| > \sqrt{2\lambda/\beta}\}}, \dots, w_d^t \chi_{\{|w_d| > \sqrt{2\lambda/\beta}\}})$, where $H_\alpha(\cdot)$ is the hard-thresholding operator with parameter α . Second, the update of \mathbf{w}^t is a gradient descent step itself, so the theoretical guarantees will not deviate from practice. Third, without the Lagrange multiplier term \mathbf{z}^t , there will be a gap between \mathbf{w}^t and \mathbf{u}^t at the limit (finally trained DNNs). However, the limit of \mathbf{u}^t is much sparser than that in the case of ADMM. At the end of each training epoch, we replace \mathbf{w}^t by \mathbf{u}^t for the validation process. Numerical results in Section 5 will show that the AT DNN with parameters \mathbf{u}^t usually outperforms the traditional ADMM in both accuracy and robustness.

Algorithm 1 RVSM

Input: $\eta, \beta, \lambda, max_{epoch}, max_{batch}$
Initialization: \mathbf{w}^0
Define: $\mathbf{u}^0 = H_{\sqrt{2\lambda/\beta}}(\mathbf{w}^0)$
for $t = 0, 1, 2, \dots, max_{epoch}$ **do**
 for $batch = 1, 2, \dots, max_{batch}$ **do**
 $\mathbf{w}^{t+1} \leftarrow \mathbf{w}^t - \eta \nabla f(\mathbf{w}^t) - \eta \beta (\mathbf{w}^t - \mathbf{u}^t)$
 $\mathbf{u}^{t+1} \leftarrow \arg \min_{\mathbf{u}} \mathcal{L}_\beta(\mathbf{u}, \mathbf{w}^t) = H_{\sqrt{2\lambda/\beta}}(\mathbf{w}^t)$
 end for
end for

Algorithm 2 RGSM

Input: $\eta, \beta, \lambda_1, \lambda_2, max_{epoch}, max_{batch}$
Objective: $\tilde{f}(\mathbf{w}) = f(\mathbf{w}) + \lambda_2 \|\mathbf{w}\|_{GL}$
Initialization: Initialize \mathbf{w}^0 , define \mathbf{u}^0
for $g = 1, 2, \dots, G$ **do**
 $\mathbf{u}_g^0 = Prox_{\lambda_1}(\mathbf{w}_g^0)$
end for
for $t = 0, 1, 2, \dots, max_{epoch}$ **do**
 for $batch = 1, 2, \dots, max_{batch}$ **do**
 $\mathbf{w}^{t+1} = \mathbf{w}^t - \eta \nabla \tilde{f}(\mathbf{w}^t) - \eta \beta (\mathbf{w}^t - \mathbf{u}^t)$
 for $g = 1, 2, \dots, G$ **do**
 $\mathbf{u}_g^{t+1} = Prox_{\lambda_1}(\mathbf{w}_g^t)$
 end for
 end for
end for

RGSM is a method that generalizes RVSM to structured pruning, in particular, channel pruning. Let $\mathbf{w} = \{\mathbf{w}_1, \dots, \mathbf{w}_g, \dots, \mathbf{w}_G\}$ be the grouped weights of convolutional layers of a DNN, where G is the total number of groups. Let I_g be the indices of \mathbf{w} in group g . The group Lasso (GLasso) penalty and group- ℓ_0

penalty [40] are defined as

$$\|\mathbf{w}\|_{GL} := \sum_{g=1}^G \|\mathbf{w}_g\|_2, \quad \|\mathbf{w}\|_{G\ell_0} := \sum_{g=1}^G \mathbf{1}_{\|\mathbf{w}_g\|_2 \neq 0} \quad (5)$$

and the corresponding Proximal (projection) operators are

$$\text{Prox}_{GL,\lambda}(\mathbf{w}_g) := \text{sgn}(\mathbf{w}_g) \max(\|\mathbf{w}_g\|_2 - \lambda, 0), \quad \text{Prox}_{G\ell_0,\lambda}(\mathbf{w}_g) := \mathbf{w}_g \mathbf{1}_{\|\mathbf{w}_g\|_2 \neq \sqrt{2\lambda}} \quad (6)$$

where $\text{sgn}(\mathbf{w}_g) := \mathbf{w}_g / \|\mathbf{w}_g\|_2$. The RGSM method is described in Algorithm 2, which improves on adding group Lasso penalty directly in the objective function [35] for natural DNN training [36].

4.2 Theoretical Guarantees

We propose a convergence analysis of the RVSM algorithm to minimize the Lagrangian (4). Consider the following empirical adversarial risk minimization (EARM)

$$\min_{f \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n \max_{\|\mathbf{x}'_i - \mathbf{x}_i\|_\infty \leq \epsilon} L(F(\mathbf{x}'_i, \mathbf{w}), y_i) \quad (7)$$

where the classifier $F(\cdot, \mathbf{w})$ is a function in the hypothesis class \mathcal{H} , e.g. ResNet and its ensembles, parametrized by \mathbf{w} . Here, $L(F(\mathbf{x}_i, \mathbf{w}), y_i)$ is the appropriate loss function associated with F on the data-label pair (\mathbf{x}_i, y_i) , e.g. cross-entropy for classification and root mean square error for regression problem. Since our model is trained using PGD AT, let

$$f(\mathbf{w}) = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\max_{\mathbf{x}'} L(F(\mathbf{x}', \mathbf{w}), y)] \quad (8)$$

where \mathbf{x}' is obtained by applying the PGD attack to the clean data \mathbf{x} [32,11,23,25]. In a nutshell, $f(\mathbf{w})$ is the population adversarial loss of the network parameterized by $\mathbf{w} = (\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_M)$. Before proceeding, we first make the following assumption:

Assumption 1. *Let $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_M$ be the weights in the M layers of the given DNN, then there exists a positive constant L such that for all t ,*

$$\|\nabla f(\cdot, \mathbf{w}_j^{t+1}, \cdot) - \nabla f(\cdot, \mathbf{w}_j^t, \cdot)\| \leq L \|\mathbf{w}_j^{t+1} - \mathbf{w}_j^t\|, \text{ for } j = 1, 2, \dots, M. \quad (9)$$

Assumption 1 is a weaker version of that made by [33,29], in which the empirical adversarial loss function is smooth in both the input \mathbf{x} and the parameters \mathbf{w} . Here we only require the population adversarial loss f to be smooth in each layer of the DNN in the region of iterations. An important consequence of Assumption 1 is

$$f(\cdot, \mathbf{w}_j^{t+1}, \cdot) - f(\cdot, \mathbf{w}_j^t, \cdot) \leq \langle \nabla f(\cdot, \mathbf{w}_j^t, \cdot), (0, \dots, \mathbf{w}_j^{t+1} - \mathbf{w}_j^t, 0, \dots) \rangle + \frac{L}{2} \|\mathbf{w}_j^{t+1} - \mathbf{w}_j^t\|^2 \quad (10)$$

Theorem 1. *Under the Assumption 1, suppose also that the RVSM algorithm is initiated with a small stepsize η such that $\eta < \frac{2}{\beta+L}$. Then the Lagrangian $\mathcal{L}_\beta(\mathbf{w}^t, \mathbf{u}^t)$ decreases monotonically and converges sub-sequentially to a limit point $(\bar{\mathbf{w}}, \bar{\mathbf{u}})$.*

The proof of Theorem 1 is provided in the Appendix. From the descent property of $\mathcal{L}_\beta(\mathbf{w}^t, \mathbf{u}^t)$, classical results from optimization [26] can be used to show that after $T = O(1/\epsilon^2)$ iterations, we have $\nabla_{\mathbf{w}^t} \mathcal{L}_\beta(\mathbf{w}^t, \mathbf{u}^t) = O(\epsilon)$, for some $t \in (0, T]$. The term $\|\mathbf{u}\|_0$ promotes sparsity and $\frac{\beta}{2}\|\mathbf{w} - \mathbf{u}\|^2$ helps keep \mathbf{w} close to \mathbf{u} . Since $\mathbf{u} = H_{\sqrt{2\lambda/\beta}}(\mathbf{w})$, it follows that $\bar{\mathbf{w}}$ will have lots of very small (and thus negligible) components. This explains the sparsity in the limit $\bar{\mathbf{u}}$.

5 Numerical Results

In this section, we verify the following advantages of the proposed algorithms:

- RVSM/RGSM is efficient for unstructured/channel-wise pruning for the AT DNNs.
- After pruning by RVSM and RGSM, EnResNet’s weights are significantly sparser than the baseline ResNet’s, and more accurate in classifying both natural and adversarial images.

These two merits lead to the fact that a synergistic integration of RVSM/RGSM with the Feynman-Kac formula principled EnResNet enables sparsity to meet robustness.

We perform AT by PGD integrated with RVSM, RGSM, or other sparsification algorithms on-the-fly. For all the experiments below, we run 200 epochs of the PGD (10 iterations of the iterative fast gradient sign method (IFGSM¹⁰) with $\alpha = 2/255$ and $\epsilon = 8/255$, and an initial random perturbation of magnitude ϵ). The initial learning rate of 0.1 decays by a factor of 10 at the 80th, 120th, and 160th epochs, and the RVSM/RGSM/ADMM sparsification takes place in the back-propagation stage. We split the training data into 45K/5K for training and validation, and the model with the best validation accuracy is used for testing. We test the trained models on the clean images and attack them by FGSM, IFGSM²⁰, and C&W with the same parameters as that used in [32,41,23]. We denote the accuracy on the clean images and under the FGSM, IFGSM²⁰ [12], C&W [4], and NAttack [20]³ attacks as \mathbf{A}_1 , \mathbf{A}_2 , \mathbf{A}_3 , \mathbf{A}_4 , and \mathbf{A}_5 , resp. A brief introduction of these attacks is available in the appendix. We use both sparsity and channel sparsity to measure the performance of the pruning algorithms, where the sparsity is defined to be the percentage of zero weights; the channel sparsity is the percentage of channels whose weights’ ℓ_2 norm is less than $1E - 15$.

³ For NAttack, we use the default parameters in <https://github.com/cmhcbb/attackbox>.

5.1 Model Compression for AT ResNet and EnResNets

First, we show that RVSM is efficient to sparsify ResNet and EnResNet. Table 1 shows the accuracies of ResNet20 and $\text{En}_2\text{ResNet20}$ under the unstructured sparsification with different sparsity controlling parameter β . We see that after the unstructured pruning by RVSM, $\text{En}_2\text{ResNet20}$ has much sparser weights than ResNet20. Moreover, the sparsified $\text{En}_2\text{ResNet20}$ is remarkably more accurate and robust than ResNet20. For instance, when $\beta = 0.5$, $\text{En}_2\text{ResNet20}$'s weights are 16.42% sparser than ResNet20's (56.34% vs. 39.92%). Meanwhile, $\text{En}_2\text{ResNet20}$ boost the natural and robust accuracies of ResNet20 from 74.08%, 50.64%, 46.67%, and 57.24% to 78.47%, 56.13%, 49.54%, and 65.57%, resp. We perform a few independent trials, and the random effects is small.

Table 1. Accuracy and sparsity of ResNet20 and $\text{En}_2\text{ResNet20}$ under different attacks and β , with $\lambda = 1E - 6$. (Unit: %, n/a: do not perform sparsification. Same for all the following tables.)

β	ResNet20					$\text{En}_2\text{ResNet20}$				
	\mathbf{A}_1	\mathbf{A}_2	\mathbf{A}_3	\mathbf{A}_4	Sparsity	\mathbf{A}_1	\mathbf{A}_2	\mathbf{A}_3	\mathbf{A}_4	Sparsity
n/a	76.07	51.24	47.25	59.30	0	80.34	57.11	50.02	66.77	0
0.01	70.26	46.68	43.79	55.59	80.91	72.81	51.98	46.62	63.10	89.86
0.1	73.45	49.48	45.79	57.72	56.88	77.78	55.48	49.26	65.56	70.55
0.5	74.08	50.64	46.67	57.24	39.92	78.47	56.13	49.54	65.57	56.34

Second, we verify the effectiveness of RGSM in channel pruning. We lists the accuracy and channel sparsity of ResNet20, $\text{En}_2\text{ResNet20}$, and $\text{En}_5\text{ResNet20}$ in Table 2. Without any sparsification, $\text{En}_2\text{ResNet20}$ improves the four type of accuracies by 4.27% (76.07% vs. 80.34%), 5.87% (51.24% vs. 57.11%), 2.77% (47.25% vs. 50.02%), and 7.47% (59.30% vs. 66.77%), resp. When we set $\beta = 1$, $\lambda_1 = 5e - 2$, and $\lambda_2 = 1e - 5$, after channel pruning both natural and robust accuracies of ResNet20 and $\text{En}_2\text{ResNet20}$ remain close to the unsparsified models, but $\text{En}_2\text{ResNet20}$'s weights are 33.48% (41.48% vs. 8%) sparser than that of ResNet20's. When we increase the channel sparsity level by increasing λ_1 to $1e - 1$, both the accuracy and channel sparsity gaps between ResNet20 and $\text{En}_2\text{ResNet20}$ are enlarged. $\text{En}_5\text{ResNet20}$ can future improve both natural and robust accuracies on top of $\text{En}_2\text{ResNet20}$. For instance, at $\sim 55\%$ (53.36% vs. 56.74%) channel sparsity, $\text{En}_5\text{ResNet20}$ can improve the four types of accuracy of $\text{En}_2\text{ResNet20}$ by 4.66% (80.53% vs. 75.87%), 2.73% (57.38% vs. 54.65%), 2.86% (50.63% vs. 47.77%), and 1.11% (66.52% vs. 65.41%), resp.

Third, we show that an ensemble of small ResNets via the Feynman-Kac formalism performs better than a larger ResNet of roughly the same size in accuracy, robustness, and sparsity. We AT $\text{En}_2\text{ResNet20}$ ($\sim 0.54\text{M}$ parameters) and ResNet38 ($\sim 0.56\text{M}$ parameters) with and without channel pruning. As shown in Table 3, under different sets of parameters, after RGSM pruning, $\text{En}_2\text{ResNet20}$

Table 2. Accuracy and sparsity of different EnResNet20. (Ch. Sp.: Channel Sparsity)

Net	β	λ_1	λ_2	A_1	A_2	A_3	A_4	A_5	Ch. Sp.
ResNet20	n/a	n/a	n/a	76.07	51.24	47.25	59.30	45.88	0
	1	5.E-02	1.E-05	75.91	51.52	47.14	58.77	45.02	8.00
	1	1.E-01	1.E-05	71.84	48.23	45.21	57.09	43.84	25.33
En ₂ ResNet20	n/a	n/a	n/a	80.34	57.11	50.02	66.77	49.35	0
	1	5.E-02	1.E-05	78.28	56.53	49.58	66.56	49.11	41.48
	1	1.E-01	1.E-05	75.87	54.65	47.77	65.41	46.77	56.74
En ₅ ResNet20	n/a	n/a	n/a	81.41	58.21	51.60	66.48	50.21	0
	1	1.E-02	1.E-05	81.46	58.34	51.35	66.84	50.07	19.76
	1	2.E-02	1.E-05	80.53	57.38	50.63	66.52	48.23	53.36

always has much more channel sparsity than ResNet38, also much more accurate and robust. For instance, when we set $\beta = 1$, $\lambda_1 = 5e - 2$, and $\lambda_2 = 1e - 5$, the AT ResNet38 and En₂ResNet20 with channel pruning have channel sparsity 17.67% and 41.48%, resp. Meanwhile, En₂ResNet20 outperforms ResNet38 in the four types of accuracy by 0.36% (78.28% vs. 77.92%), 3.02% (56.53% vs. 53.51%), 0.23% (49.58% vs. 49.35%), and 6.34% (66.56% vs. 60.32%), resp. When we increase λ_1 , the channel sparsity of two nets increase.. As shown in Fig. 3, En₂ResNet20’s channel sparsity growth much faster than ResNet38’s, and we plot the corresponding four types of accuracies of the channel sparsified nets in Fig. 4.

Table 3. Performance of En₂ResNet20 and ResNet38 under RVSM.

Net	β	λ_1	λ_2	A_1	A_2	A_3	A_4	Ch. Sp.
En ₂ ResNet20	n/a	n/a	n/a	80.34	57.11	50.02	66.77	0
ResNet38	n/a	n/a	n/a	78.03	54.09	49.81	61.72	0
En ₂ ResNet20	1	5.E-02	1.E-05	78.28	56.53	49.58	66.56	41.48
ResNet38	1	5.E-02	1.E-05	77.92	53.51	49.35	60.32	17.67
En ₂ ResNet20	1	1.E-01	1.E-05	76.30	54.65	47.77	65.41	56.74
ResNet38	1	1.E-01	1.E-05	72.95	49.78	46.48	57.92	43.80




Fig. 3. Sparsity of En₂ResNet20 and ResNet38 under different λ_1 . (5 runs)




Fig. 4. Accuracy of En₂ResNet20 and ResNet38 under different λ_1 . (5 runs)

Table 4. Contrasting ADMM versus RVSM for the AT ResNet20.

	Unstructured Pruning					Channel Pruning				
	A_1	A_2	A_3	A_4	Sp.	A_1	A_2	A_3	A_4	Ch. Sp.
RVSM	70.26	46.68	43.79	55.59	80.91	71.84	48.23	45.21	57.09	25.33
ADMM	71.55	47.37	44.30	55.79	10.92	63.99	42.06	39.75	51.90	4.44

5.2 RVSM/RGSM versus ADMM

In this subsection, we will compare RVSM, RGSM, and ADMM [42]⁴ for unstructured and channel pruning for the AT ResNet20, and we will show that RVSM and RGSM iterations can promote much higher sparsity with less natural and robust accuracies degradations than ADMM. We list both natural/robust accuracies and sparsities of ResNet20 after ADMM, RVSM, and RGSM pruning in Table 4. For unstructured pruning, ADMM retains slightly better natural ($\sim 1.3\%$) and robust ($\sim 0.7\%$, $\sim 0.5\%$, and 0.2% under FGSM, IFGSM²⁰, and C&W attacks) accuracies. However, RVSM gives much better sparsity (80.91% vs. 10.89%). In the channel pruning scenario, RVSM significantly outperforms ADMM in all criterion including natural and robust accuracies and channel sparsity, as the accuracy gets improved by at least 5.19% and boost the channel sparsity from 4.44% to 25.33%. Part of the reason for ADMM’s inefficiency in sparsifying DNN’s weights is due to the fact that the ADMM iterations try to close the gap between the weights w^t and the auxiliary variables u^t , so the final result has a lot of weights with small magnitude, but not small enough to be regarded as zero (having norm less than $1e-15$). The RVSM does not seek to

⁴ We use the code from: github.com/KaiqiZhang/admm-pruning

close this gap, instead it replaces the weight \mathbf{w}^t by \mathbf{u}^t , which is sparse, after each epoch. This results in a much sparser final result, as shown in Figure 5: ADMM does result in a lot of channels with small norms; but to completely prune these off, RVSM does a better job. Here, the channel norm is defined to be the ℓ_2 norm of the weights in each channel of the DNN [35].

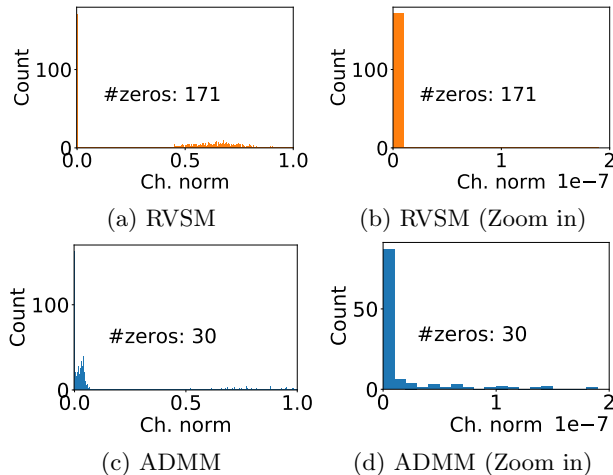


Fig. 5. Channel norms of the AT ResNet20 under RVSM and ADMM.

5.3 Beyond ResNet Ensemble and Beyond CIFAR10

5.4 Beyond ResNet Ensemble

In this part, we show that the idea of average over noise injected ResNet generalizes to other types of DNNs, that is, the natural and robust accuracies of DNNs can be improved by averaging over noise injections. As a proof of concept, we remove all the skipping connections in ResNet20 and $\text{En}_2\text{ResNet20}$, in this case, the model no longer involves a TE and CDE. As shown in Table 5, removing the skip connection degrades the performance significantly, especially for $\text{En}_2\text{ResNet20}$ (all the four accuracies of the AT $\text{En}_2\text{ResNet20}$ reduces more than the AT ResNet20’s). These results confirm that skip connection is crucial, without it the TE model assumption breaks down, and the utility of EnResNets reduces. However, the ensemble of noise injected DNNs still improves the performance remarkably.

Next, let us consider the sparsity, robustness, and accuracies of ResNet20 and $\text{En}_2\text{ResNet20}$ without any skip connection, when they are AT using weights sparsification by either RVSM or RGSM. For RVSM pruning, we set $\beta = 5E - 2$ and $\lambda = 1E - 6$; for RGSM channel pruning, we set $\beta = 5E - 2$, $\lambda_1 = 5E - 2$, and $\lambda_2 = 1E - 5$. We list the corresponding results in Table 6. We see that under both RVSM and RGSM pruning, $\text{En}_2\text{ResNet20}$ is remarkably more accurate on both

Table 5. Accuracies of ResNet20 and En₂ResNet20 with and without skip connections.

Net	Type	\mathbf{A}_1	\mathbf{A}_2	\mathbf{A}_3	\mathbf{A}_4
ResNet20	Base Net	76.07	51.24	47.25	59.30
	Without skip connection	75.45	51.03	47.22	58.44
En ₂ ResNet20	Base net	80.34	57.11	50.02	66.77
	Without skip connection	79.12	55.76	49.92	66.26

clean and adversarial images, and significantly sparser than ResNet20. When we compare the results in Table 6 with that in Tables 1 and 2, we conclude that once we remove the skip connections, the sparsity, robustness, and accuracy degrades dramatically.

Table 6. Sparsity and accuracies of ResNet20 and En₂ResNet20 without skip connection under different pruning algorithms.

Net	Pruning Algorithm	\mathbf{A}_1	\mathbf{A}_2	\mathbf{A}_3	\mathbf{A}_4	Sp.	Ch. Sp.
ResNet20 (no skip connection)	RVSM	73.72	50.46	46.98	58.28	0.05	0.15
	RGSM	74.63	50.44	46.86	58.05	1.64	9.04
En ₂ ResNet20 (no skip connection)	RVSM	76.95	55.17	49.28	58.35	10.87	9.48
	RGSM	78.51	56.55	49.71	67.08	7.95	15.48

5.5 Beyond CIFAR10

Besides CIFAR10, we further show the advantage of EnResNet + RVSM/RGSM in compressing and improving accuracy/robustness of the AT DNNs for CIFAR100 classification. We list the natural and robust accuracies and channel sparsities of the AT ResNet20 and En₂ResNet20 with different RGSM parameters (n/a stands for do not perform channel pruning) in Table 7. For $\lambda_1 = 0.05$, RGSM almost preserves the performance of the DNNs without channel pruning, while improving channel sparsity by 7.11% for ResNet20, and 16.89% for En₂ResNet20. As we increase λ_1 to 0.1, the channel sparsity becomes 18.37% for ResNet20 and 39.23% for En₂ResNet20. Without any channel pruning, En₂ResNet20 improves natural accuracy by 4.66% (50.68% vs. 46.02%), and robust accuracies by 5.25% (30.2% vs. 24.77%), 3.02% (26.25% vs. 23.23%), and 7.64% (40.06% vs. 32.42%), resp., under the FGSM, IFGSM²⁰, and C&W attacks. Even in very high channel sparsity scenario ($\lambda_1 = 0.05$), En₂ResNet20 still dramatically increase \mathbf{A}_1 , \mathbf{A}_2 , \mathbf{A}_3 , and \mathbf{A}_4 by 2.90%, 4.31%, 1.89%, and 5.86%, resp. These results are similar to the one obtained on the CIFAR10 in Table 2. These results further confirm that RGSM together with the Feynman-Kac formalism principled ResNets ensemble can significantly improve both natural/robust accuracy and sparsity on top of the baseline ResNets.

Table 7. Accuracy and sparsity of different Ensembles of ResNet20’s on the CIFAR100.

Net	β	λ_1	λ_2	A_1	A_2	A_3	A_4	Ch. Sp.
ResNet20	n/a	n/a	n/a	46.02	24.77	23.23	32.42	0
	1	5.E-02	1.E-05	45.74	25.34	23.55	33.53	7.11
	1	1.E-01	1.E-05	44.34	24.46	23.12	32.38	18.37
En ₂ ResNet20	n/a	n/a	n/a	50.68	30.2	26.25	40.06	0
	1	5.E-02	1.E-05	50.56	30.33	26.23	39.85	16.89
	1	1.E-01	1.E-05	47.24	28.77	25.01	38.24	39.23

6 Concluding Remarks

The Feynman-Kac formalism principled AT EnResNet’s weights are much sparser than the baseline ResNet’s. Together with the relaxed augmented Lagrangian based unstructured/channel pruning algorithms, we can compress the AT DNNs much more efficiently, meanwhile significantly improves both natural and robust accuracies of the compressed model. As future directions, we propose to quantize EnResNets and to integrate neural ODE into our framework.

7 Acknowledgements

The work was partially supported by NSF grants IIS-1632935 (TD and JX), DMS-1737770 (AB and SO), ATD-1924935 (BW), and DMS-1924548 (TD and JX); ONR Grant N00014181257 (BW and SO); AFOSR MURI FA9550-18-1-0502 (BW and SO); and Google Research Grant W0KPVLW22T86GB8M (TD).

References

1. Alvarez, J.M., Salzmann, M.: Learning the number of neurons in deep networks. In: *Advances in Neural Information Processing Systems*. pp. 2270–2278 (2016)
2. Athalye, A., Carlini, N., Wagner, D.: Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. *arXiv preprint arXiv:1802.00420* (2018)
3. Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J., et al.: Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning* **3**(1), 1–122 (2011)
4. Carlini, N., Wagner, D.: Towards evaluating the robustness of neural networks. *IEEE European Symposium on Security and Privacy* pp. 39–57 (2016)
5. Chen, T., Rubanova, Y., Bettencourt, J., Duvenaud, D.: Neural ordinary differential equations. In: *Advances in neural information processing systems*. pp. 6571–6583 (2018)
6. Cheng, Y., Wang, D., Zhou, P., Zhang, T.: A survey of model compression and acceleration for deep neural networks. *arXiv preprint arXiv:1710.09282* (2017)
7. Courbariaux, M., Hubara, I., Soudry, D., El-Yaniv, R., Bengio, Y.: Binarynet: Training deep neural networks with weights and activations constrained to +1 or -1. *ArXiv abs/1602.02830* (2016)
8. Denil, M., Shakibi, B., Dinh, L., Ranzato, M., de Freitas, N.: Predicting parameters in deep learning. In: *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2*. pp. 2148–2156. NIPS’13, Curran Associates Inc., USA (2013), <http://dl.acm.org/citation.cfm?id=2999792.2999852>
9. Dinh, T., Xin, J.: Convergence of a relaxed variable splitting method for learning sparse neural networks via ℓ_1 , ℓ_0 , and transformed- ℓ_1 penalties. *arXiv preprint arXiv:1812.05719* (2018)
10. Goldstein, T., Osher, S.: The split bregman method for l_1 -regularized problems. *SIAM journal on imaging sciences* **2**(2), 323–343 (2009)
11. Goodfellow, I., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. *CoRR abs/1412.6572* (2014)
12. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6275* (2014)
13. Gui, S., Wang, H., Yu, C., Yang, H., Wang, Z., Liu, J.: Adversarially trained model compression: When robustness meets efficiency. *arXiv preprint arXiv:1902.03538* (2019)
14. Guo, Y., Zhang, C., Zhang, C., Chen, Y.: Sparse dnns with improved adversarial robustness. In: *Advances in neural information processing systems*. pp. 242–251 (2018)
15. Han, S., Pool, J., Tran, J., Dally, W.: Learning both weights and connections for efficient neural network. In: *Advances in neural information processing systems*. pp. 1135–1143 (2015)
16. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
17. He, Y., Zhang, X., Sun, J.: Channel pruning for accelerating very deep neural networks. In: *The IEEE International Conference on Computer Vision (ICCV)* (Oct 2017)
18. Howard, A., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017)

19. Ladyženskaja, O., Solonnikov, V., Ural'ceva, N.: Linear and quasi-linear equations of parabolic type, vol. 23. American Mathematical Soc. (1988)
20. Li, Y., Li, L., Wang, L., Zhang, T., Gong, B.: Nattack: Learning the distributions of adversarial examples for an improved black-box attack on deep neural networks. arXiv preprint arXiv:1905.00441 (2019)
21. Li, Z., Shi, Z.: Deep residual learning and pdes on manifold. arXiv preprint arXiv:1708.05115 (2017)
22. Liu, B., Wang, M., Foroosh, H., Tappen, M., Pensky, M.: Sparse convolutional neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 806–814 (2015)
23. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. In: International Conference on Learning Representations (2018), <https://openreview.net/forum?id=rJzIBfZAb>
24. Mohammadi, M., Al-Fuqaha, A., Sorour, S., Guizani, M.: Deep learning for iot big data and streaming analytics: A survey. IEEE Communications Surveys & Tutorials **20**(4), 2923–2960 (2018)
25. Na, T., Ko, J., Mukhopadhyay, S.: Cascade adversarial machine learning regularized with a unified embedding. In: International Conference on Learning Representations (2018), <https://openreview.net/forum?id=HyRVBzap->
26. Nesterov, Y.: Introductory Lectures on Convex Optimization: A Basic Course. Springer Publishing Company, Incorporated, 1 edn. (2014)
27. Polino, A., Pascanu, R., Alistarh, D.: Model compression via distillation and quantization. arXiv preprint arXiv:1802.05668 (2018)
28. Rakin, A., He, Z., Yang, L., Wang, Y., Wang, L., Fan, D.: Robust sparse regularization: Simultaneously optimizing neural network robustness and compactness. arXiv preprint arXiv:1905.13074 (2019)
29. Sinha, A., Namkoong, H., Duchi, J.: Certifiable distributional robustness with principled adversarial training. In: International Conference on Learning Representations (2018), <https://openreview.net/forum?id=Hk6kPgZA->
30. Srinivas, S., Babu, R.V.: Data-free parameter pruning for deep neural networks. arXiv preprint arXiv:1507.06149 (2015)
31. Wang, B., Luo, X., Li, Z., Zhu, W., Shi, Z., Osher, S.: Deep neural nets with interpolating function as output activation. In: Advances in Neural Information Processing Systems. pp. 743–753 (2018)
32. Wang, B., Yuan, B., Shi, Z., Osher, S.: ResNet ensemble via the Feynman-Kac formalism to improve natural and robust accuracies. In: Advances in Neural Information Processing Systems (2019)
33. Wang, Y., Ma, X., Bailey, J., Yi, J., Zhou, B., Gu, Q.: On the convergence and robustness of adversarial training. In: Chaudhuri, K., Salakhutdinov, R. (eds.) Proceedings of the 36th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 97, pp. 6586–6595. PMLR, Long Beach, California, USA (09–15 Jun 2019), <http://proceedings.mlr.press/v97/wang19i.html>
34. Wang, Y., Zeng, J., Yin, W.: Global Convergence of ADMM in Nonconvex Nonsmooth Optimization. Journal of Scientific Computing, online (2018). <https://doi.org/10.1007/s10915-018-0757-z>
35. Wen, W., Wu, C., Wang, Y., Chen, Y., Li, H.: Learning structured sparsity in deep neural networks. In: Advances in neural information processing systems. pp. 2074–2082 (2016)
36. Yang, B., Lyu, J., Zhang, S., Qi, Y.Y., Xin, J.: Channel pruning for deep neural networks via a relaxed group-wise splitting method. In Proc. of 2nd International Conference on AI for Industries (AI4I), Laguna Hills, CA (2019)

37. Yao, S., Zhao, Y., Zhang, A., Su, L., Abdelzaher, T.: Deepiot: Compressing deep neural network structures for sensing systems with a compressor-critic framework. In: Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems. p. 4. ACM (2017)
38. Ye, S., Xu, K., Liu, S., Cheng, H., Lambrechts, J., Zhang, H., Zhou, A., Ma, K., Wang, Y., Lin, X.: Second rethinking of network pruning in the adversarial setting. arXiv preprint arXiv:1903.12561 (2019)
39. Yin, P., Zhang, S., Lyu, J., Osher, S., Qi, Y., Xin, J.: Blended coarse gradient descent for full quantization of deep neural networks. *Research in the Mathematical Sciences* **6**(1), 14 (Jan 2019). <https://doi.org/10.1007/s40687-018-0177-6>, <https://doi.org/10.1007/s40687-018-0177-6>
40. Yuan, M., Lin, Y.: Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society, Series B* **68**(1), 49–67 (2007)
41. Zhang, H., Yu, Y., Jiao, J., Xing, E., Ghaoui, L., Jordan, M.: Theoretically principled trade-off between robustness and accuracy. arXiv preprint arXiv:1901.08573 (2019)
42. Zhang, T., Ye, S., Zhang, K., Tang, J., Wen, W., Fardad, M., Wang, Y.: A systematic dnn weight pruning framework using alternating direction method of multipliers. arXiv preprint 1804.03294 (Jul 2018), <https://arxiv.org/abs/1804.03294>
43. Zhang, X., Zhou, X., Lin, M., Sun, J.: Shufflenet: An extremely efficient convolutional neural network for mobile devices. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6848–6856 (2018)
44. Zhou, A., Yao, A., Guo, Y., Xu, L., Chen, Y.: Incremental network quantization: Towards lossless cnns with low-precision weights. arXiv preprint arXiv:1702.03044 (2017)

A Proof of Theorem 1

By the arg min update of \mathbf{u}^t , $\mathcal{L}_\beta(\mathbf{w}^{t+1}, \mathbf{u}^{t+1}) \leq \mathcal{L}_\beta(\mathbf{w}^{t+1}, \mathbf{u}^t)$. It remains to show $\mathcal{L}_\beta(\mathbf{w}^{t+1}, \mathbf{u}^t) \leq \mathcal{L}_\beta(\mathbf{w}^t, \mathbf{u}^t)$. To this end, we show

$$\mathcal{L}_\beta(\mathbf{w}_1^{t+1}, \mathbf{w}_2^t, \dots, \mathbf{w}_M^t, \mathbf{u}^t) \leq \mathcal{L}_\beta(\mathbf{w}_1^t, \mathbf{w}_2^t, \dots, \mathbf{w}_M^t, \mathbf{u}^t)$$

the conclusion then follows by a repeated argument. Notice the only change occurs in the first layer \mathbf{w}_1^t . For simplicity of notation, only for the following chain of inequality, let $\mathbf{w}^t = (\mathbf{w}_1^t, \dots, \mathbf{w}_M^t)$ and $\mathbf{w}^{t+1} = (\mathbf{w}_1^{t+1}, \mathbf{w}_2^t, \dots, \mathbf{w}_M^t)$. Then for a fixed $\mathbf{u} := \mathbf{u}^t$, we have

$$\begin{aligned} & \mathcal{L}_\beta(\mathbf{w}^{t+1}, \mathbf{u}) - \mathcal{L}_\beta(\mathbf{w}^t, \mathbf{u}) \\ &= f(\mathbf{w}^{t+1}) - f(\mathbf{w}^t) + \frac{\beta}{2} (\|\mathbf{w}^{t+1} - \mathbf{u}\|^2 - \|\mathbf{w}^t - \mathbf{u}\|^2) \\ &\leq \langle \nabla f(\mathbf{w}^t), \mathbf{w}^{t+1} - \mathbf{w}^t \rangle + \frac{L}{2} \|\mathbf{w}^{t+1} - \mathbf{w}^t\|^2 + \frac{\beta}{2} (\|\mathbf{w}^{t+1} - \mathbf{u}\|^2 - \|\mathbf{w}^t - \mathbf{u}\|^2) \\ &= \frac{1}{\eta} \langle \mathbf{w}^t - \mathbf{w}^{t+1}, \mathbf{w}^{t+1} - \mathbf{w}^t \rangle - \beta \langle \mathbf{w}^t - \mathbf{u}, \mathbf{w}^{t+1} - \mathbf{w}^t \rangle \\ &\quad + \frac{L}{2} \|\mathbf{w}^{t+1} - \mathbf{w}^t\|^2 + \frac{\beta}{2} (\|\mathbf{w}^{t+1} - \mathbf{u}\|^2 - \|\mathbf{w}^t - \mathbf{u}\|^2) \\ &= \frac{1}{\eta} \langle \mathbf{w}^t - \mathbf{w}^{t+1}, \mathbf{w}^{t+1} - \mathbf{w}^t \rangle + \left(\frac{L}{2} + \frac{\beta}{2} \right) \|\mathbf{w}^{t+1} - \mathbf{w}^t\|^2 \\ &\quad + \frac{\beta}{2} \|\mathbf{w}^{t+1} - \mathbf{u}\|^2 - \frac{\beta}{2} \|\mathbf{w}^t - \mathbf{u}\|^2 - \beta \langle \mathbf{w}^t - \mathbf{u}, \mathbf{w}^{t+1} - \mathbf{w}^t \rangle - \frac{\beta}{2} \|\mathbf{w}^{t+1} - \mathbf{w}^t\|^2 \\ &= \left(\frac{L}{2} + \frac{\beta}{2} - \frac{1}{\eta} \right) \|\mathbf{w}^{t+1} - \mathbf{w}^t\|^2 \end{aligned}$$

Thus, when $\eta \leq \frac{2}{\beta+L}$, we have $\mathcal{L}_\beta(\mathbf{w}^{t+1}, \mathbf{u}) \leq \mathcal{L}_\beta(\mathbf{w}^t, \mathbf{u})$. Apply the above argument repeatedly, we arrive at

$$\begin{aligned} \mathcal{L}_\beta(\mathbf{w}^{t+1}, \mathbf{u}^{t+1}) &\leq \mathcal{L}_\beta(\mathbf{w}^{t+1}, \mathbf{u}^t) \\ &\leq \mathcal{L}_\beta(\mathbf{w}_1^{t+1}, \mathbf{w}_2^{t+1}, \dots, \mathbf{w}_{M-1}^{t+1}, \mathbf{w}_M^t, \mathbf{u}^t) \\ &\leq \dots \\ &\leq \mathcal{L}_\beta(\mathbf{w}^t, \mathbf{u}^t) \end{aligned}$$

This implies $\mathcal{L}_\beta(\mathbf{w}^t, \mathbf{u}^t)$ decreases monotonically. Since $\mathcal{L}_\beta(\mathbf{w}^t, \mathbf{u}^t) \geq 0$, $(\mathbf{w}^t, \mathbf{u}^t)$ must converge sub-sequentially to a limit point $(\bar{\mathbf{w}}, \bar{\mathbf{u}})$. This completes the proof.

B Adversarial Attacks Used in This Work

We focus on the ℓ_∞ norm based untargeted attack. For a given image-label pair $\{\mathbf{x}, y\}$, a given ML model $F(\mathbf{x}, \mathbf{w})$, and the associated loss $L(\mathbf{x}, y) := L(F(\mathbf{x}, \mathbf{w}), y)$:

- Fast gradient sign method (FGSM) searches an adversarial, \mathbf{x}' , within an ℓ_∞ -ball as

$$\mathbf{x}' = \mathbf{x} + \epsilon \cdot \text{sign}(\nabla_{\mathbf{x}} L(\mathbf{x}, y)).$$

- Iterative FGSM (IFGSM^M) [12] iterates FGSM and clip the range as

$$\mathbf{x}^{(m)} = \text{Clip}_{\mathbf{x}, \epsilon} \left\{ \mathbf{x}^{(m-1)} + \alpha \cdot \text{sign} \left(\nabla_{\mathbf{x}^{(m-1)}} L(\mathbf{x}^{(m-1)}, y) \right) \right\}$$

with $\mathbf{x}^{(0)} = \mathbf{x}$, $m = 1, \dots, M$.

- C&W attack [4] searches the minimal perturbation (δ) attack as

$$\min_{\delta} \|\delta\|_\infty, \text{ subject to } F(\mathbf{w}, \mathbf{x} + \delta) = t, \mathbf{x} + \delta \in [0, 1]^d, \text{ for } \forall t \neq y.$$

- NAttack [20] is an effective gradient-free attack.

C More Visualizations of the DNNs' Weights

In section 3, we showed some visualization results for part of the weights of a randomly selected convolutional layer of the AT ResNet20 and En₅ResNet20. The complete visualization results of this selected layer are shown in Figs. 6 and 7, resp., for ResNet20 and En₅ResNet20. These plots further verifies that:

- The magnitude of the weights of the adversarially trained En₅ResNet20 is significantly smaller than that of the robustly trained ResNet20.
- The overall pattern of the weights of the adversarially trained En₅ResNet20 is more regular than that of the robustly trained ResNet20.




Fig. 6. Weights of a randomly selected convolutional layer of the PGD AT ResNet20.




Fig. 7. Weights of the PGD AT En_5 ResNet20 at the same layer as that shown in Fig. 6.