

DNR: A Tunable Robust Pruning Framework Through Dynamic Network Rewiring of DNNs

Souvik Kundu

souvikku@usc.edu

University of Southern California
Los Angeles, California

Peter A. Beerel

pabeerel@usc.edu

University of Southern California
Los Angeles, California

Mahdi Nazemi

mnazemi@usc.edu

University of Southern California
Los Angeles, California

Massoud Pedram

pedram@usc.edu

University of Southern California
Los Angeles, California

ABSTRACT

This paper presents a dynamic network rewiring (DNR) method to generate pruned deep neural network (DNN) models that are robust against adversarial attacks yet maintain high accuracy on clean images. In particular, the disclosed DNR method is based on a unified constrained optimization formulation using a hybrid loss function that merges ultra-high model compression with robust adversarial training. This training strategy dynamically adjusts inter-layer connectivity based on per-layer normalized momentum computed from the hybrid loss function. In contrast to existing robust pruning frameworks that require multiple training iterations, the proposed learning strategy achieves an overall target pruning ratio with only a single training iteration and can be tuned to support both irregular and structured channel pruning. To evaluate the merits of DNR, experiments were performed with two widely accepted models, namely VGG16 and ResNet-18, on CIFAR-10, CIFAR-100 as well as with VGG16 on Tiny-ImageNet. Compared to the baseline uncompressed models, DNR provides over 20 \times compression on all the datasets with no significant drop in either clean or adversarial classification accuracy. Moreover, our experiments show that DNR consistently finds compressed models with better clean and adversarial image classification performance than what is achievable through state-of-the-art alternatives. Our models and test codes are available at https://github.com/ksouvik52/DNR_ASP_DAC2021.

CCS CONCEPTS

- Adversarial training → Robustness; • DNN Model Pruning → Irregular and Structured pruning.

KEYWORDS

Pruning, adversarial robustness, model compression.

This work was partly supported by NSF including grant #1763747.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ASPDAC '21, January 18–21, 2021, Tokyo, Japan

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-7999-1/21/01...\$15.00

<https://doi.org/10.1145/3394885.3431542>

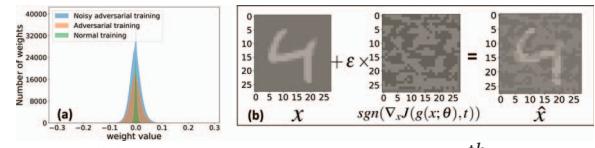


Figure 1: (a) Weight distribution of the 14th convolution layer of ResNet18 model for different training schemes: normal, adversarial [23], and noisy adversarial [14]. (b) An adversarially generated image (\hat{x}) obtained through FGSM attack, which is predicted to be the number 5 instead of 4 (x). ACM Reference Format:

Souvik Kundu, Mahdi Nazemi, Peter A. Beerel, and Massoud Pedram. 2021. DNR: A Tunable Robust Pruning Framework Through Dynamic Network Rewiring of DNNs. In *26th Asia and South Pacific Design Automation Conference (ASPDAC '21), January 18–21, 2021, Tokyo, Japan*. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3394885.3431542>

1 INTRODUCTION

In recent years, deep neural networks (DNNs) have emerged as critical components in various applications, including image classification [17], speech recognition [15], medical image analysis [21] and autonomous driving [3]. However, despite the proliferation of deep learning-powered applications, machine learning models have raised significant security concerns due to their vulnerability to *adversarial examples*, i.e., maliciously generated images which are perceptually similar to clean ones with the ability to fool classifier models into making wrong predictions [2, 8]. Various recent work have proposed associated defense mechanisms including adversarial training [8], hiding gradients [30], adding noise to the weights [14], and several others [24].

Meanwhile, large model sizes have high inference latency, computation, and storage costs that represent significant challenges in deployment on IoT devices. Thus reduced-size models [7, 19] and model compression techniques e.g., *pruning* [4, 5, 12], have gained significant traction. In particular, earlier work showed that without a significant accuracy drop, pruning can remove more than 90% of the model parameters [4, 5] and that ensuring the pruned models have structure can yield observed performance improvements on a broad range of compute platforms [13]. However, adversarial training that increases network robustness generally demands more non-zero parameters than needed for only clean data [23] as illustrated in Fig. 1(a). Thus a naively compressed model performing well on clean images, can become vulnerable to adversarial images.

Unfortunately, despite a plethora of work on compressed model performance on clean data, there have been only a few studies on the robustness of compressed models under adversarial attacks.

In particular, some prior works [9, 31] have tried to design a compressed yet robust model through a unified constrained optimization formulation by using the alternating direction method of multipliers (ADMM) in which dynamic L_2 regularization is the key to outperforming state of the art pruning techniques [27]. However, these efforts require the network designer to specify layer-wise sparsity ratios, which requires prior knowledge of an effective compressed model. This knowledge may not be available and thus training may require multiple iterations to determine good layer-sparsity ratios. In other schemes like Lasso [26], a target compression ratio cannot be set because the final compression ratio is not determined until training is completed. Moreover, Lasso requires separate re-training to increase the accuracy after the assignment of non-significant weights to zero, resulting in costly training.

In contrast, this paper presents dynamic network rewiring (DNR), a unified training framework to find a compressed model with increased robustness that does not require individual per-layer target sparsity ratios. In particular, we introduce a hybrid loss function for robust compression which has three major components: a clean image classification loss, a dynamic L_2 -regularizer term inspired by a relaxed version of ADMM [6], and an adversarial training loss. Inspired by *sparse-learning*-based training scheme of [4], we then propose a single-shot training framework to achieve a robust pruned DNN using the proposed loss. In particular, DNR dynamically arranges per layer pruning ratios using normalized momentum, maintaining the target pruning every epoch, without requiring any fine tuning. In summary, our key contributions are:

- Given only a global pruning ratio, we propose a single-shot (non-iterative) training framework that simultaneously achieves ultra-high compression ratio, state-of-the-art accuracy on clean data, and robustness to perturbed images.
- We extend the approach to support structured pruning technique, namely *channel pruning*, enabling benefits on a broader class of compute platforms. As opposed to conventional sparse-learning [4] that can perform only irregular pruning, models generated through structured DNR can significantly speed up inference. To the best of our knowledge, we are the first to propose a non-iterative robust training framework that supports both irregular and channel pruning.
- We provide a comprehensive investigation of adversarial robustness for both channel and irregular pruning, and obtain insightful observations through evaluation on an extensive set of experiments on CIFAR-10 [16], CIFAR-100 [16], and Tiny-ImageNet [10] using variants of ResNet18 [11] and VGG16 [29]. Our proposed method consistently outperforms state-of-the-art (SOTA) [26, 31] approaches with negligible accuracy drop compared to the unpruned baselines.

We further empirically demonstrate the superiority of our scheme when used to target model compression on clean-only image classification task compared to SOTA non-iterative pruning mechanisms [4, 5, 12, 20].¹

¹This paper targets low-cost training, thus comparisons to iterative pruning methods (e.g., [27]) are out of scope.

The remainder of this paper is structured as follows. In section 2 we present necessary background work. Section 3 describes proposed DNR based training method. We present our experimental results in Section 4 and conclude in Section 5.

2 BACKGROUND WORK

2.1 Adversarial Attacks

Recently, various adversarial attacks have been proposed to find fake images, i.e., adversarial examples, which have barely-visible perturbations from real images but still manage to fool a trained DNN. One of the most common attacks is the fast gradient sign method (FGSM) [8]. Given a vectorized input \mathbf{x} of the real image and corresponding label \mathbf{t} , FGSM perturbs each element x in \mathbf{x} along the sign of the associated element of the gradient of the inference loss w.r.t. \mathbf{x} as shown in Eq. 1 and illustrated in Fig. 1(b). Another common attack is the projected gradient descent (PGD) [23]. The PGD attack is a multi-step variant of FGSM where $\hat{\mathbf{x}}^{k=1} = \mathbf{x}$ and the iterative update of the perturbed data $\hat{\mathbf{x}}$ in k^{th} step is given in Eq. 2.

$$\hat{\mathbf{x}} = \mathbf{x} + \epsilon \times \text{sgn}(\nabla_{\mathbf{x}} J(g(\mathbf{x}; \boldsymbol{\theta}), \mathbf{t})) \quad (1)$$

$$\hat{\mathbf{x}}^k = \text{Proj}_{P_\epsilon(\mathbf{x})}(\hat{\mathbf{x}}^{k-1} + \alpha \times \text{sgn}(\nabla_{\mathbf{x}} J(g(\hat{\mathbf{x}}^{k-1}; \boldsymbol{\theta}), \mathbf{t}))) \quad (2)$$

Here, the scalar ϵ corresponds to the perturbation constraint that determines the severity of the perturbation. $g(\mathbf{x}; \boldsymbol{\theta})$ generates the output of the DNN, parameterized by $\boldsymbol{\theta}$. Here, Proj projects the updated adversarial sample onto the projection space $P_\epsilon(\mathbf{x})$ which is the ϵ - L_∞ neighbourhood of the benign sample² \mathbf{x} . α is the attack step size.

Note that these two strategies assume the attacker knows the details of the DNN and are thus termed as *white-box attacks*. We will evaluate the merit of our training scheme by measuring the robustness of our trained models to the fake images generated by these attacks. We argue that this evaluation is more comprehensive than using images generated by attacks that assume limited knowledge of the DNN [28]. Moreover, we note that PGD is one of the strongest L_∞ adversarial example generation algorithms [23] and use it as part of our proposed framework.

2.2 Model Compression

ADMM is a powerful optimization method used to solve problems with non-convex, combinatorial constraints [1]. It decomposes the original optimization problem into two sub-problems and solves the sub-problems iteratively until convergence. Pruning convolutional neural networks (CNNs) can be modeled as an optimization problem where the cardinality of each layer’s weight tensor is bounded by its pre-specified pruning ratio. In the ADMM framework, such constraints are transformed to ones represented with indicator functions, such as $\mathbb{I}_{\theta}(\mathbf{z}) = 0$ for $|\mathbf{z}| \leq n$ and $+\infty$ otherwise. Here, \mathbf{z} denotes the duplicate variable [1] and n represents the target number of non-zero weights determined by pre-specified pruning ratios. Next, the original optimization problem is reformulated as:

$$\mathcal{L}_\rho(\boldsymbol{\theta}, \mathbf{z}, \lambda) = J(g(\mathbf{x}; \boldsymbol{\theta}), \mathbf{t}) + \mathbb{I}_{\theta}(\mathbf{z}) + \langle \lambda, \boldsymbol{\theta} - \mathbf{z} \rangle + \frac{\rho}{2} \|\boldsymbol{\theta} - \mathbf{z}\|_2^2 \quad (3)$$

²It is noteworthy that the generated $\hat{\mathbf{x}}$ are clipped to a valid range which for our experiments is $[0, 1]$.

where λ is the Lagrangian multiplier and ρ is the penalization factor when parameters θ and z differ. Eq. (3) is broken into two sub-problems which solve θ and z iteratively until convergence [27]. The first sub-problem uses stochastic gradient descent (SGD) to update θ while the second sub-problem applies projection to find the assignment of z that is closest to θ yet satisfies the cardinality constraint, effectively pruning weights with small magnitudes.

Not only can ADMM prune a model's weight tensors but it also has as a dynamic regularizer. Such adaptive regularization is one of the main reasons behind the success of its use in pruning. However, ADMM-based pruning has several drawbacks. First, ADMM requires prior knowledge of the per-layer pruning ratios. Second, ADMM does not guarantee the pruning ratio will be met, and therefore, an additional round of hard pruning is required after ADMM completes. Third, not all problems solved with ADMM are guaranteed to converge. Fourth, to improve the convergence, ρ needs to be progressively increased across several rounds of training, which increases training time [1].

Sparse learning [4] addresses the shortcomings of ADMM by leveraging exponentially smoothed gradients (momentum) to prune weights. It redistributes pruned weights across layers according to their mean momentum contribution. The weights that will be removed and transferred to other layers are chosen according to their magnitudes while the weights that are brought back (reactivated) are selected based on their momentum values. On the other hand, a major shortcoming of sparse learning compared to ADMM is that it does not benefit from a dynamic regularizer and thus often yields lower levels of accuracy. Furthermore, existing sparse-learning schemes only support irregular forms of pruning, limiting speed-up on many compute platforms. Finally, sparse-learning, to the best of our knowledge, has not previously been extended to robust model compression.

3 DYNAMIC NETWORK REWIRING

To tackle the shortcomings of ADMM and sparse-learning this section introduces a dynamic L_2 regularizer that enables non-iterative training to achieve high accuracy with compressed models. We then describe a hybrid loss function to provide robustness to the compressed models and an extension to support structured pruning.

3.1 Dynamic L_2 Regularizer

For a DNN parameterized by θ with L layers, we let θ_l represent the weight tensor of layer l . In our sparse-learning approach, these weight tensors are element-wise multiplied (\odot) by corresponding binary mask tensors (\mathbf{m}_l) to retain only a fraction of non-zero weights, thereby meeting a target pruning ratio. We update each layer mask in every epoch similar to [4]. The number of non-zeros is updated based on the layer's normalized momentum and the specific non-zero entries are set to favor large magnitude weights. We incorporate an ADMM dynamic L_2 regularizer [27] into this framework by introducing duplicate variable z for the non-zero weights, which is in turn updated at the start of every epoch. Unlike [27], we only penalize differences between the masked weights ($\theta_l \odot \mathbf{m}_l$) of a layer l and their corresponding duplicate variable \mathbf{z}_l . Because the total cardinality constraint of the masked parameters is satisfied, i.e. $\sum_{l=1}^L \text{card}(\theta_l \odot \mathbf{m}_l) \leq n$, the indicator penalty factor

is redundant and the loss function may be simplified as

$$\mathcal{L}_\rho(\theta, \mathbf{z}, \mathbf{m}) = J(g(\mathbf{x}; \theta, \mathbf{m}), t) + \frac{\rho}{2} \sum_{l=1}^L \|\theta_l \odot \mathbf{m}_l - \mathbf{z}_l\|_2^2 \quad (4)$$

where, ρ is the dynamic L_2 penalizing factor. This simplification is particularly important because the indicator function used in Eq. 3 is non-differentiable and its removal in Eq. 4 enables the loss function to be minimized without decomposition into two sub-problems.³ Moreover, SGD with this loss function converges similarly to the SGD with $J(g(\mathbf{x}; \theta, \mathbf{m}), t)$ and more reliably than ADMM. Intuitively, the key role of the dynamic regularizer in this simplified loss function is to encourage the DNN to not change values of the weights that have large magnitude unless the corresponding loss is large, similar to what the dynamic regularizer does in ADMM-based pruning.

3.2 Proposed Hybrid Loss Function

For a given input image \mathbf{x} , adversarial training can be viewed as a min-max optimization problem that finds the model parameters θ that minimize the loss associated with the corresponding adversarial sample $\hat{\mathbf{x}}$, as shown below:

$$\arg \min_{\theta} \{ \arg \max_{\hat{\mathbf{x}} \in P_\epsilon(\mathbf{x})} J(g(\hat{\mathbf{x}}; \theta), t) \} \quad (5)$$

In our framework, we use SGD for loss minimization and PGD to generate adversarial images. More specifically, to boost classification robustness on perturbed data we propose using a hybrid loss function that combines the proposed simplified loss function in Eq. 4 with adversarial image loss, i.e.,

$$J_{\text{tot}} = \beta \mathcal{L}_\rho(\theta, \mathbf{z}, \mathbf{m}) + (1 - \beta) J(g(\hat{\mathbf{x}}; \theta, \mathbf{m}), t) \quad (6)$$

β provides a tunable trade-off between the two loss components.

Observation 1 A DNN only having a fraction of weights active throughout the training can be trained with the proposed hybrid loss to finally converge similar to that of the un-pruned model (mask $\mathbf{m} = 1$) to provide a robust yet compressed model.

This is exemplified in Fig. 2(a) which shows similar convergence trends for both pruned and unpruned models, simultaneously achieving both the target compression and robustness while also mitigating the requirement of multiple training iterations.

3.3 Support for Channel Pruning

Let the weight tensor of a convolutional layer l be denoted as $\theta_l \in \mathbb{R}^{M \times N \times h \times w}$, where h and w are the height and width of the convolutional kernel, and M and N represent the number of filters and channels per filter, respectively. We convert this tensor to a 2D weight matrix, with M and $N \times h \times w$ being the number of rows and columns, respectively. We then partition this matrix into N sub-matrices of M rows and $h \times w$ columns, one for each channel. To compute the importance of a channel c , we find the Frobenius norm (F-norm) of corresponding sub-matrix, thus effectively compute $O_l^c = \|\theta_l^{c, \dots}\|_F^2$. Based on the fraction of non-zero weights that need to be rewired during an epoch i , denoted by the pruning rate p_i , we compute the number of channels that must be pruned from each layer, $c_l^{p_i}$, and prune the $c_l^{p_i}$ channels with the

³Note this simplified loss function also drops the term $\langle \lambda, \theta - \mathbf{z} \rangle$ because \mathbf{z} is updated with θ at the beginning of each epoch, forcing the Lagrangian multiplier λ and its contribution to the loss function to be always 0.

Algorithm 1: DNR Training.

Data: weight θ_l , momentum μ_l , binary mask $M_l, l = 0..k$
Data: density $d, i = 0..numEpochs$, pruning rate $p = p_{i=0}$
 pT: irregular or channel

- 1 **for** $l \leftarrow 0$ **to** k **do**
- 2 $\theta_l \leftarrow \text{init}(\theta_l)$ & $M_l \leftarrow \text{createMaskForWeight}(\theta_l, d)$
- 3 applyMaskToWeights(θ_l, M_l)
- 4 $z_l \leftarrow \theta_l \odot M_l$
- 5 **end**
- 6 **for** $i \leftarrow 0$ **to** $numEpochs$ **do**
- 7 **for** $j \leftarrow 0$ **to** $numBatches$ **do**
- 8 $J = \text{computeCleanLoss}(x) + \text{updateDynamicRegularizer}(\theta, z)$
- 9 $J_{adv} = \text{computePerturbedLoss}(\hat{x})$
- 10 $J_{tot} = \text{updateRobustLoss}(J, J_{adv})$
- 11 $\frac{\partial J_{tot}}{\partial \theta} = \text{computeGradients}(\theta, \text{batch})$
- 12 updateMomentumAndWeights($\frac{\partial J_{tot}}{\partial \theta}, \mu$)
- 13 **for** $l \leftarrow 0$ **to** k **do**
- 14 applyMaskToWeights(θ_l, M_l)
- 15 **end**
- 16 **end**
- 17 $tM \leftarrow \text{getTotalMomentum}(\mu)$
- 18 $pT \leftarrow \text{getTotalPrunedWeights}(\theta, p_i)$
- 19 $p_i \leftarrow \text{linearDecay}(p_i)$
- 20 **for** $l \leftarrow 0$ **to** k **do**
- 21 $\mu_l \leftarrow \text{getMomentumContribution}(\theta_l, M_l, tM, pT)$
- 22 Prune(θ_l, M_l, p_i, pT)
- 23 Regrow($\theta_l, M_l, \mu_l \cdot tM, pT$)
- 24 applyMaskToWeights(θ_l, M_l)
- 25 $z_l \leftarrow \theta_l \odot M_l$
- 26 **end**
- 27 **end**

lowest F-norms. We then compute each layer's importance based on the normalized momentum contributed by its non-zero channels. These importance measures are used to determine the number of zero-F-norm channels $r_l^i \geq 0$ that should be re-grown for each layer l . More precisely, we re-grow the r_l^i zero-F-norm channels with the highest Frobenius norms of their momentum. We note that this approach can easily be extended to enable various other forms of structured pruning. Moreover, despite supporting pruning of both convolution and linear layers, this paper focuses on reducing the computational complexity of a DNN. We thus experiment with pruning only convolutional layers because they dominate the computational complexity [18]. The detailed pseudo-code of the proposed training framework is shown in Algorithm 1.

It is noteworthy that DNR's ability to arrange per-layer pruning ratio for robust compression successfully avoids the tedious task of hand-tuning the pruning-ratio based on *layer sensitivity*. To illustrate this, we follow [5] to quantify the sensitivity of a layer by measuring the percentage reduction in classification accuracy on both clean and adversarial images caused by pruning that layer by $x\%$ without pruning other layers.

Observation 2 DNN layers' sensitivity towards clean and perturbed images are not necessarily equal, thus determining layer pruning ratios for robust models is particularly challenging.

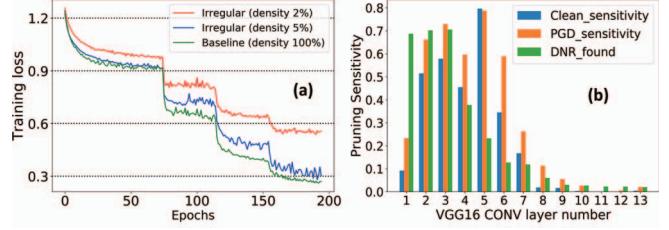


Figure 2: (a) Training loss vs. epochs and (b) Pruning sensitivity per layer for VGG16 on CIFAR-10.

As exemplified in Fig. 2(b), for $x = 95\%$ there is significant difference in the sensitivity of the layers for clean and perturbed image classification. DNR, on the contrary, automatically finds per-layer pruning ratios (overlaid as pruning sensitivity as in [5]) that serves well for both types of image classification targeting a global pruning of 95%.

4 EXPERIMENTS

In this section, we first describe the experimental setup we used to evaluate the effectiveness of the proposed robust training scheme. We then compare our method against other state-of-the-art robust pruning techniques based on ADMM [31] and L_1 lasso [26]. We also evaluate the merit of DNR as a clean-image pruning scheme and show that it consistently outperforms contemporary non-iterative model pruning techniques [4, 5, 12, 20]. We finally present an ablation study to empirically evaluate the importance of the dynamic regularizer in the DNR's loss function. We used Pytorch [25] to write the models and trained/tested on AWS P3.2x large instances that have an NVIDIA Tesla V100 GPU.

4.1 Experimental Setup

4.1.1 Models and Datasets. We selected three widely used datasets, CIFAR-10 [16] CIFAR-100 [16] and Tiny-ImageNet [10] and picked two well known CNN models, VGG16 [29] and ResNet18 [11]. Both CIFAR-10 and CIFAR-100 datasets have 50K training samples and 10K test samples with an input image size of $32 \times 32 \times 3$. Training and test data size for Tiny-ImageNet are 100k and 10k, respectively where each image size is of $64 \times 64 \times 3$. For all the datasets we used standard data augmentations (horizontal flip and random crop with reflective padding) to train the models with a batch size of 128.

4.1.2 Adversarial Attack and DNR Training Settings. For PGD, we set ϵ to $8/255$, the attack step size $\alpha = 0.01$, and the number of attack iterations to 7, the same values as in [14]. For FGSM, we choose the same ϵ value as above.

We performed DNR based training for 200/170/60 epochs for CIFAR-10/CIFAR-100/Tiny-ImageNet, with a starting learning rate of 0.1, momentum value of 0.9, and weight decay value of $5e^{-4}$. For CIFAR-10 and CIFAR-100 the learning rate (LR) was reduced by a factor of 0.2 after 80, 120, and 160 epochs. For Tiny-ImageNet we reduced the LR value after 30 and 50 epochs. In addition, we hand-tuned ρ to 10^{-4} and set the pruning rate $p = 0.5$. We linearly decreased the pruning rate every epoch by $\frac{p}{\text{total epochs}}$. Finally, to balance between the clean and adversarial loss, we set β to 0.5. Lastly, note that we performed warm-up sparse learning [4] for the

Pruning type	Compression -ratio	% Channel present	Accuracy (%)		
			Clean	FGSM	PGD
Unpruned-baseline	1×	100	50.91	18.19	13.87
Irregular Channel	20.63×	98.52	51.71	18.21	14.46
	1.45×	74	51.09	17.92	13.54

Table 1: Results on VGG16 to classify Tiny-ImageNet.

first 5 epochs with only the clean image loss function before using the hybrid loss function with dynamic regularization (see Eq. 6) for robust compression for the remaining epochs.

4.2 Results

Results on CIFAR datasets: We analyzed the impact of our robust training framework on both clean and adversarially generated images with various target compression ratios in the range [0.01, 1.0], where model compression is computed as the ratio of total weights of the model to the non-zero weights in the pruned model. As shown in Figs. 3(a-b) DNR can effectively find a robust model with high compression and negligible compromise in accuracy. In particular, for irregular pruning our method can compress up to $\sim 20\times$ with negligible drop in accuracy on clean as well as PGD and FGSM based perturbed images, compared to the baseline non-pruned models, tested with VGG16 on CIFAR-10 and ResNet18 on CIFAR-100.⁴

Observation 3 *As the target compression ratio increases, channel pruning degrades adversarial robustness more significantly than irregular pruning.*

As we can see in Fig. 3(a-b), the achievable model compression with negligible accuracy loss for structured (channel) pruned models is $\sim 10\times$ lower than that achievable through irregular pruning. This trend matches with that of the model’s performance on clean image. However, as we can see in Fig. 3(c), the percentage of channels present in our channel-pruned models can be up to $\sim 10\times$ lower than its irregular counterparts, implying a similarly large speedup in inference time on a large range of compute platforms [4].

Results on Tiny-ImageNet: As shown in Table 1, DNR can compress the model up to $20.63\times$ without any compromise in performance for both clean and perturbed image classification.

It is also noteworthy that all our accuracy results for both clean and adversarial images correspond to models that provide the best test accuracy on clean images. This is because robustness gains are typically more relevant on models in which the performance on clean images is least affected.

4.3 Comparison with State-of-the-art

Here, we compare the performance of DNR with ADMM [31] and L_1 lasso based [26] robust pruning. For ADMM based robust pruning we followed a three stage compression technique namely *pre-training*, *ADMM based pruning*, and *masked retraining*, performing pruning for 30 epochs with $\rho_{admm} = 10^{-3}$ as described in [31]. L_1 lasso based pruning adds a L_1 regularizer to its loss function to penalize the weight magnitudes, where the regularizer coefficient determines the penalty factor. Table 2 shows that our proposed method outperforms both ADMM and L_1 Lasso based approaches by a considerable margin, retaining advantages of both worlds⁵. In particular, compared to ADMM, with VGG16 (ResNet18) model

⁴A similar trend is observed for VGG16 on CIFAR-100 and ResNet18 on CIFAR-10. These are not included in the paper due to space limitations.

Model	Method	No pre-trained model	Per-layer sparsity knowledge not-needed	Target pruning met	Pruning type	Compression ratio	Accuracy (%)		
							Clean	FGSM	PGD
VGG16	ADMM [31]	✗	✗	✓	Irregular	16.78×	86.34	49.52	40.62
	ADMM naive	✗	✓	✓	Irregular	19.74×	83.87	42.46	32.87
	L_1 Lasso [26]	✓	✓	✗	Irregular	2.01×	83.24	50.32	42.01
	DNR	✓	✓	✓	Irregular	20.85×	86.74	52.92	43.21
ResNet18	ADMM [31]	✗	✗	✓	Irregular	14.6×	87.15	54.65	46.57
	ADMM naive	✗	✓	✓	Irregular	19.74×	86.10	50.49	42.24
	L_1 Lasso [26]	✓	✓	✗	Irregular	6.84×	85.92	55.20	46.80
	DNR	✓	✓	✓	Irregular	21.57×	87.32	55.13	47.35

Table 2: Comparison of DNR, ADMM based, and L_1 lasso based robust pruning schemes on CIFAR-10.

on CIFAR-10, DNR provides up to 3.4% (0.78%) increased classification accuracy on perturbed images with $1.24\times$ (1.48×) higher compression. Compared to L_1 Lasso, we achieve $10.38\times$ (3.15×) higher compression and up to 2.6% (0.55%), and 3.5% (1.4%) increased accuracy on perturbed and clean images, respectively, for VGG16 (ResNet18) on CIFAR-10 classification.

Observation 4 *Naively tuned per-layer pruning ratio degrades both robustness and clean-image classification performance of a model.*

For this, we evaluated robust compression using naive ADMM, i.e. using naively tuned per-layer pruning-ratio (all but the 1st layer $\sim x\%$ for a $x\%$ total sparsity). As shown in Table 2, this clearly degrades the performance, implying layer-sparsity tuning is necessary for ADMM to perform well.

4.4 Pruning to Classify Clean-only Images

To evaluate the merit of DNR as a clean-image only pruning scheme (DNR-C), we trained using DNR with the same loss function minus the adversarial loss term (by setting $\beta = 1.0$ in Eq. 6) to reach a target pruning ratio. Table 3 shows that our approach consistently outperforms other state-of-the-art non-iterative pruning approaches based on momentum information [4, 5], reinforcement-learning driven auto-compression (AMC) [12], and connection-sensitivity [20]⁵. The δ value in the seventh column represents the error difference from corresponding non-pruned baseline models. We also present performance on CIFAR-100 for VGG16 and ResNet18 and Tiny-ImageNet for VGG16.⁶ In particular, we can achieve up to $34.57\times$ (12.61×) compression on CIFAR-10 dataset with irregular (channel) pruning maintaining accuracy similar to the baseline. On CIFAR-100 compression of up to $22.45\times$ (5.57×) yields no significant accuracy drop (less than 2.7% in top-1 accuracy) with irregular (channel) pruning. Moreover, our evaluation shows a possible practical speed up of up to $6.06\times$ for CIFAR-10 and $2.41\times$ for CIFAR-100 can be achieved through channel pruning using DNR-C. For Tiny-ImageNet, DNR-C can provide compression and speed-up of up to $11.55\times$ and $1.53\times$, respectively with negligible accuracy drop.

4.5 Ablation Study

To understand the performance of the proposed hybrid loss function with a dynamic L_2 regularizer, we performed ablation with both VGG16 and ResNet18 on CIFAR-10 for a target parameter density of 5% and 50% using irregular and channel pruning, respectively. As

⁵Romanized numbers in the table are results of our experiments, and italicized values are directly taken from the respective original papers.

⁶To have an "apple to apple" comparison we provide results on ResNet50 model for classification on CIFAR-10. All other simulations are done on only the ResNet18 variant of ResNet.

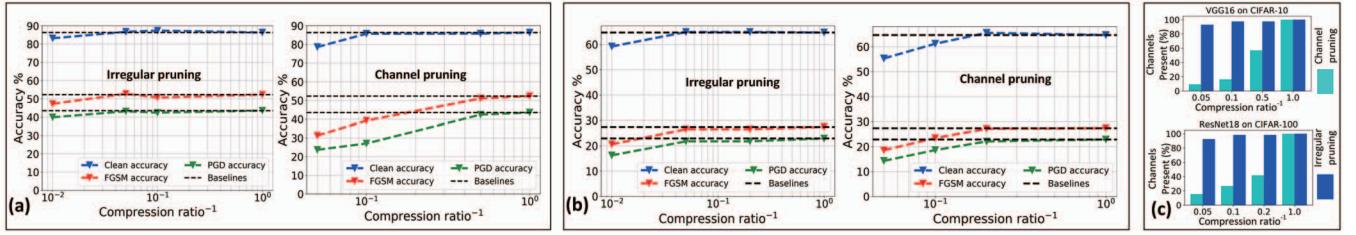


Figure 3: Model compression vs. accuracy (on both clean and adversarially generated images) for irregular and channel pruning evaluated with (a) VGG16 on CIFAR-10 and (b) ResNet18 on CIFAR-100. (c) Comparison of channel pruning with irregular pruning in terms of % of channels present. Note that the % of channels present correlates with inference time [4, 22].

Dataset	Model	Method	Pruning type	Compression ratio	Error (%) top-1	δ from baseline	Speedup
CIFAR-10	VGG16	SNIP [20]	Irregular	32.33X	8.00	-0.26	–
		Sparse-learning [4]	Irregular	32.33X	7.00	-0.5	–
		DNR-C	Irregular	34.57X	6.50	-0.09	1.29X
	ResNet50	DNR-C	Channel	12.61X	8.00	-1.5	6.06X
		GSM [5]	Irregular	10X	6.20	-0.25	–
		AMC [12]	Irregular	2.5X	6.45	+0.02	–
CIFAR-100	ResNet18	DNR-C	Irregular	20X	4.8	-0.07	1.75X
		DNR-C	Irregular	20.32X	5.19	-0.10	1.31X
	VGG16	DNR-C	Channel	5.67X	5.36	-0.27	2.43X
		DNR-C	Channel	2.76X	28.78	-2.68	2.06X
Tiny ImageNet	ResNet18	DNR-C	Irregular	22.45X	24.9	-1.17	1.13X
		DNR-C	Channel	5.57X	25.28	-1.55	2.41X

Table 3: Comparison with state-of-the-art non-iterative pruning schemes on CIFAR-10 and comparison of deviation from baseline on CIFAR-100 and Tiny-ImageNet.

Model	Method: DNR	Accuracy (%) with irregular pruning			Accuracy (%) with channel pruning		
		Clean	FGSM	PGD	Clean	FGSM	PGD
VGG16	Without dynamic L_2	87.01	50.09	40.62	86.28	49.49	41.25
	With dynamic L_2	86.74	52.92	43.21	85.83	51.03	42.36
ResNet18	Without dynamic L_2	87.45	53.52	45.33	87.97	53.10	45.91
	With dynamic L_2	87.32	55.13	47.35	87.49	56.09	48.33

Table 4: Comparison of DNR with and without the dynamic regularizer for CIFAR-10 classification.

shown in Table 4, using the dynamic regularizer improves the adversarial classification accuracy by up to 2.83% for VGG16 and $\sim 3\%$ for ResNet18 with similar clean-image classification performance.

4.6 Generalized Robustness Against PGD Attack of Different Strengths

Fig. 4 presents the performance of the pruned models as a function of the PGD attack iteration and the attack bound ϵ . In particular, we can see that, for both irregular and channel pruned models, the accuracy degrades with higher number of attack iterations. When ϵ increases, the accuracy drop is similar in both the pruning schemes. These trends suggest that our robustness is not achieved through gradient obfuscation [26].

5 CONCLUSIONS

This paper addresses the open problem of achieving ultra-high compression of DNN models while maintaining their robustness

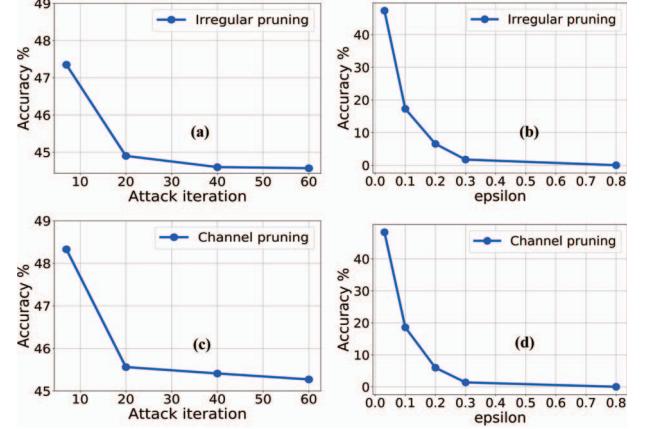


Figure 4: On CIFAR-10, the perturbed data accuracy of ResNet18 under PGD attack versus increasing (a), (c) attack iteration and (b), (d) attack bound ϵ for irregular (5% density), and channel pruned (50% density) models, respectively.

through a non-iterative training approach. In particular, the proposed DNR method leverages a novel sparse-learning strategy with a hybrid loss function that has a dynamic regularizer to achieve better trade-offs between accuracy, model size, and robustness. Furthermore, our extension to support channel pruning shows that compressed models produced by DNR can have a practical inference speed-up of up to $\sim 10\times$.

REFERENCES

- [1] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning* 3, 1 (2011), 1–122.
- [2] Nicholas Carlini and David Wagner. 2017. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy*. IEEE, 39–57.
- [3] Chenyi Chen, Ari Seff, Alain Kornhauser, and Jianxiong Xiao. 2015. Deepdriving: Learning affordance for direct perception in autonomous driving. In *Proceedings of the IEEE International Conference on Computer Vision*. 2722–2730.
- [4] Tim Dettmers and Luke Zettlemoyer. 2019. Sparse networks from scratch: Faster training without losing performance. *arXiv preprint arXiv:1907.04840* (2019).
- [5] Xiaohan Ding, Xiangxin Zhou, Yuchen Guo, Jungong Han, Ji Liu, et al. 2019. Global Sparse Momentum SGD for Pruning Very Deep Neural Networks. In *Advances in Neural Information Processing Systems*. 6379–6391.
- [6] Thu Dinh and Jack Xin. 2018. 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).

[7] Arash Fayyazi, Souvik Kundu, Shahin Nazarian, Peter A Beerel, and Massoud Pedram. 2019. CSram: Area-efficient low-power ex-situ training framework for memristive neuromorphic circuits based on clustered sparsity. In *2019 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*. IEEE, 465–470.

[8] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014).

[9] Shupeng Gui, Haotao N Wang, Haichuan Yang, Chen Yu, Zhangyang Wang, and Ji Liu. 2019. Model Compression with Adversarial Robustness: A Unified Optimization Framework. In *Advances in Neural Information Processing Systems*. 1283–1294.

[10] Lucas Hansen. 2015. Tiny ImageNet challenge submission. *CS 231N* (2015).

[11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.

[12] Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. 2018. AMC: AutoML for Model Compression and Acceleration on Mobile Devices. In *The European Conference on Computer Vision (ECCV)*.

[13] Yihui He, Xiangyu Zhang, and Jian Sun. 2017. Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*. 1389–1397.

[14] Zhezhi He, Adnan Siraj Rakin, and Deliang Fan. 2019. Parametric noise injection: Trainable randomness to improve deep neural network robustness against adversarial attack. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 588–597.

[15] Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, et al. 2012. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal processing magazine* 29, 6 (2012), 82–97.

[16] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images. *Technical report, Citeseer* (2009).

[17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. 1097–1105.

[18] Souvik Kundu, Mahdi Nazemi, Massoud Pedram, Keith M Chugg, and Peter A Beerel. 2020. Pre-defined Sparsity for Low-Complexity Convolutional Neural Networks. *IEEE Trans. Comput.* (2020).

[19] Souvik Kundu, Saurav Prakash, Haleh Akrami, Peter A. Beerel, and Keith M. Chugg. 2019. pSConv: A Pre-defined sparse Kernel Based Convolution for Deep CNNs. In *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 100–107.

[20] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip HS Torr. 2018. SNIP: Single-shot network pruning based on connection sensitivity. *arXiv preprint arXiv:1810.02340* (2018).

[21] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen Awm Van Der Laak, Bram Van Ginneken, and Clara I Sánchez. 2017. A survey on deep learning in medical image analysis. *Medical image analysis* 42 (2017), 60–88.

[22] Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. 2018. Rethinking the value of network pruning. *arXiv preprint arXiv:1810.05270* (2018).

[23] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2017. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083* (2017).

[24] Dongyu Meng and Hao Chen. 2017. Magnet: a two-pronged defense against adversarial examples. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 135–147.

[25] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. (2017).

[26] Adnan Siraj Rakin, Zhezhi He, Li Yang, Yanzhi Wang, Liqiang Wang, and Deliang Fan. 2019. Robust Sparse Regularization: Simultaneously Optimizing Neural Network Robustness and Compactness. *arXiv preprint arXiv:1905.13074* (2019).

[27] Ao Ren, Tianyun Zhang, Shaokai Ye, Jiayu Li, Wenya Xu, Xuehai Qian, Xue Lin, and Yanzhi Wang. 2019. ADMM-NN: An algorithm-hardware co-design framework of DNNs using alternating direction methods of multipliers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. 925–938.

[28] Kui Ren, Tianhang Zheng, Zhan Qin, and Xue Liu. 2020. Adversarial Attacks and Defenses in Deep Learning. *Engineering* (2020).

[29] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).

[30] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. 2017. Ensemble adversarial training: Attacks and defenses. *arXiv preprint arXiv:1705.07204* (2017).

[31] Shaokai Ye, Kaidi Xu, Sijia Liu, Hao Cheng, Jan-Henrik Lambrechts, Huan Zhang, Aojun Zhou, Kaisheng Ma, Yanzhi Wang, and Xue Lin. 2019. Adversarial robustness vs. model compression, or both. In *The IEEE International Conference on Computer Vision (ICCV)*, Vol. 2.