
TO COMPRESS OR NOT TO COMPRESS: UNDERSTANDING THE INTERACTIONS BETWEEN ADVERSARIAL ATTACKS AND NEURAL NETWORK COMPRESSION

Yiren Zhao^{*1} Ilya Shumailov^{*1} Robert Mullins¹ Ross Anderson¹

ABSTRACT

As deep neural networks (DNNs) become widely used, pruned and quantised models are becoming ubiquitous on edge devices; such compressed DNNs lower the computational requirements. Meanwhile, multiple recent studies show ways of constructing adversarial samples that make DNNs misclassify. We therefore investigate the extent to which adversarial samples are transferable between uncompressed and compressed DNNs. We find that such samples remain transferable for both pruned and quantised models. For pruning, adversarial samples at high sparsities are marginally less transferable. For quantisation, we find the transferability of adversarial samples is highly sensitive to integer precision.

1 INTRODUCTION

Deep Neural Networks (DNNs) perform well on a wide range of tasks, including image classification (Krizhevsky et al., 2012), object detection (Ren et al., 2015), reading comprehension (Seo et al., 2016) and machine translation (Bahdanau et al., 2015). They have proved to be an efficient method of harvesting information from large amounts of data and are expected to be ubiquitous in the future. Despite these successes, two questions remain crucial for deploying them in embedded systems. First, their substantial computational and memory requirements can make deployment challenging on power-limited devices. Second, as they start to appear in safety-critical applications, their reliability and security become a serious issue.

In order to compute DNNs efficiently on embedded systems, researchers have proposed various compression methods. *Pruning* directly reduces the number of parameters of DNNs – this reduction translates to fewer data movements and thus saves energy directly. *Quantisation* is another popular compression technique – it simultaneously reduces the memory footprint and decreases the energy cost of multiplications. Both compression methods are widely deployed on DNN accelerators. For instance, Efficient Inference Engines (EIE) use pruning, quantisation and encoding techniques for energy efficiency (Han et al., 2016a); the Sparse CNN (SCNN) accelerator first requires network parameters to be pruned

and encoded, then performs computations directly using the encoded data format (Parashar et al., 2017). It seems likely that pruning, quantisation and other compression techniques will be used for future DNNs on embedded devices, even where hardware accelerators are also used.

Over the past five years, research has found DNNs to be sensitive to small perturbations of the input images, with the result that they can often be fooled easily using specially-crafted adversarial inputs (Szegedy et al., 2013). Such adversarial samples are a real concern for safety-critical systems; attackers might try to manipulate autonomous vehicles (Eykholt et al., 2018) or break into smart phones by tricking the speaker recognition system (Carlini et al., 2016).

In this paper, we study the portability of adversarial samples. Might an attacker learn how to break into widely-deployed low-cost systems and then use the same adversarial samples as a springboard to attack other related systems?

We make the following contributions in this paper.

- We investigated the effects of different DNN compression mechanisms on adversarial attacks.
- We have developed the first compression-aware machine learning attack taxonomy and used it to evaluate the transferability of adversarial samples between compressed and uncompressed models.
- As for pruning, we found that adversarial samples are transferable between compressed and uncompressed models. However, adversarial samples generated from uncompressed models are marginally less effective on compressed models at preferred sparsities, and adver-

^{*}Equal contribution ¹Computer Laboratory, University of Cambridge, UK. Correspondence to: Yiren Zhao <Yiren.Zhao@cl.cam.ac.uk>.

arial samples generated from extremely sparse models are no longer effective on the baseline model.

- As for quantisation, we found that adversarial samples are transferable between compressed and uncompressed models. However, a reduction in integer precision provides clipping effects and marginally limits transferability in fast gradient based attacks.

2 RELATED WORK

2.1 Pruning

Pruning directly reduces the number of parameters in a DNN model, and thus the number of off-chip to on-chip data transfers on modern DNN accelerators (Chen et al., 2016). If the architecture allows, pruning may also reduce the computation cost (Kim et al., 2018). Consider a weight tensor (\mathbf{W}_n); fine-grained pruning is simply performing an element-wise multiplication (\odot) between a mask operator \mathbf{M}_n and the original weight tensor (\mathbf{W}_n).

$$\mathbf{W}_n' = \mathbf{W}_n \odot \mathbf{M}_n \tag{1}$$

Han et al. first proposed pruning a DNN by applying a threshold to the DNN’s parameters (Han et al., 2016b). In this case, the mask (\mathbf{M}_n) consists of thresholding by a single value α .

$$\mathbf{M}_n = h_k(\mathbf{W}_n^{(i,j)}) = \begin{cases} 0 & \text{if } \alpha > |\mathbf{W}_k^{(i,j)}| \\ 1 & \text{otherwise} \end{cases} \tag{2}$$

Using this simple one-shot pruning technique, Han et al. were able to reduce the number of parameters in AlexNet by 9x and VGG16 by 13x (Han et al., 2016b). In their implementation, the masking and fine-tuning happen iteratively but the masked values are not allowed to recover in later stages.

Guo et al. subsequently proposed dynamic network surgery (DNS), which allows pruned parameters to recover at later stages (Guo et al., 2016). The approach is to condition the mask using the following equation, where α and β are two constants.

$$\mathbf{M}_n = h_k(\mathbf{W}_n^{(i,j)}) = \begin{cases} 0 & \text{if } \alpha > |\mathbf{W}_k^{(i,j)}| \\ \mathbf{M}_n^{(i,j)} & \text{if } \alpha \leq |\mathbf{W}_k^{(i,j)}| \leq \beta \\ 1 & \text{otherwise} \end{cases} \tag{3}$$

Values that become bigger at later stages are allowed to re-join the fine-tuning process. Guo et al. demonstrated higher compression rates on a large range of networks compared

to Han et al. In this paper, we generate pruned DNNs using the DNS method.

2.2 Quantisation

Quantisation refers to using fewer bits for parameters in a DNN than the standard 32-bit single-precision floating-point representation used on modern CPUs and GPUs. Hubara et al. showed that low-precision fixed-point numbers can be used for neural-network inference with nearly no loss of accuracy (Hubara et al., 2017). In the extreme case, the parameters of a DNN can be quantised to either binary or ternary values (Courbariaux et al., 2016; Li et al., 2016). Such aggressive quantisation can greatly speed up DNN hardware accelerators but suffers from significant loss of accuracy. For resource-constrained devices, a low-precision fixed-point representation can give a balance between accuracy and performance (Lin et al., 2016). The narrower bitwidth means direct reductions in memory requirement and fixed-point multiplications are less computationally expensive compared with standard single-precision floating point. In this paper, we generate models that use fixed-point parameters at various levels of precision.

2.3 Adversarial Attacks

Szegedy et al. discovered that, despite generalising well, models trained on huge datasets are all vulnerable to adversarial samples (Szegedy et al., 2013). Misclassification can even happen with imperceptible perturbations of the data samples. All the samples they used were within the expected data distribution and only a small specially-crafted amount of noise was added. They observed that models of different configurations, trained on different datasets, misclassify the same samples. Finally, they noted that training a model on adversarial samples helps make it more robust against them. However, this defence is not always practical; their approach based on L-BFGS requires an expensive constrained optimisation with multiple iterations.

A follow-up paper by Goodfellow et al. (Goodfellow et al., 2015) explored the underlying reasons for the existence and generalisability of adversarial samples. They argue that such samples are an artefact of high-dimensional dot-products, and attacks are generalisable because different models learn similar functions when trained to perform the same task. Additionally, they presented two methods to generate adversarial samples in a white-box setting, the fast gradient method (FGM) and the fast gradient sign method (FGSM). Finally, they discovered that RBF-based networks are much more resistant to adversarial samples.

Papernot et al. (Papernot et al., 2016b) came up with another way to generate adversarial samples. They use the gradients of a network to construct saliency maps for the input to discover which input values are so sensitive that a

change can drive misclassification. They showed that their method has the flexibility of being used in both supervised and unsupervised settings and is capable of generating samples with a user-given priority on particular properties of the inputs. Finally, they also observed that adversarial attacks become harder when models have been trained with adversarial samples.

There is now a growing corpus of research on the transferability of adversarial samples (Szegedy et al., 2013; Goodfellow et al., 2015; Papernot et al., 2016a; Tramèr et al., 2017). Transferability refers to the ability of an adversarial sample to evade the correct classification on two different classifiers trained to perform a similar task.

Goodfellow et al. (Goodfellow et al., 2015) and Warde-Farley & Goodfellow (Warde-Farley & Goodfellow, 2016) empirically found that adversarial examples usually occur in large, continuous spatial regions. Tramer et al. (Tramèr et al., 2017) found out that each of the models differs in the dimensionality of its subspaces. A higher number of dimensions increases the chance that the subspaces of different models intersect, leading to transferable samples.

Transferable adversarial samples are a real hazard for model deployment, as they are ‘break-once, run-anywhere’: attacks developed on a particular type of classifier can potentially be deployed everywhere. Papernot et al. (Papernot et al., 2017) in particular have shown that an adversary can sometimes perform attacks without any knowledge of a model’s internal parameters – it can be enough to approximate a model with another known model and build adversarial samples against that instead.

3 METHODOLOGY

3.1 Attack Taxonomy and Threat Model

In this paper, we are interested in the interaction between adversarial attacks and model compression, and we investigate three specific attack scenarios. By ‘compressed models’ we will mean models that have pruned or quantised, while a ‘baseline model’ means a pretrained network without any compression that is dense and whose parameters are represented using full-precision (float32) values.

- Scenario 1: Adversarial attacks occur on each individual compressed model, with the adversarial examples generated and applied on the same model.
- Scenario 2: Adversarial samples are generated from the baseline model but applied on each compressed model.
- Scenario 3: Adversarial samples are generated from compressed models but applied on the baseline model.

In the first scenario, adversarial samples are generated from each compressed model. Attackers can access these compressed models fully, and generate adversarial samples for each one individually. This is the case where attackers buy products and figure out how to attack them.

The second scenario makes the assumption that attackers can only access the baseline model to generate adversarial samples, which are then used to attack various compressed models. Attackers are not allowed to fetch any gradients from compressed models. This is the case where firms take publicly-available models and compress them to run more efficiently on edge devices. Attackers can find the public model and craft adversarial samples to attack derived proprietary devices.

The third scenario assumes that only compressed models are visible to attackers, and attackers generate adversarial samples using compressed models to attack the hidden baseline model. In practice, companies now deploy various compressed neural-network models on edge devices that are exposed to end-users. The assumption is the attackers can access these models and create adversarial samples from them to attack the hidden baseline model. This then back leads to the second scenario; the attacker’s knowledge and toolkit can be transferred to other compressed products from the same firm. Figure 1 shows the second and third attack scenarios.

For example, modern anti-virus (AV) software uses DNNs to detect malware behaviour. Some AV modules detect such behaviours offline. When deploying a compressed model in such an application, how likely is it that malware could analyse the compressed model, work out how to evade the full model, and thus defeat the firm’s other AV products? Similarly, if an alarm company deploys a compressed model for intruder detection in consumer-grade CCTV equipment, could an intelligence agency that buys such equipment figure out how to defeat not just that product but government products derived from the same full model? The risk is that just as a new type of software attack such as Heartbleed or Meltdown can cause widespread disruption by requiring thousands of disparate systems to be patched, so portable adverse examples could force upgrades to large numbers of diverse embedded systems.

3.2 Networks and Compression Methods

We use LeNet5 (LeCun et al., 2015) and CifarNet (Zhao et al., 2018) for our experiments on MNIST (LeCun et al., 2010) and CIFAR10 (Krizhevsky et al., 2014) datasets. The LeNet5 model has 431K parameters and classifies MNIST digits with an accuracy of 99.36%. The CifarNet classifier (Zhao et al., 2018) has 1.3M parameters and achieves 85.93% classification accuracy.




Figure 1. Two different attack setups. Attackers only generate adversarial samples based on baseline model (left), or attackers generate adversarial samples based on compressed models (right).

We implemented two types of compression method:

- Fine-grained pruning on weights;
- Fixed-point quantisation of both weights and activations.

We used the Mayo tool to generate pruned and quantised models (Zhao et al., 2018), and fine-tuned these models after pruning and quantisation. For each pruning density or quantised bitwidth, we retrain 350 epochs for LeNet5 and 300 epochs for CifarNet with three scheduled learning rate decays starting from 0.01. For each decay, the learning rate decreases by a factor of 10.

Applying pruning on a pretrained model shrinks the number of parameters and thus the memory footprint of future AI ASICs. We use fixed-point quantisation on both weights and activations of a DNN. Quantising both weights and activations means that computations operate in low-precision fixed-point formats, which cut the time and energy cost both data moves and computations. For fixed-point quantisation, we use a 1-bit integer when bitwidth is 4, a 2-bit integer when bitwidth is 8, and 4-bit integers for the rest of the fixed-point quantisations.

3.3 Adversarial attacks

In the work reported in this paper we used three popular attacks developed in the research community. We now present mathematical definitions of the attacks and comments about their behaviour.

Goodfellow et al. first introduced the fast gradient method (FGM) and fast gradient sign method (FGSM) to develop attacks (Goodfellow et al., 2015). For the definitions we will use the following notation: θ represents the parameters of the model, \mathbf{X} represents the inputs, while y and y_l represents the outputs and labels respectively. We can

then use $J(\theta, \mathbf{X}, y_l)$ to represent the cost function. The original FGM and FGSM perturbations are computed as in Equation (4) and Equation (5) respectively, where ϵ is a hyperparameter and the function $\nabla_{\mathbf{X}}(\cdot)$ computes the first-order derivative with respect to input \mathbf{X} .

$$\eta = \epsilon(\nabla_{\mathbf{X}} J(\theta, \mathbf{X}, y)) \quad (4)$$

$$\eta = \epsilon \text{sign}(\nabla_{\mathbf{X}} J(\theta, \mathbf{X}, y)) \quad (5)$$

Kurakin et al. presented an iterative algorithm based on FGM and FGSM methods (Kurakin et al., 2016). In Algorithm 1, we present an iterative FGSM (IFGSM), where the adversarial samples $\mathbf{X}_n^{\text{adv}}$ are generated for the n th iteration.

During each iteration, the intermediate results get clipped to ensure that the resulting adversarial images lie within ϵ of the previous iteration.

Algorithm 1 IFGSM

Input: data X_{in}
 Initialize $X_0^{\text{adv}} = X_{in}$.
for $n = 0$ **to** $m - 1$ **do**
 $\mathbf{N} = \epsilon \text{sign}(\nabla_{\mathbf{X}} J(\theta, \mathbf{X}_n^{\text{adv}}, y_l))$
 $\mathbf{X}_{n+1}^{\text{adv}} = \text{Clip}_{\mathbf{X}, \epsilon}\{\mathbf{X}_n^{\text{adv}} + \mathbf{N}\}$
end for

Kurakin et al. also presented an iterative version of FGM where instead of just using the sign to determine the direction of a gradient, the gradient amplitudes contribute to the gradient update step. The iterative FGM (IFGM) is nearly identical to the IFGSM except that $\mathbf{N} = \epsilon \nabla_{\mathbf{X}} J(\theta, \mathbf{X}_n^{\text{adv}}, y_l)$.

Moosavi-Dezfooli et al. featured another attack called ‘Deepfool’, which is also based on iterative gradient adjustment (Moosavi-Dezfooli et al., 2016). However, Deepfool




Figure 2. Transferability properties for pruning. The green, red and cyan lines represent the first, second and third attack scenarios respectively. The blue line show the accuracies of pruned models without any attacks.

is different from IFGSM in that it does not scale and clip gradients. It is based on the idea that the separating hyper-planes in linear classifiers indicate the decision boundaries of different classes. It therefore iteratively perturbs an image \mathbf{X}_0^{adv} , linearises the classification space around \mathbf{X}_n^{adv} and moves towards the closest decision boundary. The step is chosen according to the l^0 , l^1 or even the l^p norm of \mathbf{X}_n^{adv} to the last-found decision boundary. The applied step is then used as \mathbf{X}_{n+1}^{adv} .

In practice Deepfool is found to produce smaller perturbations than the original IFGSM, which makes it a more precise attack (Moosavi-Dezfooli et al., 2016). In this paper we used an L2 norm-based version of Deepfool.

It should be noted that in this particular paper we were not interested in the absolute accuracy but the relative behaviours with a set of fixed parameters for adversarial attacks. We chose the strongest white box adversary model and picked three of the strongest iterative attacks. For all the experiments, we did not sweep all the possible hyper-parameters for the adversarial attacks, but picked empirically sensible hyper-parameters. The parameters are shown in Table 1 and were chosen in such a way that they generated perturbations of a sensible l^2 and l^0 and caused noticeable classification change.

Finally, we want to talk about how realistic the scenarios presented in this paper are and why we chose those particular attacks. First, we want to mention that not all of the attacks are transferable – as a matter of fact the attacks we

Network/Attack	I-FGSM		I-FGM		DeepFool	
	ϵ	i	ϵ	i	ϵ	i
LeNet5	0.02	12	10.0	5	0.01	5
CifarNet	0.02	12	0.02	12	0.01	3

Table 1. Attack hyper-parameters.

chose are amongst the least transferable ones. We chose those specific attacks to explore the lower bound of transferability and show how much of the subspaces actually survive the compression process (Tramèr et al., 2017). For DeepFool, we trained two models with different random initialisation and tested how transferable the adversarial samples are. For LeNet5 only 7% of the samples actually went across, whereas for CifarNet the transferability was better, but still only 60%. We now describe a slightly different scenario – the models attacking and defending are the same, just some of them are compressed.

For all of the attacks we made sure that in each of the iterations the perturbations stayed within the expected range.

4 EVALUATION

4.1 Pruning

Goodfellow et al. explained the existence of adversarial samples as follows (Goodfellow et al., 2015). Consider an adversarial sample as the original input \mathbf{x} with an addi-




Figure 3. LeNet5 accuracy with IFGSM and IFGM-generated adversarial samples with different ϵ values and number of epochs

tional noise η . When passing through multiple layers of matrix multiplication, this small noise eventually grows to a large enough value to shift the decision of the whole model. Given weights w of a particular layer of a neural network and adversarial sample $\tilde{x} = x + \eta$, the output of that particular layer is $w^T \tilde{x} = w^T x + w^T \eta$, and the adversarial perturbation causes the the output activations to grow by $w^T \eta$.

Figure 2 shows the performance of IFGSM, IFGM and DeepFool on pruned models under three different attack scenarios. The horizontal axis shows the densities of DNNs, effectively the ratio of the number of non-zero values to the total number of values. The vertical axis presents test accuracies of DNNs. Apart from showing the accuracies of pruned networks without any attacks (BASE ACC), we present the accuracies of the pruned models with three different attack scenarios. The first scenario corresponds to COMP \rightarrow COMP, the second scenario and third scenario corresponds to FULL \rightarrow COMP and COMP \rightarrow FULL respectively.

The first thing to note is that samples generated from the compressed models are transferable to the baseline model when densities are relatively large. This finding reinforces the idea that the adversarial samples are not scattered randomly but reside in large and contiguous high-dimensional spaces, enabling them to survive the effects of pruning. We suggest that pruning smooths the decision space by removing DNN weights that have little impact. This ultimately has an effect on IFGSM – with unimportant parts removed, the gradients now follow the path towards the most important and prominent parts of the space (first and fourth plots on Figure 2). As a result, relatively small perturbations based on compressed models generalise very well on the uncompressed model when networks are not heavily pruned. For all of the attacks, adversarial samples generated on networks with very small densities are not effective on the baseline

networks (increase in the red line and fall of the blue line near zero in all plots of Figure 2) Heavily-pruned networks acquire a feature space that is hugely different from the baseline models, and this limits the transfer of adversarial samples. However, low-density networks often suffer large losses in classification accuracy, making them infeasible to deploy in real life.

When one compressed model is attacking another (Comp to Comp, green line on Figure 2), we see a general trend that attacks remain transferable.

Using an uncompressed model to attack a compressed one (cyan line, Figure 2), we observe a slight increase in accuracy occurs when densities are small, but then a rapid drop when they keep decreasing; this effect occurs when the base accuracy (blue line) also starts to drop. We view pruning as a regularization method which removes local minima from the large optimization space. When the blue line is just starting to decrease, this turning point is the preferred density, where the network just stops overfitting. In other words, the preferred density of a network represents the minimal number of parameters needed to accomplish the same accuracy as the dense model. To further illustrate the effect of pruning, we present CifarNet with Deepfool and IFGSM in Figure 4. The horizontal axis shows the accuracy of the baseline model and the vertical axis has the adversarial accuracy. In terms of adversarial attacks, although limited, we observe that networks that reach preferred density show a protective nature in Figure 4.

We observe attacks perform worse on LeNet5 in comparison to CifarNet. We notice that LeNet5 inherently achieves a larger accuracy on MNIST, meaning that the loss is smaller on the evaluation dataset as well. When building attacks, we often make use of the amplitude of gradients to generate adversarial samples, and the smaller loss associated with LeNet5 implies that it is less vulnerable to attack. This phe-




Figure 4. CifarNet accuracy with IFGSM and DeepFool-generated adversarial samples with corresponding base accuracy.

nomenon becomes more apparent on Deepfool and IFGM, since these attacks employ the gradients amplitudes to generate adversarial images.

Intuitively, pruning largely preserves the feature space of a baseline CNN, so adversarial samples remain transferable. This empirical observation confirms recent pruning discoveries, that pruned networks distil feature spaces (Liu et al., 2018; Frankle & Carbin, 2018). In addition, our observation is in line with the suggestions made by Tramèr et al.: if the feature spaces are similar, adversarial samples stay highly transferable (Tramèr et al., 2017).

Summary

1. The transferability of adversarial samples between pruned and full models remains when networks are slightly pruned.
2. For compressed models attacking uncompressed models, we observe worse transferability when models are heavily pruned, but the original accuracy of these sparse models decreases significantly.
3. For uncompressed models attacking compressed models, the accuracy is maximised, and the transferability minimised, at the preferred density where the network stops overfitting.

4.2 Fixed-point Quantisation

Fixed-point quantisation refers to quantising both weights and activations to fixed-point numbers; for example, for a bitwidth of four we use a one-bit integer plus a three-bit fraction. Multiplication is much faster as we can use integer operations rather than floating point. Figure 5 shows the performance of adversarial attacks on quantised models under our three different attack scenarios. Attack performance

stays nearly constant at bitwidths higher than 8. When using fewer bits for both weights and activations, the model shows defensive behaviour, mainly because of the reduction in integer precision.

Intuitively, we have two effects when values are quantised to smaller bitwidths. First, a smaller bitwidth means fewer fractional bits causing a loss in precision, and introducing much the same effect as pruning. Second, it can mean fewer integer bits, so weight and activation values are smaller. Thus, our models in 4-bit fixed-point quantisation have smaller weight and activation values and contain more zeros than models at higher precisions. In Figure 6.a, we show the cumulative distribution function (CDF) of CifarNet with different fixed-point quantisations. There are clearly more zeros in the 4-bit CifarNet – its cumulative density reaches around 0.9 when value is at 0. The clipping effect is also more obvious in the 4-bit model, since it only has a 1-bit integer part; we can see the 4-bit model has its weights CDF reach 1.0 before all other bitwidths in Figure 6.a.

Using the adversarial examples generated by compressed models to attack the baseline model (red line), we observe both Deepfool and IFGSM methods become less effective on LeNet5 and CifarNet. The same phenomenon occurs when we use adversarial samples generated by the baseline model to attack quantised models (cyan line). We suggest that during quantisation, reducing fractional bits will not hugely impact the attacks’ performance at high bitwidths, but introduces a similar effect to pruning at low bitwidths. In addition, reducing integer bits essentially introduces large differences between the baseline model and the quantised ones for adversarial attacks. Using FGM to attack LeNet5, on the other hand, gives very different behavior. We also noticed that attacking LeNet5 require large epsilon values and more iterative runs. We suggest this is because of the accuracy issue with LeNet5 that we’ve addressed earlier – attacks that rely on gradient magnitudes struggle with

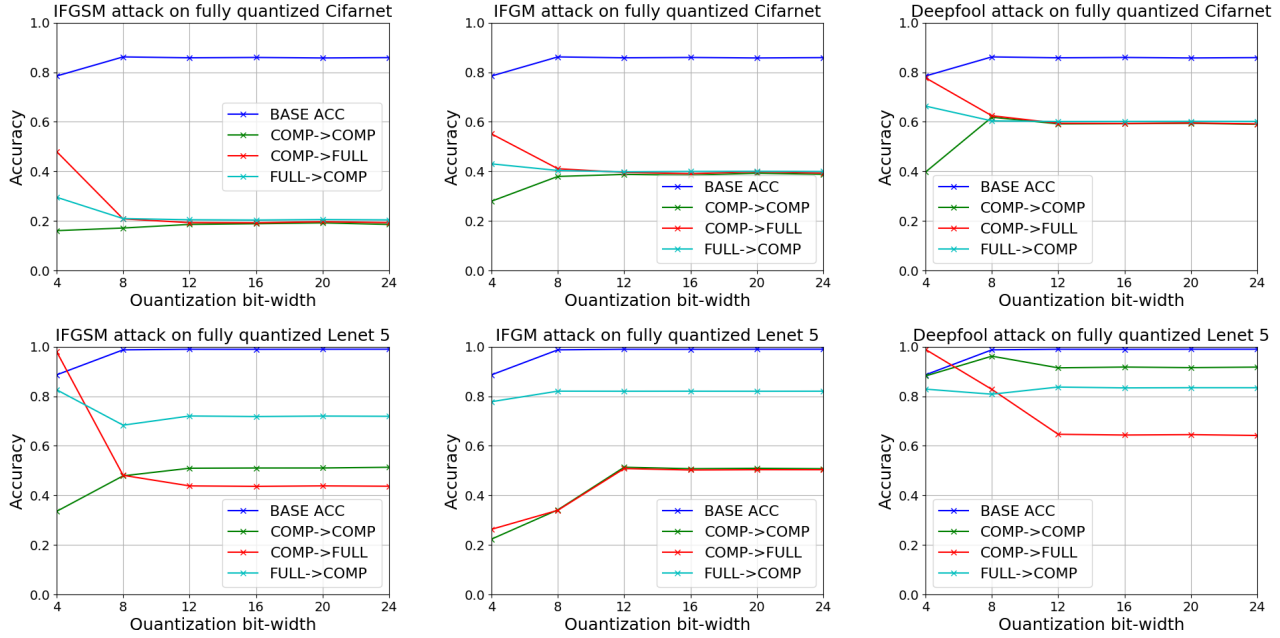


Figure 5. Transferability properties for quantising both weights and activations. The green, red and cyan lines represent the first, second and third attack scenarios respectively. The blue line show the accuracy of quantised models without any attacks.

networks that achieve high accuracy.

By reducing the length of the fraction, the rounding process of fixed-point quantisation becomes more lossy. Uniformly adding quantisation noise to each individual weight does not affect attack performance. As we can see in Figure 5, all three attack scenarios show a stable performance when bitwidths are higher than 8, where the difference lies in the length of the fraction. When the network gets quantised down to 4 bits, quantisation behaves rather like pruning – a large part of the network gets zeroed out.

By reducing integer bitwidth, we are clipping the numerical values. Theoretically, clipping weights is different from clipping activations. For the former, we first consider how to create an adversarial sample with minimal perturbation. Intuitively, the way to achieve such a perturbation with minimal changes on the input image is tweak pixels with large weight values that are connected to important activations. Thus a small change in input image pixels can have the maximal effect on activation values. When weights are clipped, adversarial attacks see more weights as having equal importance because they saturate to the same maximal value. This undermines attack transferability between quantised models and baseline models. For example, on a quantised network, an adversarial example X_i considers $w_i = \max(W_i)$ to be the largest weight associated with the important activations among all the weights (W_i) associated with activation a_i . This relationship $w_i = \max(W_i)$ might break on the baseline model and thus the adversarial sample becomes less

effective. In Figure 6, a 4-bit fixed-point quantisation clearly shows a clipping effect on weight values, which contributes to the marginal defensive nature we observed in Figure 5.

When activations are clipped to a smaller maximal value, transferability between quantised and baseline models becomes worse. Figure 6.b shows how activations are clipped to different maximum values. Consider a simple case, where an adversarial example overdrives one activation to be larger than others in the same layer to cause a misclassification. Clipping the activation values forces the attacker to find more subtle ways of achieving differential activation, which is significantly harder.

Clipping weights and clipping activations can both significantly affect attack performance. As we can see from the cyan line in Figure 5, at smaller bit widths, all our attack scenarios show an increase in accuracy. except for LeNet5 attacked by IFGM. In terms of transferability, adversarial examples remain transferable between quantised and baseline models under both IFGSM and IFGM when fractional bits are lost, but integer bits start to decrease, transferability becomes worse.

Surprisingly, we find that Deepfool, unlike IFGM and IFGSM, struggles to generate effective adversarial samples when models are quantized. Restricting values to discrete levels mean now an attack has to inject a large enough perturbation to the values to push them to the neighbouring quantization level. Since Deepfool is a very fine-grained attack, it struggles to do this.

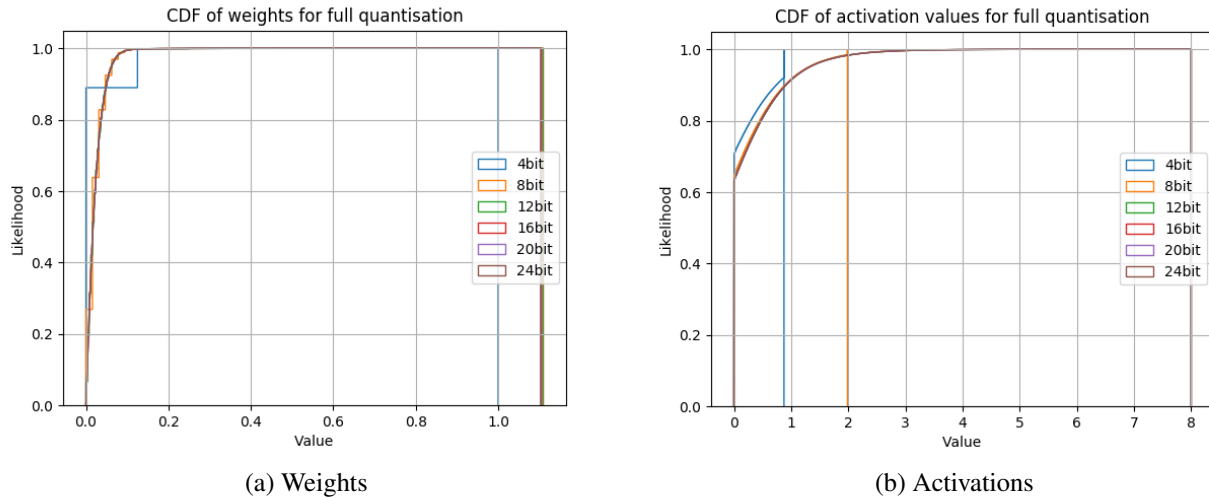


Figure 6. Cumulative distribution function (CDF) for all weights and activations in quantised CIFARNet. Ten randomly chosen input images from the validation dataset were used to generate CDF of activation values.

Although IFGSM shows slightly higher accuracy at low precision, this protective behavior is only marginal. Attacks still show good performance in all three scenarios compared to both IFGM and DeepFool if we consider the classification accuracy (Figure 5).

In summary, transferability is still a hazard for quantized models. But there is some marginal protection when activation values are quantized. Intuitively, this forces many activations to saturate and thus makes it harder for the attacker to overdrive certain values.

Summary

1. The transferability of adversarial samples between quantised and baseline models is not affected by reducing fractional bitwidth at high precision.
2. Aggressive reduction of fractional bits introduces the same effect as fine-grained pruning.
3. Smaller integer bitwidths of weights and activations make it marginally harder to attack the baseline model using adversarial samples generated from compressed models. This suggests that the network’s knowledge is contained in both activations and weights.
4. We hypothesize that clipping activations changes the feature space of CNNs and thus marginally protects models from transferability.

5 CONCLUSION

This paper reports an empirical study of the interaction between adversarial attacks and neural network compression.

Both quantisation and pruning sparsify the network, i.e. a greater number of activations and weights are zero. Attacks generated from heavily pruned models work effectively against the underlying baseline model. However, low-density DNNs are somewhat defensive when attacked by adversarial samples generated from the baseline model using fast-gradient-based methods. Quantisation is different in that adversarial samples from fast-gradient-based methods become marginally harder to transfer when models are heavily quantised. This defensive behaviour appears due to the reduction in integer bits of both weights and activations rather than to the truncation in fractional bits.

The broader implications are that attacks on DNN classifiers that involve adversarial inputs may be surprisingly portable. Even if a firm ships only a compressed version of its classifier in widely distributed products, such as IoT devices or apps, attacks that people discover on these compressed classifiers may translate fairly easily to attacks on the underlying baseline model, and thus to other compressed versions of the same model. Just as software vulnerabilities such as Heartbleed and Spectre required the patching of many disparate systems, so also a new adversarial sample may defeat many classifiers of the same heritage. Firms should be aware that while shipping a compressed classifier may give real benefits in terms of performance, it may not provide much in the way of additional safety or security.

ACKNOWLEDGEMENTS

Partially supported with funds from Bosch-Forschungsstiftung im Stifterverband.

REFERENCES

- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. *International Conference on Learning Representations (ICLR)*, 2015.
- Carlini, N., Mishra, P., Vaidya, T., Zhang, Y., Sherr, M., Shields, C., Wagner, D., and Zhou, W. Hidden Voice Commands. In *25th USENIX Security Symposium (USENIX Security 16)*. USENIX Association, 2016.
- Chen, Y.-H., Emer, J., and Sze, V. Eyeriss: A spatial architecture for energy-efficient dataflow for convolutional neural networks. In *ACM SIGARCH Computer Architecture News*, volume 44, pp. 367–379. IEEE Press, 2016.
- Courbariaux, M., Hubara, I., Soudry, D., El-Yaniv, R., and Bengio, Y. Binarized neural networks: Training deep neural networks with weights and activations constrained to+1 or-1. *arXiv preprint*, 2016.
- Eykholt, K., Evtimov, I., Fernandes, E., Li, B., Rahmati, A., Xiao, C., Prakash, A., Kohno, T., and Song, D. Robust Physical-World Attacks on Deep Learning Visual Classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1625–1634, 2018.
- Frankle, J. and Carbin, M. The lottery ticket hypothesis: Finding sparse, trainable neural networks. 2018.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *International Conference on Learning Representations (ICLR)*, 2015.
- Guo, Y., Yao, A., and Chen, Y. Dynamic network surgery for efficient DNNs. In *Advances in Neural Information Processing Systems*, 2016.
- Han, S., Liu, X., Mao, H., Pu, J., Pedram, A., Horowitz, M. A., and Dally, W. J. EIE: Efficient Inference Engine on Compressed Deep Neural Network. *SIGARCH Comput. Archit. News*, 44(3):243–254, June 2016a. ISSN 0163-5964. doi: 10.1145/3007787.3001163.
- Han, S., Mao, H., and Dally, W. J. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *International Conference on Learning Representations (ICLR)*, 2016b.
- Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., and Bengio, Y. Quantized neural networks: Training neural networks with low precision weights and activations. *J. Mach. Learn. Res.*, pp. 6869–6898, 2017.
- Kim, D., Ahn, J., and Yoo, S. Zena: Zero-aware neural network accelerator. *IEEE Design & Test*, 35(1):39–46, 2018.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Krizhevsky, A., Nair, V., and Hinton, G. The CIFAR-10 dataset. 2014.
- Kurakin, A., Goodfellow, I. J., and Bengio, S. Adversarial examples in the physical world. *CoRR*, abs/1607.02533, 2016.
- LeCun, Y., Cortes, C., and Burges, C. MNIST handwritten digit database. 2, 2010.
- LeCun, Y. et al. LeNet-5, convolutional neural networks. pp. 20, 2015.
- Li, F., Zhang, B., and Liu, B. Ternary weight networks. *1st NIPS Workshop on Efficient Methods for Deep Neural Networks (EMDNN)*, 2016.
- Lin, D., Talathi, S., and Annapureddy, S. Fixed point quantization of deep convolutional networks. In *International Conference on Machine Learning*, pp. 2849–2858, 2016.
- Liu, Z., Sun, M., Zhou, T., Huang, G., and Darrell, T. Rethinking the value of network pruning. *arXiv preprint arXiv:1810.05270*, 2018.
- Moosavi-Dezfooli, S., Fawzi, A., and Frossard, P. DeepFool: a simple and accurate method to fool deep neural networks. 2016.
- Papernot, N., McDaniel, P., and Goodfellow, I. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *arXiv preprint*, 2016a.
- Papernot, N., McDaniel, P., Jha, S., Fredrikson, M., Celik, Z. B., and Swami, A. The limitations of deep learning in adversarial settings. pp. 372–387, 2016b.
- Papernot, N., McDaniel, P., Goodfellow, I., Jha, S., Celik, Z. B., and Swami, A. Practical black-box attacks against machine learning. pp. 506–519, 2017.
- Parashar, A., Rhu, M., Mukkara, A., Puglielli, A., Venkatesan, R., Khailany, B., Emer, J., Keckler, S. W., and Dally, W. J. SCNN: An accelerator for compressed-sparse convolutional neural networks. In *2017 ACM/IEEE 44th Annual International Symposium on Computer Architecture (ISCA)*, pp. 27–40, June 2017. doi: 10.1145/3079856.3080254.
- Ren, S., He, K., Girshick, R., and Sun, J. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pp. 91–99, 2015.

Seo, M., Kembhavi, A., Farhadi, A., and Hajishirzi, H. Bidirectional attention flow for machine comprehension. *arXiv preprint*, 2016.

Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I. J., and Fergus, R. Intriguing properties of neural networks. *CoRR*, abs/1312.6199, 2013.

Tramèr, F., Papernot, N., Goodfellow, I., Boneh, D., and McDaniel, P. The space of transferable adversarial examples. *arXiv preprint*, 2017.

Warde-Farley, D. and Goodfellow, I. Adversarial perturbations of deep neural networks. *Perturbations, Optimization, and Statistics*, pp. 311, 2016.

Zhao, Y., Gao, X., Mullins, R., and Xu, C. Mayo: A framework for auto-generating hardware friendly deep neural networks. 2018.