Stochastic-Shield: A Probabilistic Approach Towards Training-Free Adversarial Defense in Quantized CNNs

Lorena Qendro*
University of Cambridge

René de Jong
Arm ML Research Lab

Sangwon Ha
Arm ML Research Lab

Partha Maji
Arm ML Research Lab

ABSTRACT

Quantized neural networks (NN) are the common standard to efficiently deploy deep learning models on tiny hardware platforms. However, we notice that quantized NNs are as vulnerable to adversarial attacks as the full-precision models. With the proliferation of neural networks on small devices that we carry or surround us, there is a need for efficient models without sacrificing trust in the prediction in presence of malign perturbations. Current mitigation approaches often need adversarial training or are bypassed when the strength of adversarial examples is increased.

In this work, we investigate how a probabilistic framework would assist in overcoming the aforementioned limitations for quantized deep learning models. We explore **Stochastic-Shield**: a flexible defense mechanism that leverages an input filtering layer and a probabilistic deep learning approach materialized via Monte Carlo dropout. We show that it is possible to jointly achieve efficiency and robustness by accurately enabling each module without the burden of re-training or ad hoc fine-tuning.

CCS CONCEPTS

• **Security and Privacy** → **Intrusion/anomaly detection and malware mitigation**; • **Computing methodologies** → **Artificial intelligence**.

KEYWORDS

Adversarial Attack, Probabilistic Deep Learning, Adversarial Mitigation, Quantized CNNs

ACM Reference Format:

Lorena Qendro, Sangwon Ha, René de Jong, and Partha Maji. 2021. Stochastic-Shield: A Probabilistic Approach Towards Training-Free Adversarial Defense in Quantized CNNs. In *1st Workshop on Security and Privacy for Mobile AI (MAISP'21)*, June 24, 2021, Virtual, WI, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3469261.3469404>

1 INTRODUCTION

Quantized neural networks have facilitated the deployment of deep learning models in real-time applications by reducing computation

*This work was done when the author was affiliated with Arm ML Research Lab.

This work is licensed under a Creative Commons Attribution International 4.0 License.

MAISP'21, June 24, 2021, Virtual, WI, USA

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8601-2/21/06.

<https://doi.org/10.1145/3469261.3469404>

and memory costs. However, these networks are as vulnerable to adversarial attacks as traditional non-optimized deep learning networks [16]. Although, the reduced bit-width¹ should help denoising small perturbations, there is negligible or no improvement to any perturbation strength from current adversarial attacks. As shown in Figure 1 in the presented image classification tasks, the prediction accuracy of the quantized models drops significantly with the increasing strength of the malign perturbation. The difference to the floating-point precision models in terms of accuracy is negligible, suggesting that restricting the range of the possible values through quantization only does not provide any protection against adversarial attacks. This vulnerability is a huge concern in safety-critical scenarios, such as health applications, autonomous driving, face and fingerprint identification, and voice recognition [3].

Adversarial attacks can be particularly designed to be virtually indistinguishable to human perception, or still visible to the human eye but not perceived as an attack (e.g. graffiti or stickers on road signs [13, 14]). Therefore, it is important to provide robustness and adversarial defense to real scenarios against attacks in a wide range of the strength scale. Adversarial perturbations do not stop at the feature level but they are propagated through the network, by easily bypassing techniques which focus on the feature space only [27]. While, other state-of-the-art mitigation techniques based on adversarial training are able to mitigate these attacks [11, 17], the multiplicity of attacks grows continuously, therefore, these approaches add the significant burden of several cycles of retraining and re-deployment. Other non adversarial-specific works focus on profiling the path benign samples follow during training and use this information to detect malign examples that perform a different path during inference [10, 22]. The downside of these methods is in the fact that the attacker can profile the behavior too, moreover, they would fail in cases of data shift.

Given these limitations, we investigate how and to what extent can a probabilistic approach help this ecosystem. Our intuition suggests that by providing randomness in the network, it will be more difficult for the attack to follow a specific path towards misclassification, increasing its cost of production. Through Stochastic-Shield, we study how to obtain robustness against adversarial attacks via a modular defense mechanism which can be applied to already trained deep learning models. Moreover, since the majority of these models are and will be deployed on ubiquitous embedded devices which need a reduced computation demand by design, we consider 8-bit quantized deep learning models.

¹Number of bits used to represent the weights and activations of the neural network (8-bit integer in this work).

Stochastic-Shield adopts an *input filtering layer* to create a first shielding in addition to a probabilistic approach. This layer consists of two blocks: the input quantization and median smoothing block, which restrict the degrees of adversarial freedom in the input space. Additionally, the probabilistic framework via *Monte Carlo dropout* (MCDrop) [9] introduces stochasticity to the network and, therefore, provides an implicit ensemble of models at runtime to make it more difficult for the attacker to aim at a specific set of neurons. Our experiments, which aim to measure robustness in presence of attacks at different strength levels, show that the use of the probabilistic framework becomes crucial when the attack strength increases, and operating at both feature and models space, it is possible achieve higher robustness for quantized models.

2 BACKGROUND

In this section, we navigate through common adversarial attacks, current mitigation approaches and, finally, we present the background on the probabilistic approach adopted in our framework.

2.1 Adversarial Attacks

Given a clean image X , an adversarial attack introduces small perturbations ∇ such that the class prediction for X and X^{adv} differs.

Fast Gradient Sign Method (FGSM). Introduced by [11], FGSM is a single step attack which aims to find the adversarial perturbations by moving in the opposite direction to the gradient of the loss function $L(X, y)$ w.r.t. the image (∇):

$$X^{adv} = X + \epsilon * \text{sign}(\nabla_X L(X, y)), \quad (1)$$

where ϵ is the step size which restricts the l_∞ of the perturbation.

Projected Gradient Descend (PGD). A stronger variant of FGSM [12] consists on applying it iteratively introducing a small step α :

$$X_0^{adv} = X', X_{n+1}^{adv} = \text{clip}_X^\epsilon \{X_n^{adv} + \alpha \text{sign}(\nabla_X L(X_n^{adv}, y))\} \quad (2)$$

where

$$X' = X + \epsilon_1 * \text{sign}(\mathcal{N}(\mathbf{0}^d, \mathbf{I}^d)) \quad (3)$$

(with parameters ϵ_1, ϵ such as $\epsilon_1 < \epsilon$) is an additional prepended random step [24] which avoids going towards a false direction of ascent. Both steps, 2 and 3, make PGD [17] which proves to be a universal first-order attack.

Carlini Wagner (CW). CW [6] is a very strong iterative attack which works with various l_p norms. If we consider CW $L2$ which aims to minimize the following function:

$$\|\nabla\|_2^2 + cf(X + \nabla) \quad (4)$$

The constant c is adjusted through line search in order to pick a point where $f(X + \nabla)$ becomes negative and $f(\cdot)$ is defined as:

$$f(X^{adv}) = Z(X^{adv})_y - \max\{Z(X_{y_a}^{adv}) : y_a \neq y\} \quad (5)$$

Where $Z(X_{y_a}^{adv})$ gives the pre softmax predictions for class y_a on the adversarial image X^{adv} , instead, y is the correct class which the attack wants to diverge from.

2.2 Adversarial detection and defense

Adversarial training [11, 17, 28] minimizes the risk that a perturbed sample can be misclassified, although its side-effects are not negligible. Classifiers trained with adversarial examples learn fundamentally different representations compared to standard classifiers reducing accuracy [25] or they can cause disparity on accuracy for both clean and adversarial samples between different classes [26]. In addition, they are extremely resource-consuming and cannot be applied to already deployed or trained network suggesting a very limited generalization and applicability in the real world.

Feature Squeezing [27] aims to reduce the degrees of freedom available to an adversary by “squeezing” out unnecessary input features e.g. color-depth bit reduction. The technique consists on running the network with both the original and transformed input and detect the adversarial if the difference in the prediction passes a certain α threshold. Although, its computational overhead is minimal, it can be easily bypassed by increasing adversarial strength [21]). In Stochastic-Shield, we adopt the promising transformations (input quantization and local smoothing) used in this approach. We use it differently as an additional filtering layer to the NNs for mitigation in reducing the chance of perturbed pixels to propagate in the quantized model.

The taboo trap [22] trains the model applying restrictions on activations, and during inference any sample violating these restrictions is considered adversarial. This solution is not generic as it would fail in presence of covariate shifts, and in addition, the attack can discover the limit of activations used in the restriction too.

Ptolemy [10] detects adversarial samples at runtime based on the observation that malign samples tend to activate distinctive paths from those of benign inputs. Its main drawback having a backdoor in the model since the attack can learn which specific paths get activated by profiling too. In contrast, Stochastic-Shield does not depend on any specific path, therefore, profiling its activation path will give no information due to its stochastic properties.

Defensive quantization (DQ) [16] observes that the quantization operation amplifies the perturbation noise when passing through the deep NN. To overcome the aforementioned issue, this method proposes controlling the Lipschitz constant of the network during quantization, such that the magnitude of the adversarial noise remains non-expansive during inference. Nevertheless, DQ needs significant quantization-aware training and the quantization is applied to activations only, while the weights are still represented in 32-bit floating-point which is not a real world scenario.

Stochastic approaches assess adversarial attacks by looking at the likelihood of perturbed examples and the distance between perturbed and clean examples for each classifier [2, 15, 20]. [5] propose a Gaussian Process (GP) hybrid deep NN to help mitigating adversarial attacks, however, a closed-form solution for GPs has a $O(n^3)$ computational complexity. In addition, GPs are very sensitive to quantization and thus implementation on embedded hardware is challenging. A promising direction, Monte Carlo Dropout (MCDrop) [8], casts dropout training as approximate inference in Bayesian CNNs. [7] investigate model confidence on adversarial samples through MCDrop by looking at uncertainty estimates. They operate a two-feature approach similar to ours, however, they need to observe the density estimates during training to calculate to

Figure 1: Accuracy under adversarial attack for floating-point precision and quantized neural networks. 8-bit integer-quantized NNs and 32-bit floating-point NNs are equally vulnerable to attacks. Setup: VGG16 and MobileNetV2 on CIFAR10 under FGSM and PGD (10 iterations) attacks.

Figure 2: Stochastic-Shield = Input filtering layer + A stochastic neural network (Bernoulli masks to the weight matrix via the dropout-based stochastic layer). Input filtering is performed only once, instead the model is run T times to perform the forward stochastic passes needed for MCDrop.

detect the points that lie far from the data manifold, which might lead to issues adapting to data shifts. While these techniques study the uncertainty for adversarial detection, we envision using the probabilistic framework to prevent the attack from changing the prediction by making it more difficult for the attacker to overcome the randomness introduced by dropout and finding a deterministic path towards the wrong class.

2.3 Monte Carlo Dropout

Standard dropout [23] is regularization technique used to avoid overfitting. It consists of dropping random units with a certain probability p which allows a shift from a deterministic to a stochastic model constituted of implicit ensemble of networks. Let a fully connected layer with dropout be:

$$\begin{aligned}
 \mathbf{z}_{[i]}^{(l)} &\sim \text{Bernoulli}(\cdot | \mathbf{p}_{[i]}^{(l)}) \\
 \tilde{\mathbf{W}}^{(l)} &= \text{diag}(\mathbf{z}^{(l)}) \mathbf{W}^{(l)} \\
 \mathbf{y}^{(l)} &= \mathbf{x}^{(l)} \tilde{\mathbf{W}}^{(l)} + \mathbf{b}^{(l)} \\
 \mathbf{x}^{(l+1)} &= f^{(l)}(\mathbf{y}^{(l)})
 \end{aligned} \tag{6}$$

where $\mathbf{x}^{(l)}$ and $\mathbf{y}^{(l)}$ are the input and output of that layer, and $f^{(l)}(\cdot)$ is the nonlinear activation function. $\mathbf{W}^{(l)}$ is the weight matrix of l with dimensions $K^{(l)} \times K^{(l-1)}$ and $\mathbf{b}^{(l)}$ is the bias vector of dimensions $K^{(l)}$. Here $\mathbf{z}_{[i]}^{(l)}$ are Bernoulli distributed random variables with some probabilities $\mathbf{p}_{[i]}^{(l)}$. The $\text{diag}(\cdot)$ maps vectors to diagonal matrices.

[9] proved the equivalence between dropout training in a neural network and approximate inference in a deep Gaussian Process. They showed that the objective converges to a minimization of the Kullback-Leibler divergence between an approximate distribution and the posterior of a deep Gaussian process marginalized over its covariance function parameters. The true posterior distribution is, therefore, approximated by the variational distribution $q(\tilde{\mathbf{W}}^{(l)})$ where $\tilde{\mathbf{W}}^{(l)}$ represents the random variables used in dropout operations as described in (1).

$$\begin{aligned}
 \mathbf{z}_{[i]}^{(l)} &\sim \text{Bernoulli}(\cdot | \mathbf{p}_{[i]}^{(l)}) \\
 q(\tilde{\mathbf{W}}^{(l)}) &= \text{diag}(\mathbf{z}^{(l)}) \mathbf{W}^{(l)}
 \end{aligned} \tag{7}$$

Finally, they perform Monte Carlo (MC) sampling of the random variables $\mathcal{W}$,

$$q(\mathbf{y}|\mathbf{x}) = \frac{1}{T} \sum_{t=1}^T q(\mathbf{y}|\mathbf{x}, \mathcal{W}_t), \quad (8)$$

where T is the number of MC samples. This method is called Monte Carlo Dropout (MCDrop) and is equivalent to performing T stochastic passes. Similarly for convolution layers [8], the sampled Bernoulli random variables $\mathbf{z}_{i,j,k}$ are applied as masks to the weight matrix $\mathbf{W}_i \cdot \text{diag}([\mathbf{z}_i, j, k])$ which is equivalent to setting weights to 0 for different elements of the input.

3 STOCHASTIC-SHIELD

Adversarial attacks increase the uncertainty in the model prediction by adding perturbations that trick the model to change the prediction class. A probabilistic framework can capture the uncertainty in the data and model parameters, which are spaces where the adversarial attacks particularly aim to operate. Therefore, in order to offer robustness against adversarial attacks, our intuition led to exploring Stochastic-Shield.

Stochastic-Shield (Figure 2) initially adopts an *input filtering* layer to create a first shielding which aims to filter out some of the perturbations before they are propagated in the network. Input filtering, which represents quantization and median smoothing in the input space, can capture the perturbations and offers a solution to reduce the degrees of freedom available to the adversary by filtering out unnecessary features. This layer can reduce the effect of the perturbations in the input but it cannot help against stronger attacks which are propagated in the deeper layers.

To solve the latter, we adopt Monte Carlo dropout (MCDrop), which creates an implicit ensemble of quantized models by running the same model multiple times with dropout layers activated during inference. Compared to deterministic deep ensembles which, once trained, are fixed; the stochasticity via MCDrop helps mitigating the adversarial attack by making it more difficult to profile the specific activation path. This would mean attacking an implicit ensemble of networks that dynamically changes at every run (i.e. different path towards the prediction class) significantly increasing the production cost of the attack. While operating in two different core parts of the network, we observe these two techniques together to be very effective on detecting and mitigating adversarial attacks for 8-bit integer-quantized deep learning models.

Finally, our approach is flexible and adaptable to the energy envelope as both shielding levels can be applied separately or co-operatively to provide the best defense.

To answer our initial question on how and to what extent can a probabilistic approach help mitigating adversarial attacks in quantized models, we test Stochastic-Shield on 3 networks: MobileNetV2 and VGG16 for CIFAR10 and a CNN based one for Street View House Number (SVHN) (See 4 for details). For each of the networks and datasets, we measure the accuracy and the expected calibration error (ECE) such as

$$p(y|x) \approx \frac{1}{T} \sum_{t=1}^T \text{Softmax}(\hat{y}_t) \quad (9)$$

where $\hat{y}_t$ is the prediction for sample t . ECE measures the correspondence between predicted probabilities and empirical accuracy. It is computed as the average gap between bin accuracy and bin predicted probability for B bins

$$ECE = \sum_{b=1}^B \frac{n_b}{N} |\text{accuracy}(b) - \text{confidence}(b)| \quad (10)$$

where n_b is the number of predictions falling in bin b .

Figure 3 indicates the accuracy of the pre-trained 8-bit quantized models under various adversarial attacks (FGSM and PGD) and strengths. The mitigation techniques are applied separately and together to show the efficacy of each individually and in unison. We see that input filtering is able to keep an acceptable accuracy when the adversarial strength is low, but when the attacks are stronger it continuously fails to achieve the results of MCDrop. In these experiments, MCDrop uses 5 samples as recommended by previous work [19] suggesting the optimal number of MC samples for a well-balance performance and computation tradeoff. However, in all cases combination of both techniques that make Stochastic-Shield achieve the desired robustness presenting a considerably higher accuracy (improvements up to 3-folds compared to Vanilla models and up to 2-folds compared to input filtering only) in presence of attacks and better calibrated models (see Figure 4). As expected, the vanilla model achieves high ECE indicating that it does not guarantee calibration under adversarial perturbation. Interestingly, for MobileNetV2, both input filtering and MCDrop used separately do not achieve good calibration when used separately, however when used together in Stochastic-Shield, the model is well-calibrated. In all scenarios, using Stochastic-Shield leads to better calibrated models. Higher the adversarial strength, bigger the need for the probabilistic approach and Stochastic-Shield.

Table 1 shows the results when using a CW L2 non-targeted attack reinforcing this finding. We see here that although our approach is the best performing, the difference to the other approaches is not big since the attack is not very strong (compared to PGD($\epsilon = 32$) e.g. which drops the accuracy to less than 10%). CW is very costly for embedded devices, however, a CW L_∞ attack would have a greater effect on accuracy and Stochastic-Shield would be a must to mitigate it.

Table 1: Accuracy and Expected Calibration Error (ECE) of models under CW L2 attack with default parameters from [6]. [Best performing in bold.]

	Vanilla		Input Filter		MCDrop 5		Stochastic-Shield	
	Acc $\uparrow$	ECE $\downarrow$	Acc $\uparrow$	ECE $\downarrow$	Acc $\uparrow$	ECE $\downarrow$	Acc $\uparrow$	ECE $\downarrow$
MobileNetV2	51.9%	0.75	52.5%	0.74	53.3%	0.73	54.1%	0.75
VGG16	66.7%	0.8	65.7%	0.8	70.1%	0.73	70.7%	0.73
CNN	69.7%	0.68	80.8%	0.75	83.6%	0.71	83.6%	0.7

4 IMPLEMENTATION DETAILS

We consider 3 networks: MobileNetV2 and VGG16 for CIFAR10 and a CNN based one for Street View House Number (SVHN). All networks have been trained with dropout layers which is kept activated at inference time keeping its training dropout rate. The training

Figure 3: Accuracy on clean ($\epsilon = 0$) and perturbed input for FGSM and PGD attacks with different perturbation strength for Vanilla (no defence) DNN, IF (Input Filter, no_bits=5, median_filter=2x2), MCDrop5 (5 samples), and Stochastic-Shield. All models are quantized with 8-bit weights and activations.

Figure 4: Expected calibration error for Vanilla (no defence) DNN, IF (Input Filter, no_bits=5, median_filter=2x2), MC5 (MCDrop 5 samples), and Stochastic-Shield. Lower error indicates better calibrated model, less prone to falling into the adversarial trap.

does not include any adversarial training or hyper-parameter tuning to increase robustness. The already trained networks are quantized to use 8-bit weights and activations via TensorFlow Model Optimization Toolkit [1].

We implement the adversarial attacks (FGSM, PGD, and CW) in Keras using the library *adversarial-robustness-toolbox* [18]. For FGSM and PGD, we set the perturbation strength $\epsilon = \{0, 2, 8, 16, 32\}$ and colors are represented from 0-255. PGD is an iterative attack and in our experiments we use 10 iterations (attack applied in iterative manner as explained in Section 2.1). We use the L2 norm CW attack which aims to minimize the objective function by using

the gradient descent. We used 1000 iterations, 9 binary search steps and confidence 0.0 as in [6]. These adversarial perturbations are applied to the whole test set and all the metrics represent the overall accuracy when an adversarial attack is applied to each sample.

For all experiments, we are performing 5 forward passes for MCDrop (MC5). This choice was made considering previous work [4, 19] and the fact that although less computation heavy than pure Bayesian approaches or Gaussian Processes, MCDrop can introduce some overhead to the system given the multiple forward passes. Moreover, keeping a low number of samples allows us to see what the technique can achieve keeping the overhead at a minimum.

Table 2: Architecture and implementation details of the considered deep learning models.

VGG16		MobileNetV2		CNN based	
Layer	Details	Layer	Details	Layer	Details
Conv2D(BN, ReLU)	3x3x64	Conv2d(BN, ReLU)	3x3x32	Conv2d(BN, ReLU)	3x3x32
Dropout	0.3	InvertedResidual	16	Conv2d(BN, ReLU)	3x3x32
Conv2D(BN, ReLU)	3x3x64	InvertedResidual (x2)	24	MaxPooling2D	2x2
MaxPooling2D	2x2	InvertedResidual (x3)	32	Dropout	0.3
Conv2D(BN, ReLU)	3x3x128	InvertedResidual (x4)	64	Conv2d(BN, ReLU)	3x3x64
Dropout	0.4	Dropout	0.25	Conv2d(BN, ReLU)	3x3x64
Conv2D(BN, ReLU)	3x3x128	InvertedResidual (x3)	96	MaxPooling2D	2x2
MaxPooling2D	2x2	Dropout	0.25	Dropout	0.3
Conv2D(BN, ReLU)	3x3x256	InvertedResidual (x3)	160	Conv2d(BN, ReLU)	3x3x128
Dropout	0.4	Dropout	0.25	Conv2d(BN, ReLU)	3x3x128
Conv2D(BN, ReLU)	3x3x256	InvertedResidual	320	MaxPooling2D	2x2
Dropout	0.4	Dropout	0.25	Dropout	0.3
Conv2D(BN, ReLU)	3x3x256	Conv2D(BN, ReLU)		Flatten	
MaxPooling2D	2x2	GlobalAveragePooling2D		Dense(ReLU)	512
Conv2D(BN, ReLU)	3x3x512	Dense	10	Dropout	0.3
Dropout	0.4	Softmax		Dense	10
Conv2D(BN, ReLU)	3x3x512	where InvertedResidual is composed of Conv2D(BN, ReLU) DepthwiseConv2D(BN, ReLU) Conv2D(BN)		Softmax	
Dropout	0.4				
Conv2D(BN, ReLU)	3x3x512	(filters)			
MaxPooling2D	2x2				
Conv2D(BN, ReLU)	3x3x512				
Dropout	0.4				
Conv2D(BN, ReLU)	3x3x512				
Dropout	0.4				
Conv2D(BN, ReLU)	3x3x512				
MaxPooling2D	2x2				
Dropout	0.5				
Flatten					
Dense	512				
Dropout	0.5				
Dense	10				
Softmax					

5 DISCUSSION AND FUTURE WORK

Quantized deep learning models are very susceptible to adversarial attacks questioning the trust in them especially when dealing with safety-critical applications. In this work, we investigate to what extent a probabilistic approach can help mitigating adversarial examples. To this purpose, we build Stochastic-Shield a training-free multi-module shielding methodology which consists on adopting *input filtering* and *Monte Carlo Dropout* to provide models which are better calibrated and less sensitive to adversarial attacks.

Our experiments show that although each module contributes in mitigating attacks at different perturbation strength, when used together they consistently deliver the best performance on 8-bit quantized neural networks. This framework is easily parallelizable by running the independent Monte Carlo samples simultaneously creating a batch of the same input during inference or by running each of the forward pass on different devices of a federated system. This parallelization would guaranty a minimal delay in latency and offer robustness even on the tiniest devices.

An interesting avenue for future work could focus on studying how an approach like Stochastic-Shield would aid the robustness of models quantized with less than 8 bits (uniform 4-bit or mixed precision 8/4bit) models for even tinier devices that are used in IoT applications. Moreover, the findings in this work open a lot of opportunities in studying the effect of the level of randomness or dropout rate in adversarial mitigation.

REFERENCES

- [1] [n.d.]. TensorFlow Model Optimization. https://www.tensorflow.org/model_optimization/guide/quantization/training.
- [2] Alexander A Alemi, Ian Fischer, Joshua V Dillon, and Kevin Murphy. 2016. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410* (2016).
- [3] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. 2016. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565* (2016).
- [4] Tiago Azevedo, René de Jong, and Partha Maji. 2020. Stochastic-YOLO: Efficient Probabilistic Object Detection under Dataset Shifts. *arXiv preprint arXiv:2009.02967* (2020).
- [5] John Bradshaw, Alexander G de G Matthews, and Zoubin Ghahramani. 2017. Adversarial examples, uncertainty, and transfer testing robustness in Gaussian process hybrid deep networks. *arXiv preprint arXiv:1707.02476* (2017).
- [6] Nicholas Carlini and David Wagner. 2017. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 39–57.
- [7] Reuben Feinman, Ryan R Curtin, Saurabh Shintre, and Andrew B Gardner. 2017. Detecting adversarial samples from artifacts. *arXiv preprint arXiv:1703.00410* (2017).
- [8] Yarín Gal and Zoubin Ghahramani. 2015. Bayesian convolutional neural networks with Bernoulli approximate variational inference. *arXiv preprint arXiv:1506.02158* (2015).
- [9] Yarín Gal and Zoubin Ghahramani. 2016. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*. PMLR, 1050–1059.
- [10] Yiming Gan, Yuxian Qiu, Jingwen Leng, Minyi Guo, and Yuhao Zhu. 2020. Ptolemy: Architecture support for robust deep learning. In *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 241–255.
- [11] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014).
- [12] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. 2016. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236* (2016).
- [13] Mark Lee and Zico Kolter. 2019. On physical adversarial patches for object detection. *arXiv preprint arXiv:1906.11897* (2019).
- [14] Juncheng Li, Frank Schmidt, and Zico Kolter. 2019. Adversarial camera stickers: A physical camera-based attack on deep learning systems. In *International Conference on Machine Learning*. PMLR, 3896–3904.
- [15] Yingzhen Li and Yarín Gal. 2017. Dropout inference in Bayesian neural networks with alpha-divergences. In *International conference on machine learning*. PMLR, 2052–2061.
- [16] Ji Lin, Chuang Gan, and Song Han. 2019. Defensive quantization: When efficiency meets robustness. *arXiv preprint arXiv:1904.08444* (2019).
- [17] 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).
- [18] Maria-Irina Nicolae, Mathieu Sinn, Minh Ngoc Tran, Beat Buesser, Amrith Rawat, Martin Wistuba, Valentina Zantedeschi, Nathalie Baracaldo, Bryant Chen, Heiko Ludwig, et al. 2018. Adversarial Robustness Toolbox v1. 0.0. *arXiv preprint arXiv:1807.01069* (2018).
- [19] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua V Dillon, Balaji Lakshminarayanan, and Jasper Snoek. 2019. Can you trust your model’s uncertainty? Evaluating predictive uncertainty under dataset shift. *arXiv preprint arXiv:1906.02530* (2019).
- [20] Nicolas Papernot, Patrick McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. 2016. Distillation as a defense to adversarial perturbations against deep neural networks. In *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, 582–597.
- [21] Yash Sharma and Pin-Yu Chen. 2018. Bypassing feature squeezing by increasing adversary strength. *arXiv preprint arXiv:1803.09868* (2018).
- [22] Ilia Shumailov, Yiren Zhao, Robert Mullins, and Ross Anderson. 2018. The taboo trap: Behavioural detection of adversarial samples. *arXiv preprint arXiv:1811.07375* (2018).
- [23] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of machine learning research* 15, 1 (2014), 1929–1958.
- [24] 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).
- [25] Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. 2018. Robustness may be at odds with accuracy. *arXiv preprint arXiv:1805.12152* (2018).
- [26] Han Xu, Xiaorui Liu, Yaxin Li, and Jiliang Tang. 2020. To be Robust or to be Fair: Towards Fairness in Adversarial Training. *arXiv preprint arXiv:2010.06121* (2020).
- [27] Weilin Xu, David Evans, and Yanjun Qi. 2017. Feature squeezing: Detecting adversarial examples in deep neural networks. *arXiv preprint arXiv:1704.01155* (2017).
- [28] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. 2019. Self-attention generative adversarial networks. In *International conference on machine learning*. PMLR, 7354–7363.