



Title	Ramifications of Approximate Posterior Inference for Bayesian Deep Learning in Adversarial and Out-of-Distribution Settings
Authors(s)	Mitros, John (Ioannis), Pakrashi, Arjun, MacNamee, Brian
Publication date	2020-08-23
Publication information	Mitros, John (Ioannis), Arjun Pakrashi, and Brian MacNamee. "Ramifications of Approximate Posterior Inference for Bayesian Deep Learning in Adversarial and Out-of-Distribution Settings." Springer, August 23, 2020. https://doi.org/10.1007/978-3-030-66415-2_5 .
Conference details	The 16th European Conference on Computer Vision (ECCV 2020), Online Conference, 23-28 August 2020
Series	Lecture Notes in Computer Science, 12535
Publisher	Springer
Item record/more information	http://hdl.handle.net/10197/12575
Publisher's statement	The final publication is available at www.springerlink.com .
Publisher's version (DOI)	10.1007/978-3-030-66415-2_5

Downloaded 2026-05-02 00:29:52

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)



© Some rights reserved. For more information

Ramifications of Approximate Posterior Inference for Bayesian Deep Learning in Adversarial and Out-of-Distribution Settings

John Mitros, Arjun Pakrashi, and Brian Mac Namee

School of Computer Science, University College Dublin

`ioannis.mitros@insight-centre.org`

`{arjun.pakrashi,brian.macnamee}@ucd.ie`

Abstract. Deep neural networks have been successful in diverse discriminative classification tasks, although, they are poorly calibrated often assigning high probability to misclassified predictions. Potential consequences could lead to trustworthiness and accountability of the models when deployed in real applications, where predictions are evaluated based on their confidence scores. Existing solutions suggest the benefits attained by combining deep neural networks and Bayesian inference to quantify uncertainty over the models' predictions for ambiguous data points. In this work we propose to validate and test the efficacy of likelihood based models in the task of out of distribution detection (OoD). Across different datasets and metrics we show that Bayesian deep learning models indeed outperform conventional neural networks but in the event of minimal overlap between in/out distribution classes, even the best models exhibit a reduction in AUC scores in detecting OoD data. We hypothesise that the sensitivity of neural networks to unseen inputs could be a multi-factor phenomenon arising from the different architectural design choices often amplified by the curse of dimensionality. Preliminary investigations indicate the potential inherent role of bias due to choices of initialisation, architecture or activation functions. Furthermore, we perform an analysis on the effect of adversarial noise resistance methods regarding in and out-of-distribution performance when combined with Bayesian deep learners.

Keywords: OoD detection, Bayesian deep learning, uncertainty quantification, generalisation, anomalies, outliers

1 Introduction

Anomaly detection is concerned with the detection of unexpected events. The goal is to effectively detect rare or novel patterns which neither comply to the norm nor follow a specific trend present in the existing data distribution. The task of anomaly detection is often described as open category classification (Liu et al., 2018) and anomalies can be referred to as *outliers*, *novelties*, *noise deviations*, *exceptions* or *out-of-distribution* (OoD) examples depending on the

context of the application (Shafaei et al., 2018). Anomaly detection raises significant challenges among which some stem from the nature of generated data arising from different sources e.g. multi-modality from different domains such as manufacturing, industrial, health sectors etc. (Chalapathy and Chawla, 2019; Hendrycks et al., 2018a; Choi et al., 2018).




Fig. 1: An example of the OoD detection problem. At each angle $0^\circ \leq \theta \leq 180^\circ$ we record the image classifier’s response. The classifier is able to predict the true label, *dog*, with approximately 100% confidence for an angle $\theta \leq 30^\circ$ degrees, while it fails to predict the correct label for any rotation angle $\theta > 30^\circ$. (Inspired by an image in (Gal and Ghahramani, 2016).)

In this paper we are concerned with a variant of anomaly detection typically referred to as out-of-distribution (OoD) data detection. Figure 1 shows an illustrative example of an OoD problem. A model is trained to classify images among the following classes: $\{horse, dog, cat, bird, automobile\}$. The training procedure involved augmentation, including random rotations. When a new image \mathbf{x} is provided containing any of these entities, the model should be able to predict the result and assign a probability to each prediction. In the case of Figure 1 we record the classifier’s response $\hat{y} = f_{\mathbf{W}}(\mathbf{x})$ while rotating an image belonging to the class *dog* in the range of $0^\circ \leq \theta \leq 180^\circ$ degrees. From Figure 1 it is evident that the classifier is able to predict the true label y , which in this instance is *dog*, with approximately 100% confidence for angles y up to $\theta \leq 30^\circ$ degrees, while it fails to predict the correct label for any rotation angle $\theta > 30^\circ$. Moreover, in the cases where the model misclassifies the rotated image, it can assign a very high confidence to the incorrect class membership. For example at $\theta = 40^\circ$ and $\theta = 80^\circ$ the model classifies the rotated image to be a *cat* and an *automobile* respectively with very high confidence.

This uncovers two related issues. First, the example demonstrates the incapability of the model to detect OoD samples previously unseen by it. Second, the classifier misclassifies predictions with high confidence. For instance, the image of the dog was classified as a *cat* or an *automobile* (i.e. black and orange

curves in Figure 1) with high confidence across different rotation angles. This can potentially lead to problems related with trustworthiness, transparency and privacy of the prediction (Schulam and Saria, 2019). This phenomenon is not new (Szegedy et al., 2013), and since being identified has captivated significant research effort to understand its underpinnings and provide related solutions.

Many of the proposed solutions have their roots in different fields such as differential privacy, information theory, robust high dimensional statistics, robust control theory and robust optimisation. Recently attention has been concentrated to methods that provide a principled approach to quantifying uncertainty through Bayesian deep learning (Gal and Ghahramani, 2016; Maddox et al., 2019; Malinin and Gales, 2018; Grathwohl et al., 2019). In spite of their sophistication, questions remain unanswered regarding the effectiveness of these approaches.

In addition to the OoD prediction problem, there is another issue with which deep neural networks struggle: adversarial noise (Akhtar and Mian, 2018; Ozdag, 2018). For a trained deep neural network model, adversarial noise can be added to an input image in a way such that the model classifies it to a different class from its true one. If this noise is chosen carefully, then the image can appear unchanged to a human observer. As many deep learning based image recognition systems are being deployed in real life, such adversarial attacks can result in serious consequences. Therefore, developing strategies to mitigate against them is an important open problem.

In this paper we describe an empirical evaluation of the effectiveness of Bayesian deep learning for OoD detection on a number of different methods using a selection of image classification datasets. This experiment benchmarks the performance of current state-of-the-art methods, and is aligned in principle to the inquisitive nature of Goldblum et al. (2020). Examining why such powerful models exhibit properties of miscalibration and an incapacity to detect OoD samples will allow us to better understand the major implications in critical domain applications since it implies that their predictions cannot be trusted (Hill, 2019; Kumar et al., 2019; Alemi et al., 2018; Guo et al., 2017; Marin et al., 2012). Furthermore, we provide discussions into what we consider to be critical components attributing to this phenomenon which suggest directions for future work, inspired by evidence in recent literature. Finally, we investigate the potential benefit of using adversarial noise defence mechanisms for the OoD detection problem as well as the in-sample adversarial attacks for Bayesian neural networks.

The main contributions of this work include:

- A benchmark study demonstrating that Bayesian deep learning methods do indeed outperform point estimate neural networks in OoD detection.
- An empirical validation showing that from the considered approaches in this work, *Stochastic Weight Averaging of Gaussian Samples* (SWAG) (Maddox et al., 2019) is most effective (it achieves the best balance between performance in the OoD detection task and the original image classification task).

- We demonstrate that OoD detection is harder when there exists even minimal overlap between in-distribution and out-of-distribution examples and outlier exposure (Hendrycks et al., 2019) might not be beneficial in such scenarios.
- We investigate the efficacy of adversarial defence techniques and noise robustness showing that randomised smoothing can be quite effective in combination with Bayesian neural networks.

This paper is structured as follows: Section 2 describes related work; Section 3 describes an experiment to benchmark the ability of Bayesian neural networks to detect OoD examples; Section 4 investigates the effects of adversarial noise defence mechanisms on classification models and the ability of Bayesian neural networks with and without these defences to identify adversarial examples; finally, Section 5 concludes the paper.

2 Related Work

The goal of anomaly detection (Chalapathy and Chawla, 2019), is to detect any inputs that do not conform to in-distribution data—referred as *out-of-distribution data* (OoD). While there are approaches that are solely targeted at the detection of anomalies (for example (Chalapathy and Chawla, 2019; Erfani et al., 2016; Marchi et al., 2015)), classification models that can identify anomalous inputs as part of the classification process are attractive. For instance, it is desirable that a model trained for image classification (such as that described in Figure 1) would be capable of not only classifying input images accurately, but also identifying that does not belong to any of the classes that it was trained to recognise. Not only can this be useful for detecting ambiguous inputs (Schulam and Saria, 2019), but it can also be useful for detecting more insidious adversarial examples (Carlini and Wagner, 2017). This type of outlier detection as part of the classification process is most commonly referred to as *OoD detection*.

The majority of OoD detection approaches utilise the uncertainty of classifier’s outputs as an indicator to identify whether an input instance is OoD. Such methods are mostly based on density estimation techniques utilising generative models or ensembles. For example, Choi et al. (2018) presented an ensemble of generative models for density-based OoD detection by estimating epistemic uncertainty of the likelihood, while Hendrycks and Gimpel (2017) focused on detecting malformed inputs, \mathbf{x} , from their conditional distribution $p(y|\mathbf{x})$, assigning OoD inputs lower confidence than normal inputs. Hybrid models for uncertainty quantification are also common. For example, the work of Zhu and Laptev (2017) that describes an end-to-end system using recurrent networks and probabilistic models. The overall architecture is composed of an encoder-decoder recurrent model coupled with a multilayer perceptron (MLP) on top of the encoder allowing to estimate the uncertainty of predictions.

The recent resurgence of interest in combining Bayesian methods with deep neural networks (see (Wang and Yeung, 2016)) has given rise to different training mechanisms for classification models that output local likelihood estimates of

class membership, providing better estimates of classification uncertainty compared to standard deep neural network (DNN) approaches producing point-estimates. Models whose outputs better represent classification uncertainty can be directly applied to the OoD detection problem—intuitively predictions with high uncertainty suggest OoD examples—and recent approaches that use a Bayesian formulation have claimed to be effective at this. One early example is *MC-Dropout* (Gal and Ghahramani, 2016) which views the use of dropout at test time as approximate Bayesian inference. Monte Carlo sampling from the model at test time for different subsets of the weights is used to achieve this. MC-Dropout is based on prior work (Damianou and Lawrence, 2013) which established a relationship between neural networks with dropout and Gaussian Processes (GP). *Stochastic Weight Averaging of Gaussian Samples* (SWAG) (Maddox et al., 2019), being an extension of *stochastic weight averaging* (SWA) (Izmailov et al., 2018), is another example of an approximate posterior inference approach that has been shown to be useful for OoD detection. In SWAG, the weights of a neural network are averaged during different SGD iterations (Mandt et al., 2017) which allows approximates of output distributions to be calculated.

The *Dirichlet Prior Network* (DPN) (Malinin and Gales, 2018) based on recent work focuses explicitly on detecting OoD instances. While alternative approaches using Bayesian principles aim at constructing an implicit conditional distribution with certain desirable properties—e.g., appropriate choice of prior and inference mechanism—DPN strives to explicitly parameterise a distribution using an arithmetic mixture of Dirichlet distributions. DPN also introduces a separate source of uncertainty through an OoD dataset used at training time as a uniform prior over the OoD data (i.e. similar to outlier exposure (Hendrycks et al., 2019)). Training a DPN model involves optimising the weights of a neural network by minimising the Kullback-Leibler Divergence between in-distribution and out-of-distribution data, where in-distribution data has been modelled with a sharp prior Dirichlet distribution while out-of-distribution data is modelled with a flat prior Dirichlet distribution.

The *Joint Energy Model* (JEM) (Grathwohl et al., 2019) is an alternative approach that uses generative models to improve the calibration of a discriminative model by reinterpreting an existing neural network architecture as an Energy Based Model (EBM). Specifically they use an EBM (LeCun et al., 2006), which is a type of generative model that parameterises a probability density function using an unnormalised log-density function. It was also shown in (Grathwohl et al., 2019) that JEM is effective on OoD problems.

All of these methods perform OoD detection by converting their class membership outputs into a score indicating the likelihood of an out-of-distribution input instance and then applying a threshold to this score. There are four common approaches to calculating these scores:

- *Max probability*: This is the maximum value of the softmax outputs. If $\hat{\mathbf{y}} = [y_1, y_3, \dots, y_k]$ is the prediction, then $\max_k p(y_k | \mathbf{x})$ is used as the score.
- *Entropy*: Calculates information entropy $\mathbf{H}(Y) = -\sum_{k=1}^K p(\hat{y}_k) \log p(\hat{y}_k)$ over the class memberships predicted by a model (Malinin and Gales, 2018).

- *Mutual Information*: $\mathbb{E}_{p(\mathbf{x},y)}[\frac{p(\mathbf{x},y)}{p(\mathbf{x})p(y)}]$ measures the amount of information obtained about a random variable X by observing some other random variable Y (Tschannen et al., 2019). Here X could describe the entropy of the predictions and Y the entropy of a Dirichlet mixture encompassing the final predictions.
- *Differential Entropy*: Also known as the continuous entropy is defined as follows $\mathbf{h}(Y) = -\int p(\hat{y}) \log p(\hat{y}) dx$. This is used to measure distributional uncertainty between in-out distributions.

These techniques can also be applied to the outputs of a simple deep neural network to perform OoD detection in the same manner.

Despite the resurgence of Bayesian methods for deep learning, there are still a number of challenges that need to be addressed before fully harnessing their benefits for OoD detection. Foong et al. (2019) draws attention to the poorly understood approximations due to computational tractability upon which Bayesian neural networks rely, indicating that common approximations such as factorised Gaussian assumption and MC-Dropout lead to pathological estimates of predictive uncertainty. Posteriors obtained via mean field variation inference (MFVI) (Blei et al., 2017) are unable to represent uncertainty between data clusters but have no issue representing uncertainty outside the data clusters. Similarly, MC-Dropout cannot represent uncertainty in-between data clusters, being more confident in the midpoint of the clusters rather than the centres. In addition, Wenzel et al. (2020) questioned the efficacy of accurate posterior approximation in Bayesian deep learning, demonstrating through MCMC sampling that the predictions induced by a Bayes posterior systematically produced worse results than point estimates obtained from SGD.

Bias in models has been also observed due to the extreme over-parameterised regime of neural networks. Predictions in the vicinity of the training samples are regularised by the continuity of smoothness constraints, while predictions away from the training samples determine the generalisation performance. The different behaviours reflect the inductive bias of the model and training algorithm. Zhang et al. (2020) showed that shallow networks learn to generalise but as their size increases they are more biased into memorising training samples. Increasing the receptive field of convolutional networks introduces bias towards memorising samples, while increasing the width and number of channels does not. This seems to be in accordance with equivalent findings from Azulay and Weiss (2019) identifying that subsampling operations in combination with non-linearities introduce a bias, producing representations that are not shiftable therefore causing the network to lose its invariance properties. Finally, Fetaya et al. (2020) showed that conditional generative models have the tendency to assign higher likelihood to interpolated images due to the class-unrelated entropy indicating that likelihood-based density estimation and robust classification might be at odds with each other.

Prior work has focused on investigating the sensitivity of point estimate DNNs for OoD detection on corrupted inputs as well as comparing local vs global methods (i.e. ensembles) (Hendrycks and Gimpel, 2016; Chen et al., 2020;

Hendrycks and Dietterich, 2019; Hendrycks et al., 2018b). In this work, however, we focus on investigating the effectiveness of local likelihood methods when inputs at test time are considered corrupted with noise by an adversary or simply constitute OoD. In spite of the recent emergence of multiple approaches to the OoD detection problem based on Bayesian methods, and the uncertainty surrounding their effectiveness, no objective benchmark regarding their relative performance at this task currently exists in the literature to best of our knowledge. This paper provides such a benchmark.

To "fool" a deep neural network, examples can be specifically generated in such a way that a specific image is forced to be assigned to a class, from which it does not belong. This is done by adding noise such that it looks identical to the human eye, but it either 1) classifies the example to a specific target class to which the example does not belong (targetted), or 2) to any class other than its original one (untargetted). Such induced noise is called adversarial noise, and can be generated using several methods (Ozdag, 2018; Akhtar and Mian, 2018). As deep learning based vision systems are being deployed in real world applications, consequently, research in defending against adversarial noise is increasing in popularity and importance. For this pupose, there have been developed several adversarial noise defence mechanisms in the literature (Pang et al., 2020; Xiao et al., 2020; Cohen et al., 2019).

3 Experiment 1: Benchmarking OoD Detection

In this experiment, using a number of well-known image classification datasets, we evaluate the ability of four state-of-the-art methods based on Bayesian deep learning—DPN, MC-Dropout, SWAG, and JEM—to detect OoD examples, and compare this with the performance of a standard deep neural network (DNN).

3.1 Experimental Setup

We use a 28 layers wide and 10 layers deep WideResNet (Zagoruyko and Komodakis, 2016) as the DNN model, and this is also the base model used by the other approaches. Whenever it was appropriate and possible for specific datasets, we utilised pre-trained models provided by the original authors to avoid any discrepancies in our results. For instance, (Grathwohl et al., 2019) made available a pre-trained JEM model for the CIFAR-10 dataset, and we use this in our experiments rather than training our own. For the remaining models we trained each for 300 epochs using a validation set for hyper-parameter tuning and rolling back to the best network to avoid overfitting.

The optimiser used during the experiments was Stochastic Gradient Descent (SGD) (Goodfellow et al., 2016) with momentum set to 0.9 and weight decay in the range $[3e^{-4}, 5e^{-4}]$. Additionally, every dataset was split into three distinct sets $\{train, validation, test\}$ with augmentations such as random rotation, flip, cropping and pixel distortion applied on the training set. The source code and the pre-trained models are available online.

A significant distinction between DPN and the other approaches used is that DPN uses an idea similar to outlier exposure (Hendrycks et al., 2019) and it requires two datasets during training: one to represent the in-distribution data and the other to simulate out-of-distribution data.

Five well-known image classification datasets are used in this experiment: *CIFAR-10* (Krizhevsky et al., 2009), *CIFAR-100* (Krizhevsky et al., 2009), *SVHN* (Netzer et al., 2011), *FashionMNIST* (Xiao et al., 2017), and *LSUN* (Yu et al., 2015) (We use only the *bedroom* scene in this experiment only as OoD data for testing.)

To comprehensively explore the performance of the different models we perform experiments using each dataset to train a model and use all other datasets to measure the ability of the model to perform OoD detection. The ability of each model to perform the image classification task it was trained for, was first evaluated using the test set associated with each dataset. All datasets have balanced class distributions, therefore we use classification accuracy to measure this performance.

To measure the ability of models to recognise OoD examples we make predictions for the test portion of the dataset used to train the model (in-distribution data) and then also make predictions for the test portion of three other datasets (out-of-distribution data): CIFAR-100, SVHN, and LSUN. This means that for each training set we have three different evaluations of OoD detection effectiveness. For example, when SVHN is used as the in-distribution training set, CIFAR-10, LSUN, and CIFAR-100 are used as the out-of-distribution test sets. One of the available out-of-distribution datasets is selected for use at training time for DPN which requires this extra data. Only the training portions of this dataset is used for this purpose, while the test portions are used for evaluation.

The predictions provided by the models are converted into OoD scores using the four alternative approaches described in Section 2: *max probability*, *entropy*, *mutual information*, and *differential entropy*. To avoid having to set detection thresholds on these scores we measure the separation between the scores generated for instances on the in-distribution and out-of-distribution test sets using the area under the curve (AUC-ROC) based on the four different approaches for calculating scores. We do this individually for each approach to generate the AUC scores.

3.2 Results & Discussion

Table 1 shows the overall performance of each model for the basic in-distribution image classification tasks they were trained for, measured using classification accuracy. These results indicate that, for most cases, the Bayesian methods are performing comparably to the DNN baseline (the cases where this is not true will be discussed shortly).

Table 2 shows the results of the OoD detection experiments. The scores represent AUC-ROC scores calculated using entropy of the model outputs for OoD detection. Values inside parenthesis indicate the percentage improvement with respect to DNN, which is treated as a baseline. The \uparrow indicates an improvement and the \downarrow indicates a degradation. The last row of the table shows the average

Table 1: Accuracy of models on in-distribution dataset image classification tasks.

Model	CIFAR-10	SVHN	FashionMNIST	CIFAR100
DNN	95.06	96.67	95.27	77.44
DPN	88.10	90.10	93.20	79.34
MC-Dropout	96.22	96.90	95.40	78.39
SWAG	96.53	97.06	93.80	78.61
JEM	92.83	96.13	83.21	77.86

percentage improvement across the dataset combinations for each approach with respect to the DNN baseline. Similar tables based on OoD scores calculated using max probability, mutual information, and differential entropy are available in the supplementary material. Table 3 summarises the results from these tables and presents the average performance increase with respect to the DNN baseline based on AUC-ROC calculated using each OoD scoring method.

Table 2: Out-of-distribution experiment results. Scores are Entropy based AUC-ROC in percentage. The values in bracket are % improvement of the corresponding algorithm wrt. DNN, taken as a baseline. An \uparrow indicates improvement and \downarrow degradation wrt. the baseline (DNN). The asterisks (*) next to each dataset indicates out-distribution datasets used to train DPN.

Data		(baseline)	Entropy AUC-ROC score (% gain wrt. baseline)			
In-distribution	OoD	DNN	DPN	MC-Dropout	SWAG	JEM
CIFAR-10	CIFAR-100*	86.27	85.60 (\downarrow 0.78%)	89.92 (\uparrow 4.23%)	91.89 (\uparrow 6.51%)	87.35 (\uparrow 1.25%)
	SVHN	89.72	98.90 (\uparrow 10.23%)	96.25 (\uparrow 7.28%)	98.62 (\uparrow 9.92%)	89.22 (\downarrow 0.56%)
	LSUN	88.83	83.30 (\downarrow 6.23%)	92.04 (\uparrow 3.61%)	95.12 (\uparrow 7.08%)	89.84 (\uparrow 1.14%)
SVHN	CIFAR-100	93.19	99.10 (\uparrow 6.34%)	94.33 (\uparrow 1.22%)	95.97 (\uparrow 2.98%)	92.34 (\downarrow 0.91%)
	CIFAR-10*	94.58	99.60 (\uparrow 5.31%)	94.97 (\uparrow 0.41%)	96.03 (\uparrow 1.53%)	92.85 (\downarrow 1.83%)
FashionMNIST	LSUN	92.97	99.70 (\uparrow 7.24%)	93.31 (\uparrow 0.37%)	95.71 (\uparrow 2.95%)	91.82 (\downarrow 1.24%)
	CIFAR-100	91.20	99.50 (\uparrow 9.10%)	93.75 (\uparrow 2.80%)	96.19 (\uparrow 5.47%)	62.79 (\downarrow 31.15%)
	CIFAR-10*	94.59	99.60 (\uparrow 5.30%)	96.06 (\uparrow 1.55%)	94.28 (\downarrow 0.33%)	64.76 (\downarrow 31.54%)
CIFAR-100	LSUN	93.34	99.80 (\uparrow 6.92%)	97.40 (\uparrow 4.35%)	99.05 (\uparrow 6.12%)	65.38 (\downarrow 29.96%)
	CIFAR-10	78.25	85.15 (\uparrow 8.82%)	80.70 (\uparrow 3.13%)	84.92 (\uparrow 8.52%)	77.64 (\downarrow 0.78%)
	SVHN*	81.52	92.64 (\uparrow 13.64%)	85.59 (\uparrow 4.99%)	94.16 (\uparrow 15.51%)	81.22 (\downarrow 0.37%)
Avg % improvement			86.38 (\uparrow 11.86%)	76.58 (\downarrow 0.83%)	87.22 (\uparrow 12.95%)	77.54 (\uparrow 0.41%)
			(\uparrow 6.48%)	(\uparrow 2.76%)	(\uparrow 6.60%)	(\downarrow 7.96%)

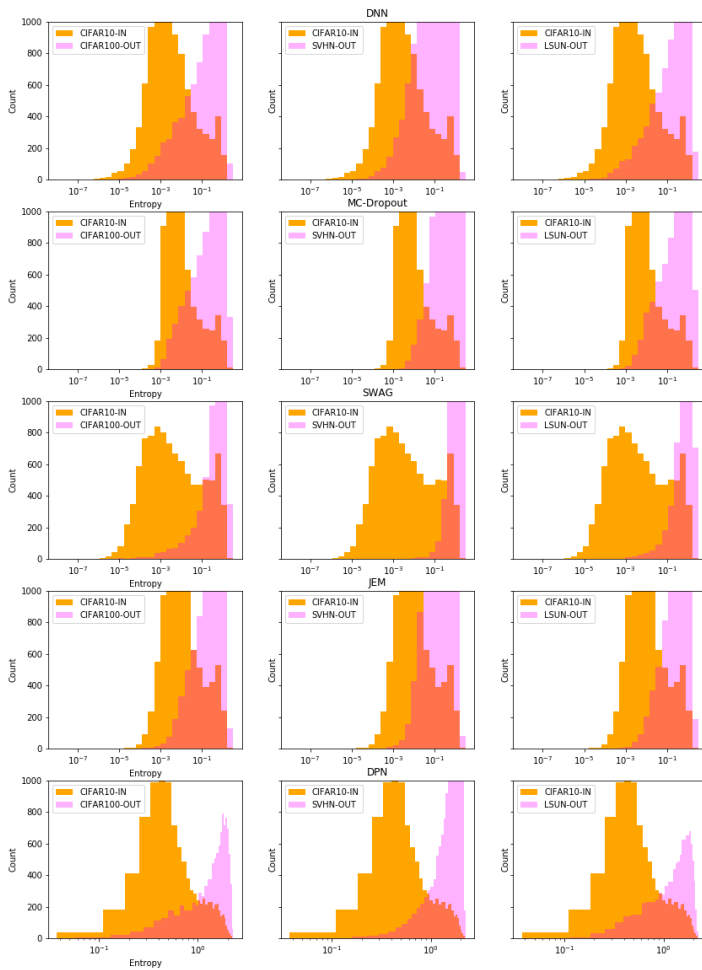


Fig. 2: Histograms of in-out distribution across models and datasets. The x-axis represents entropy in log scale.

It is evident from Table 2 that most of the Bayesian methods— DPN, MC-Dropout and SWAG—almost always improve OoD detection performance over

Table 3: Percent performance increase wrt. baseline (DPN) for all the evaluation score. The values in parenthesis are the relative ranks of improvement for the corresponding evaluation score. Last row indicates the improvement ranking of the algorithms, averaged through the evaluation scores.

	DPN	MC-Dropout	SWAG	JEM
Max prob.	5.895% (2)	2.618% (3)	6.067% (1)	-8.350% (4)
Mutual Info.	7.262% (1)	6.243% (2)	5.337% (3)	-8.034% (4)
Entropy	6.480% (2)	6.480% (3)	6.601% (1)	-7.960% (4)
Diff. Entropy	9.232% (1)	4.386% (3)	5.495% (2)	-10.217% (4)
Avg. rank	(1.50)	(2.75)	(1.75)	(4.00)

the DNN baseline. Interestingly, in our experiment JEM consistently performed poorly with respect to the DNN baseline. We attribute this mostly to a failure of robust selection of hyper-parameter combinations. Despite our best efforts we were not able to achieve classification accuracy (above 83%) on the FashionM-NIST dataset using this approach, and the OoD performance suffered from this. However, even for the datasets in which the JEM models performed well for the image classification task (SVHN and CIFAR-100) its use did not lead to an increase in OoD detection performance ¹. Overall we found that both JEM and DPN can be quite unstable and extremely sensitive to hyper-parameter choice.

The results in Table 3 indicate that among the different evaluation metrics, DPN performs best for OoD detection, with SWAG following closely behind. MC-Dropout is third, but comparatively worse than the other two. All three Bayesian methods on average increase OoD detection performance with respect to the DNN baseline regardless of the approach used to calculate OoD scores. DPN used in combination with differential entropy is especially effective. For other methods the best performance is achieved using entropy as the score for the OoD detection task. This indicates that overall, Bayesian methods improve the performance on OoD detection, with DPN and SWAG achieving similar performance (although JEM does not achieve improvements over the DNN baseline).

However, it is worth noting from Table 1 that the DPN models do not perform well on the base image classification tasks compared to the other approaches. Therefore, we conclude that of the approaches compared in this experiment SWAG is the most effective as it achieves strong OoD detection performance, without compromising basic in-distribution classification performance. It also does not require the use of an OoD dataset at training time. Although, if the objective is achieving the best possible OoD detection performance then DPN should be considered instead.

¹ Even after suggestions kindly provided by (Grathwohl et al., 2019) we found that the model diverged instead of closely approximating its original results. There was another instance we noticed a degradation in performance compared to the original results, it was in the case of DPN when trained with CIFAR-10 for in-distribution and CIFAR-100 for OoD (Malinin and Gales, 2018), due to the overlap between some classes in CIFAR-10 & CIFAR-100.

It is also worth mentioning that the performance of models used in combination with the four approaches to calculating OoD scores (described in Section 2), varies slightly. The *max probability*, *entropy* and *mutual information* approaches are sensitive to class overlap between in-distribution and out-of-distribution data, while *differential entropy* seems to be more sensitive to the target precision parameter of the Dirichlet distribution, controlling how it is concentrated over the simplex of possible outputs.

4 Experiment 2: Robustness Analysis

Having identified in the previous experiment that Bayesian neural network models can be utilised for OoD detection, we proceed by verifying whether they can withstand untargeted adversarial attacks. We design a new experiment including the usual DNN acting as a control baseline and two Bayesian models (*MC-Dropout* and *SWAG*), in order to answer the following questions:

1. Are Bayesian neural networks capable of detecting adversarial examples?
2. Could simple defence mechanisms $\{Top-k, Randomised Smoothing, MMLDA\}$, outperform BNN or improve their overall performance?

4.1 Experimental Setup

First, we evaluated the ability of each model against untargeted adversarial examples generated with *Projected Gradient Descent (PGD)* (Madry et al., 2017; Ozdag, 2018) with the following hyper-parameter values $\epsilon = 0.1$ and $\alpha = 0.01$ for ten iterations. Second, we introduced three recently proposed defence techniques against adversarial examples and evaluated their efficacy on the (i) *clean test set*, (ii) *adversarially corrupted test set*, and finally on the (iii) *OoD detection task*. The defence methods evaluated against adversarial examples are: *Randomised Smoothing* (Cohen et al., 2019), *Sparsify k-winners take all (Top-k)* (Xiao et al., 2020), and *Max-Mahalanobis Linear Discriminant Analysis (MMLDA)* (Pang et al., 2020). We use only the CIFAR-10 dataset in these experiments.

In order to train models with MMLDA and RandSmooth we utilised the same number of hyper-parameters and values as they are depicted in the original papers. For instance, for MMLDA we utilised a variance of value 10 for CIFAR-10 in order to compute the Max-Mahalanobis centres on \mathbf{z} , where \mathbf{z} denotes the features regarding the penultimate layer of each model $\mathbf{z} = f(x)$. We should also notice the possibility for randomised smoothing to achieve better results at the expense of increased prediction time. A key combination for obtaining successful results is the number of samples $n = 55$, or randomised copies for each x in the test set, together with the confidence level $\alpha = 0.001$ (i.e. there is a 0.001 probability that the answer will be wrong) and standard deviation $\sigma = 0.56$ of the isotropic Gaussian noise.

4.2 Results & Discussion

We first illustrate the effect of adversarial noise on the ability of different model types to perform the underlying CIFAR-10 classification task in the presence of adversarial noise. Table 4 shows the performance of each model type on the basic CIFAR-10 classification problem on a clean dataset with no adversarial noise. We see that, apart from randomise smoothing, the addition of noise defence strategies does not negatively affect performance at this task. Randomised smoothing had a negative impact on the in-distribution performance across all methods indicating that robustness might be at odds with accuracy (Tsipras et al., 2019). MMLDA and Top-k seem to have a positive effect on DNN and MC-Dropout but not on SWAG.

Table 4: Accuracy on CIFAR-10 clean test set for each defence technique. Values in parenthesis indicate the percentage of abstained predictions.

Model	No defence	Top-k	RandSmooth	MMLDA
DNN	95.06	94.52	86.25 (13.00)	95.18
MC-Dropout	96.22	94.43	86.98 (13.39)	95.21
SWAG	96.53	91.73	79.68 (20.32)	91.30




Fig. 3: Evaluation of defence techniques vs non-defence on clean test set of CIFAR-10.

Table 5 summarises the performance of the same models when adversarial noise is added to the test set. The first column in Table 5 corresponds to the accuracy on the adversarially corrupted test set for each model without applying any adversarial defence mechanism. Lower values indicate that the adversarial attack was successful and managed to force the model to misclassify instances

with high confidence. When Randomised smoothing is used, the values in parentheses denote the percentage of abstained predictions on the total number of instances from the test set. The impact of adversarial noise is evident in these results. When there is no defence mechanism used then classification accuracy plummets to less than 2% for all models. The addition of defence techniques, especially MMLDA and randomised smoothing indicate improvement in performance.

Table 5: Accuracy on CIFAR-10 test set corrupted with adversarial noise.

Model	No Defence	Top-k	RandSmooth	MMLDA
DNN	1.15	11.41	62.65 (27.75)	45.71
MC-Dropout	1.94	7.28	88.85 (17.98)	47.90
SWAG	0.55	0.79	36.80 (20.97)	47.95

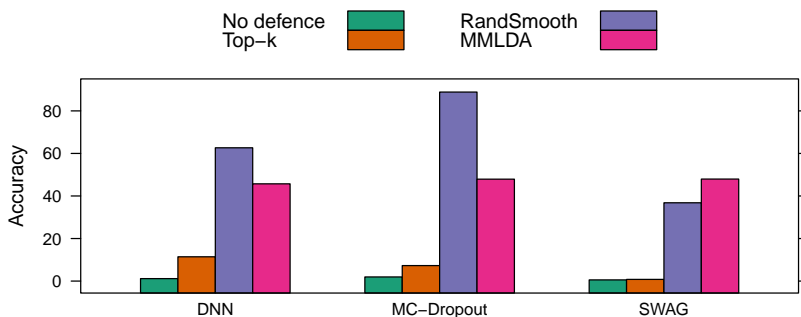


Fig. 4: Evaluation of defence techniques vs non-defence on the adversarially corrupted test set of CIFAR-10.

In addition, Table 6, shows that model predictions can be used to identify adversarial examples in the same manner that OoD examples were identified by measuring the entropy over the final predictions. It compares the entropy scores generated for an in-distribution test set versus those generated for adversarial examples and uses AUC-ROC scores to measure the ability of a model to distinguish them. First, notice that in line with the results in the previous section the Bayesian models perform better at this task than the baseline DNN when no defence against adversarial examples is used. SWAG, in particular, seems quite effective. When adversarial defence mechanisms are used, the performance increases dramatically. Top-k in particular seems effective in this occasion.

Table 6: Out-of-distribution detection results for all defence methods on clean vs adversarially corrupted CIFAR-10. Scores represent entropy based AUC-ROC in percentage.

Data		Entropy AUC-ROC score		
In-distribution	OoD	DNN	MC-Dropout	SWAG
CIFAR-10/No defence	CIFAR-10 Adv.	13.67	36.54	73.81
CIFAR-10/Topk	CIFAR-10 Adv.	99.83	99.83	98.90
CIFAR-10/MMLDA	CIFAR-10 Adv.	53.15	85.65	70.38
CIFAR-10/RandSmooth	CIFAR-10 Adv.	62.68	57.06	48.29

Finally, Table 7 shows whether the defence techniques provide any improvements in regard to the OoD detection task in which the defence techniques are incorporated within the models. The values represent AUC-ROC scores computed based on entropy of the predictions obtained from each model after having introduced the different defence techniques to each model. Again the top-k and MMLDA approaches improve performance over the no defence option.

Table 7: Out-of-distribution detection results. Scores represent entropy based AUC-ROC in percentage.

Data		Entropy AUC-ROC score		
In-distribution	OoD	DNN	MC-Dropout	SWAG
CIFAR-10/Topk	CIFAR-100	90.59	90.45	84.72
	SVHN	91.20	92.62	94.61
	LSUN	92.42	92.39	89.81
CIFAR-10/MMLDA	CIFAR-100	99.87	99.24	79.78
	SVHN	99.74	99.75	84.56
	LSUN	99.93	99.67	81.72
CIFAR-10/RandSmooth	CIFAR-100	62.91	69.33	60.21
	SVHN	42.62	63.22	66.48
	LSUN	61.94	69.10	61.23

5 Conclusion

This work investigates if Bayesian methods improve OoD detection for deep neural networks used in image classification. To do so, we investigated four recent methods and compared them with a baseline point estimate neural network on four datasets.

Our findings show that the Bayesian methods overall improve performance at the OoD detection task over the DNN baseline. Of the considered methods

examined DPN performed best, with SWAG following closely behind. However, DPN requires additional data during training, is sensitive to hyper-parameter tuning, and led to worse in-distribution performance. Therefore, we conclude that SWAG seems more effective of the approaches examined considering both in and out-distribution performance.

Furthermore we showed that, despite being better than simple DNN models on OoD, Bayesian neural networks do not possess the ability to cope with adversarial examples. Although adversarial defence techniques overall degrade accuracy on the clean test set, at the same time, they robustify against adversarial examples across models and occasionally improve OoD detection. Randomised smoothing performed the best, followed by MMLDA. In the case of the OoD detection task it seems that Top-k slightly improved performance for DNN while MMLDA improved performance for DNN and MC-Dropout.

The broader scope of our work is not only to identify whether Bayesian methods could be used for OoD detection and to mitigate against adversarial examples, but also to understand how much of the proposed methods are affected by the inductive bias in the choices of model architecture and objective functions. For our future work we look forward to examine these components and their inherent role regarding the sensitivity of models against outliers.

Acknowledgements

This research was supported by Science Foundation Ireland (SFI) under Grant number SFI/12/RC/2289_P2.

Bibliography

- Si Liu, Risheek Garrepalli, Thomas Dietterich, Alan Fern, and Dan Hendrycks. Open category detection with PAC guarantees. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 3169–3178, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.
- Alireza Shafaei, Mark Schmidt, and James J. Little. Does Your Model Know the Digit 6 Is Not a Cat? A Less Biased Evaluation of "Outlier" Detectors. *ArXiv: 1809.04729*, 2018.
- Raghavendra Chalapathy and Sanjay Chawla. Deep Learning for Anomaly Detection: A Survey. *ArXiv190103407*, 2019.
- Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep Anomaly Detection with Outlier Exposure. *ArXiv181204606*, 2018a.
- Hyunsun Choi, Eric Jang, and Alexander A. Alemi. WAIC, but Why? Generative Ensembles for Robust Anomaly Detection. *ArXiv: 1810.01392*, 2018.
- Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. *arXiv: 1506.02142*, 2016.
- Peter Schulam and Suchi Saria. Can You Trust This Prediction? Auditing Pointwise Reliability After Learning. In *Proc. of Artificial Intelligence and Statistics*, 2019.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv: 1312.6199*, 2013.
- Wesley J Maddox, Pavel Izmailov, Timur Garipov, Dmitry P Vetrov, and Andrew Gordon Wilson. A simple baseline for bayesian uncertainty in deep learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d' Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 13153–13164. Curran Associates, Inc., 2019.
- Andrey Malinin and Mark Gales. Predictive uncertainty estimation via prior networks. In *Advances in Neural Information Processing Systems*, pages 7047–7058, 2018.
- Will Grathwohl, Kuan-Chieh Wang, Jörn-Henrik Jacobsen, David Duvenaud, Mohammad Norouzi, and Kevin Swersky. Your classifier is secretly an energy based model and you should treat it like one. *arXiv preprint arXiv:1912.03263*, 2019.
- Naveed Akhtar and Ajmal Mian. Threat of adversarial attacks on deep learning in computer vision: A survey. *IEEE Access*, 6:14410–14430, 2018.
- Mesut Ozdag. Adversarial attacks and defenses against deep neural networks: a survey. *Procedia Computer Science*, 140:152–161, 2018.
- Micah Goldblum, Jonas Geiping, Avi Schwarzschild, Michael Moeller, and Tom Goldstein. Truth or backpropaganda? an empirical investigation of deep learning theory. In *International Conference on Learning Representations*, 2020.

- Brian Hill. Confidence in Belief, Weight of Evidence and Uncertainty Reporting. In *International Society for Imprecise Probability: Theories and Applications*, 2019.
- Ananya Kumar, Percy Liang, and Tengyu Ma. Verified Uncertainty Calibration. In *arXiv: 1909.10155*, 2019.
- Alexander A. Alemi, Ian Fischer, and Joshua V. Dillon. Uncertainty in the Variational Information Bottleneck. *ArXiv: 1807.00906*, 2018.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On Calibration of Modern Neural Networks. *ICML*, 2017.
- Jean-Michel Marin, Pierre Pudlo, Christian P. Robert, and Robin J. Ryder. Approximate Bayesian computational methods. *Stat Comput*, 22(6):1167–1180, 2012. ISSN 0960-3174, 1573-1375. <https://doi.org/10.1007/s11222-011-9288-2>.
- Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. In *International Conference on Learning Representations*, 2019.
- Sarah M. Erfani, Sutharshan Rajasegarar, Shanika Karunasekera, and Christopher Leckie. High-dimensional and large-scale anomaly detection using a linear one-class svm with deep learning. *Pattern Recognition*, 58:121 – 134, 2016. ISSN 0031-3203. <https://doi.org/https://doi.org/10.1016/j.patcog.2016.03.028>.
- Erik Marchi, Fabio Vesperini, Felix Weninger, Florian Eyben, Stefano Squartini, and Björn Schuller. Non-linear prediction with lstm recurrent neural networks for acoustic novelty detection. 2015. <https://doi.org/10.1109/IJCNN.2015.7280757>.
- Nicholas Carlini and David Wagner. Adversarial Examples Are Not Easily Detected: Bypassing Ten Detection Methods. *arXiv:1705.07263*, 2017.
- Dan Hendrycks and Kevin Gimpel. A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks. *ICLR*, 2017.
- Lingxue Zhu and Nikolay Laptev. Deep and Confident Prediction for Time Series at Uber. *2017 IEEE Int. Conf. Data Min. Workshop ICDMW*, pages 103–110, 2017. <https://doi.org/10.1109/ICDMW.2017.19>.
- Hao Wang and Dit-Yan Yeung. Towards bayesian deep learning: A framework and some existing methods. *IEEE Transactions on Knowledge and Data Engineering*, 28(12):3395–3408, Dec 2016. ISSN 1041-4347. <https://doi.org/10.1109/tkde.2016.2606428>. URL <http://dx.doi.org/10.1109/TKDE.2016.2606428>.
- Andreas Damianou and Neil Lawrence. Deep gaussian processes. In Carlos M. Carvalho and Pradeep Ravikumar, editors, *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, volume 31 of *Proceedings of Machine Learning Research*, pages 207–215, Scottsdale, Arizona, USA, 29 Apr–01 May 2013. PMLR.
- Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging Weights Leads to Wider Optima and Better Generalization. *arXiv:1803.05407 [cs, stat]*, March 2018.

- S. Mandt, M. D. Hoffman, and D. M. Blei. Stochastic Gradient Descent as Approximate Bayesian Inference. *arXiv e-prints*, April 2017.
- Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, and F Huang. A tutorial on energy-based learning. *Predicting structured data*, 1(0), 2006.
- Michael Tschannen, Josip Djolonga, Paul K Rubenstein, Sylvain Gelly, and Mario Lucic. On mutual information maximization for representation learning. *arXiv preprint arXiv:1907.13625*, 2019.
- Andrew Y. K. Foong, David R. Burt, Yingzhen Li, and Richard E. Turner. Pathologies of Factorised Gaussian and MC Dropout Posteriors in Bayesian Neural Networks. *arXiv:1909.00719*, September 2019.
- David M. Blei, Alp Kucukelbir, and Jon D. McAuliffe. Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association*, 112(518):859–877, April 2017. ISSN 0162-1459, 1537-274X. <https://doi.org/10.1080/01621459.2017.1285773>.
- Florian Wenzel, Kevin Roth, Bastiaan S. Veeling, Jakub Świątkowski, Linh Tran, Stephan Mandt, Jasper Snoek, Tim Salimans, Rodolphe Jenatton, and Sebastian Nowozin. How Good is the Bayes Posterior in Deep Neural Networks Really? *arXiv:2002.02405*, February 2020.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Michael C. Mozer, and Yoram Singer. Identity crisis: Memorization and generalization under extreme overparameterization. In *International Conference on Learning Representations*, 2020.
- Aharon Azulay and Yair Weiss. Why do deep convolutional networks generalize so poorly to small image transformations? *Journal of Machine Learning Research*, 20(184):1–25, 2019.
- Ethan Fetaya, Joern-Henrik Jacobsen, Will Grathwohl, and Richard Zemel. Understanding the limitations of conditional generative models. In *International Conference on Learning Representations*, 2020.
- Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016.
- Jiefeng Chen, Xi Wu, Yingyu Liang, Somesh Jha, et al. Robust out-of-distribution detection in neural networks. *arXiv preprint arXiv:2003.09711*, 2020.
- Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019.
- Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. *arXiv preprint arXiv:1812.04606*, 2018b.
- Tianyu Pang, Kun Xu, Yinpeng Dong, Chao Du, Ning Chen, and Jun Zhu. Rethinking softmax cross-entropy loss for adversarial robustness. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=Byg9A24tvB>.
- Chang Xiao, Peilin Zhong, and Changxi Zheng. Enhancing adversarial defense by k-winners-take-all. In *International Conference on Learning Representations*, 2020.

- Jeremy M Cohen, Elan Rosenfeld, and J. Zico Kolter. Certified adversarial robustness via randomized smoothing, 2019.
- Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=SyxAb30cY7>.