

On Attacking Out-Domain Uncertainty Estimation in Deep Neural Networks

Huimin Zeng^{1*}, Zhenrui Yue¹, Yang Zhang², Ziyi Kou¹, Lanyu Shang¹, Dong Wang¹

¹University of Illinois at Urbana-Champaign

²University of Notre Dame

{huiminz3, zhenrui3, ziyikou2, lshang3, dwang24}@illinois.edu, yzhang42@nd.edu

Abstract

In many applications with real-world consequences, it is crucial to develop reliable uncertainty estimation for the predictions made by the AI decision systems. Targeting at the goal of estimating uncertainty, various deep neural network (DNN) based uncertainty estimation algorithms have been proposed. However, the robustness of the uncertainty returned by these algorithms has not been systematically explored. In this work, to raise the awareness of the research community on robust uncertainty estimation, we show that state-of-the-art uncertainty estimation algorithms could fail catastrophically under our proposed adversarial attack despite their impressive performance on uncertainty estimation. In particular, we aim at attacking the out-domain uncertainty estimation: under our attack, the uncertainty model would be fooled to make high-confident predictions for the out-domain data, which they originally would have rejected. Extensive experimental results on various benchmark image datasets show that the uncertainty estimated by state-of-the-art methods could be easily corrupted by our attack.

1 Introduction

Deep neural networks (DNNs) have been thriving in various applications, such as computer vision, natural language processing and decision making. However, in many applications with real-world consequences (e.g. autonomous driving, disease diagnosis, loan granting), it is not sufficient to merely pursue the high accuracy of the predictions made by the AI models, since the deterministic wrong predictions without any uncertainty justification could lead to catastrophic consequences [Galil and El-Yaniv, 2021]. Therefore, to address the issue of producing over-confident wrong predictions from DNNs, great efforts have been made to quantify the predictive uncertainty of the models, so that ambiguous or low-confident predictions could be rejected or deferred to an expert.

Indeed, state-of-the-art algorithms for uncertainty estimation in DNNs have shown their impressive performance in

terms of quantifying the confidence of the model predictions. Handling either in-domain data (generated from the training distribution) or out-domain data under domain shift, these algorithms could successfully assign low confidence scores for ambiguous predictions. However, the robustness of such estimated uncertainty/confidence is barely studied. The robustness of uncertainty estimation in this paper is defined as: to which extent would the predictive confidence be affected when the input is deliberately perturbed? Consider an example of autonomous driving [Feng *et al.*, 2018], where the visual system of a self-driving car is trained using collected road scene images. When an extreme weather occurs, to avoid over-confident wrong decisions, the visual system would show high uncertainty for the images captured by the sensors, since the road scenes observed by the sensors under the extreme weather (out-domain) could be drastically different from the training scenes (in-domain). However, a malicious attacker might attempt to perturb the images captured by the sensors in such a way, that the visual system would regard the out-domain images as in-domain images, and make completely wrong decisions with a high confidence, leading to catastrophic consequences.

Therefore, with the intention of raising the attention of the research community to systematically investigate the *robustness of the uncertainty estimation*, we show that SoTA DNN-based uncertainty estimation algorithms could fail easily under our proposed adversarial attack. In particular, we focus on attacking the uncertainty estimation for *out-domain* data in an image classification problem. That is, under our attack, the uncertainty estimation models would be fooled to assign extremely high confidence scores for the out-domain images, which originally would have been rejected by these models due to the low confidence scores (e.g. Softmax score [Galil and El-Yaniv, 2021]). As shown in Figure 1, the attacker will perturb the out-domain images into the high-confidence region of the victim system. To show vulnerability of the SoTA DNN based uncertainty estimation algorithms under our threat model, we launched our proposed out-domain adversarial attack on various algorithms, including Deep Ensemble [Lakshminarayanan *et al.*, 2016], RBF-based Deterministic Uncertainty Quantification (DUQ) [Van Amersfoort *et al.*, 2020], Gaussian process based Deterministic Uncertainty Estimation (DUE) [van Amersfoort *et al.*, 2021] and Spectral-Normalized Gaussian Process (SNGP) [Liu *et al.*,

*Contact Author

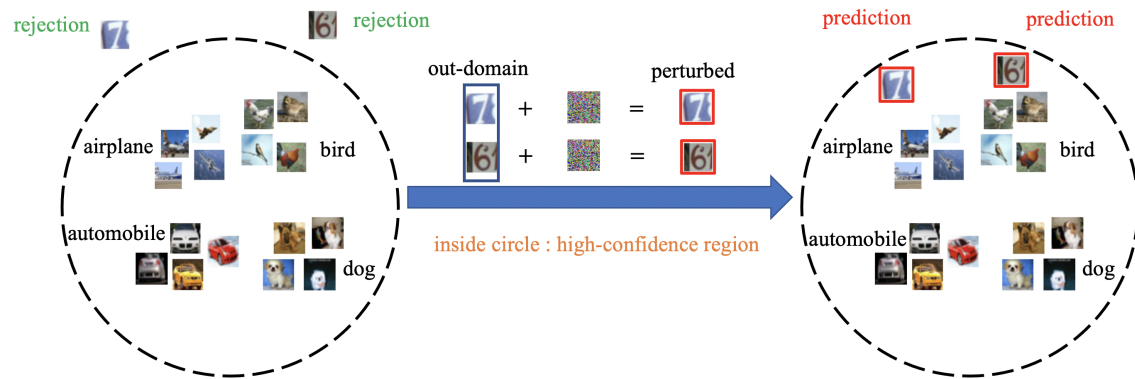


Figure 1: Consider an uncertain model trained on images of airplanes, automobile, birds and dogs. On the left, the image of '7' and '6' would be rejected by the uncertain model, since they are out-domain and geometrically lie outside the confidence region (the circle). In comparison, on the right, our proposed attack could move the out-domain images into the high-confidence region of the uncertain model (inside the circle), so that the uncertain model would make wrong predictions with high confidence.

2020]. Extensive experimental results in various benchmark image datasets show that the uncertainty estimated by these algorithms could be drastically corrupted under our attack.

2 Related Work

Uncertainty Estimation in DNNs. A significant amount of algorithms have been proposed to quantify the uncertainty in DNNs [Lakshminarayanan *et al.*, 2016; van Amersfoort *et al.*, 2021; Van Amersfoort *et al.*, 2020; Liu *et al.*, 2020; Alarab and Prakoonwit, 2021; Gal and Ghahramani, 2016]. For instance, in [Lakshminarayanan *et al.*, 2016], a set of neural networks were trained to construct an ensemble model for uncertainty estimation using the training data. In [Van Amersfoort *et al.*, 2020], RBF (radial basis function) kernel was used to quantify the uncertainty expressed by the deep neural networks. In addition, a further thread of studies built their algorithm based on the Gaussian process (GP), which has been proved to a powerful tool for uncertainty estimation theoretically. However, GP suffers from low expressive power, resulting in poor model accuracy [Bradshaw *et al.*, 2017]. Therefore, based on [Bradshaw *et al.*, 2017], [van Amersfoort *et al.*, 2021] proposed to regularize the feature extractor, so that the deep kernel learning could be stabilized. Similarly, in [Liu *et al.*, 2020], the predictive uncertainty is computed with a spectral-normalized feature extractor and a Gaussian process. However, the *robustness* of these state-of-the-art algorithms is not well studied. As shown in this paper, the uncertain DNN models trained using these algorithms could be fooled easily under our attack.

Adversarial Examples. Despite the impressive performance on clean data, it has been shown that DNNs could be extremely vulnerable to adversarial examples [Szegedy *et al.*, 2013; Goodfellow *et al.*, 2014; Kannan *et al.*, 2018; Carlini and Wagner, 2017]. That is, deep classifiers could be fooled to make completely wrong predictions for the input images that are deliberately perturbed. However, our proposed adversarial attack is different from the traditional adversarial examples in two aspects. Firstly, traditional adversarial examples are crafted by the attacker to corrupt the

accuracy of classifiers, whereas our threat model is aimed at attacking uncertainty estimation in DNNs and increasing its confidence of the wrong predictions. Moreover, the traditional adversarial attack is defined under the close-world assumption [Reiter, 1981; Han *et al.*, 2021], where all possible predictive categories are covered by the training data. The in-domain property limits the power of adversary, since all possible perturbing directions could be simulated by performing untargeted adversarial training [Carlini and Wagner, 2017]. In comparison, our proposed attack is out-domain: perturbing the out-domain data into the in-domain data distribution that the uncertainty estimation model is trained to fit. As shown in our experiments, traditional in-domain adversarial defense mechanism could not withstand our attack: even adversarially trained models could be fooled by our out-domain adversarial examples.

Adversarial Attacks on Uncertainty Estimation. The idea of attacking uncertainty estimation in DNNs have been introduced in [Galil and El-Yaniv, 2021] for the first time. In [Galil and El-Yaniv, 2021], the authors tried to manipulate the confidence score over the predictions by perturbing the input images. By heuristically moving the correctly classified input images towards the decision boundary of the model, the attacker could reduce the confidence of the model on the correct predictions. However, the threat model in [Galil and El-Yaniv, 2021] could only attack in-domain images, whereas our attack is formulated for an out-domain scenario.

3 Problem Statement

In this section, we first define a set of key concepts and present the notations for our problem. Then, we show the mathematical formulation of our proposed attack.

3.1 Key Concept Definition

Definition 1 (Data Domain). We define two data domains in this paper, namely \mathcal{P}_{IN} (representing in-domain distribution) and \mathcal{P}_{OUT} (for out-domain distribution).

An excellent uncertainty estimation model is expected to make high-confident predictions for the in-domain data, but

produce low-confident predictions for out-domain data.

Definition 2 (Dataset). *In our problem, the training set \mathcal{D}_{train} consists of images sampled from \mathcal{P}_{IN} : $\mathcal{D}_{train} = \{(\mathbf{x}, y) | (\mathbf{x}, y) \sim \mathcal{P}_{IN}\}$, where \mathbf{x} refers to input images and y represents the ground truth labels for \mathbf{x} . The test set contains images sampled from \mathcal{P}_{OUT} , since we mainly evaluate our attack on out-domain data: $\mathcal{D}_{test} = \{\mathbf{x} | \mathbf{x} \sim \mathcal{P}_{OUT}\}$.*

Definition 3 (Uncertainty Estimation). *Most uncertainty estimation models f_θ are defined as a mapping from the input image \mathbf{x} to the predictive confidence distributed over all possible categories $\hat{\mathbf{p}}(\mathbf{y})$: $f_\theta : \mathbf{x} \rightarrow \hat{\mathbf{p}}(\mathbf{y}) \in \mathcal{R}^C$, where C is the total number of classes of the in-domain data.*

Regarding $\hat{\mathbf{p}}(\mathbf{y})$, an ideal uncertainty estimation model will assign a uniform confidence over all classes for the out-domain data, since it would be too uncertain to make any predictions but random guess. Moreover, dependent on uncertainty estimation algorithm, $\hat{\mathbf{p}}(\mathbf{y})$ could have different physical implications. For instance, in Deep Ensemble, $\hat{\mathbf{p}}(\mathbf{y})$ is the Softmax uncertainty score. In DUQ, $\hat{\mathbf{p}}(\mathbf{y})$ is interpreted as the distances between the test data point and the centroids.

Definition 4 (Evaluation Metrics). *In terms of evaluating the model's uncertainty on out-domain data, we use two metrics: Entropy $\mathcal{H}(f_\theta)$:*

$$\begin{aligned} \mathcal{H}(f_\theta) &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{test}} \left[\sum_{i=1}^C -[f_\theta(\mathbf{x})]_i \log [f_\theta(\mathbf{x})]_i \right] \\ &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{test}} \left[\sum_{i=1}^C -\hat{p}(y_i) \log \hat{p}(y_i) \right]. \end{aligned} \quad (1)$$

Rejection Rate $\mathcal{R}(f_\theta)$:

$$\begin{aligned} \mathcal{R}(f_\theta) &= 1 - \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{test}} \left[\mathbb{1}(\max\{f_\theta(\mathbf{x})\} \geq \tau) \right] \\ &= 1 - \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{test}} \left[\mathbb{1}(\max\{\hat{\mathbf{p}}(\mathbf{y})\} \geq \tau) \right] \\ &= 1 - \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{test}} \left[\mathbb{1}(\max\{\hat{p}(y_1), \dots, \hat{p}(y_C)\} \geq \tau) \right], \end{aligned} \quad (2)$$

where τ is the pre-defined confidence threshold, and a prediction will be rejected by the model if the predictive confidence over all categories could not achieve the threshold.

Entropy $\mathcal{H}(f_\theta)$ measures the averaged uncertainty of a random variable's possible outcomes. Therefore, an ideal uncertainty estimation model will show a high entropy over all test (out-domain) data, whereas our attacker will increase the model's confidence for the wrong predictions, resulting in a reduced entropy. Rejection rate $\mathcal{R}(f_\theta)$ measures the fraction of the rejected test samples over the test set. Given the test set and a threshold τ , a reasonable uncertain model should show a high rejection rate as it is too uncertain to make any predictions for the out-domain data.

Note that in Equation 1 and Equation 2, it is required that all $\hat{p}(y_i) \in \hat{\mathbf{p}}(\mathbf{y})$ are probabilities. Therefore, for the uncertain models which do not output a probability distribution, their outputs will be normalized into probabilities for evaluation in our experiments. For instance, since DUE outputs the RBF-based distances, we normalize all distances into a

categorical distribution (More details in Section 5), but Deep Ensemble uses averaged predicted probabilities directly [Lakshminarayanan *et al.*, 2016], meaning that there is no need for normalization for Deep Ensemble.

3.2 Problem Formulation

Threat Model. Our threat model is that the adversarial attacker tries to perturb any out-domain test image $\mathbf{x} \sim \mathcal{P}_{OUT}$ into \mathbf{x}' , so that the uncertainty estimation model will believe that \mathbf{x}' is in-domain data and produce a high-confident prediction for it. We assume a white-box attack [Goodfellow *et al.*, 2014; Galil and El-Yaniv, 2021]. Moreover, we require that \mathbf{x}' must be visually innocuous, indicating that \mathbf{x}' and \mathbf{x} must be geometrically close to each other.

To attack the uncertainty estimation on a specific out-domain test sample \mathbf{x} , we need to solve:

$$\begin{aligned} \min_{\delta} & \left[\arg \max\{f_\theta(\mathbf{x} + \delta)\} \geq \tau \right] \\ \text{s.t.} & \quad \mathbf{x} \sim \mathcal{P}_{OUT}, \quad \|\delta\| \leq \epsilon. \end{aligned} \quad (3)$$

In Equation 3, τ represents the confidence threshold defined for Equation 2. ϵ is a small non-negative scalar, bounding the norm of the adversarial perturbation, so that the perturbed image is visually innocuous. Under this formulation, the optimal adversarial perturbation could fool the uncertain model to make predictions with high-confidence (higher than the rejection threshold) without being visually detectable.

4 Algorithm

To successfully perturb out-domain test samples, our objective function (Equation 3) must be solved rationally. However, due to the non-convexity of f_θ and non-differentiability of $\arg \max$, it is intractable to compute a closed-form solution for Equation 3. Therefore, instead of solving Equation 3 analytically, we propose to approximate optimal out-domain adversarial examples using a gradient-based method.

To begin with, we replace the optimization objective of Equation 3 with a new differentiable objective. That is, for an out-domain image \mathbf{x} , the attacker approximates its optimal perturbation by solving:

$$\begin{aligned} \min_{\delta} & \left[l(f_\theta(\mathbf{x} + \delta), \hat{y}) \right] \\ \text{s.t.} & \quad \mathbf{x} \sim \mathcal{P}_{OUT}, \quad \|\delta\| \leq \epsilon, \end{aligned} \quad (4)$$

where \hat{y} is the heuristically found closest attacked label. More precisely, \hat{y} is the original prediction produced by the uncertainty estimation model for the clean out-domain image.

$$\begin{aligned} \hat{y} &= \arg \max_y \hat{p}(y) \\ &= \arg \max_i \left[\hat{p}(y_1), \dots, \hat{p}(y_i), \dots, \hat{p}(y_C) \right] \\ &= \arg \max_i \left[[f_\theta(\mathbf{x})]_1, \dots, [f_\theta(\mathbf{x})]_i, \dots, [f_\theta(\mathbf{x})]_C \right]. \end{aligned} \quad (5)$$

Note that it is determined by Equation 5 that our proposed attack is untargeted. That being said, the adversary only moves the out-domain data to its closest data manifold instead of targeted ones. This untargeted formulation could guarantee the

Algorithm 1 Perturbing out-domain data

Input: uncertainty estimation model f_θ , a clean test sample $x \sim \mathcal{D}_{\text{test}}$
Parameter: adversarial radius ϵ , number of iterations for the attack K , step size η at each iteration
Output: the adversarial test sample x'

- 1: Initialize the iteration counter $k = 0$
- 2: Initialize the adversarial example $x' = x$
- 3: **while** $k \leq K$ **do**
- 4: $\hat{y} = \arg \max_y \hat{p}(y)$
- 5: $l(f_\theta, x', \hat{y}) = l(f_\theta(x'), \hat{y})$
- 6: $x' = \prod_{\mathbb{B}(x, \epsilon)}(x' - \eta \text{sign} \nabla_{x'} l(f_\theta, x', \hat{y}))$, where \prod is the projection operator
- 7: **end while**
- 8: **return** x'

efficiency of the adversary: the path, over which the clean data is moved, could be much shorter compared to targeted attacks, since the targeted data manifolds could be further.

The concrete procedure of generating an adversarial example for a out-domain sample x is presented in Algorithm 1. In Algorithm 1, Line 3 to Line 7 demonstrate how the optimal adversarial example (i.e. the minimizer of Equation 4) will be approximated using projected gradient descent. In particular, Line 4 corresponds to Equation 5, namely the finding of the attacked label \hat{y} . Line 5 to Line 6 demonstrate how the adversarial example is updated using gradient of the loss evaluated at current steps. The projection operation \prod guarantees that the norm of the adversarial perturbation would not exceed the adversarial radius ϵ . After applying Algorithm 1 to an out-domain sample x , the perturbed sample x' will be predicted by the uncertain model f_θ as \hat{y} with a high confidence.

Finally, we would like to comment on the differentiability of our formulation for attacking out-domain uncertainty estimation. Firstly, recall the reason we introduce Equation 4 to replace Equation 3 is that the operation $\arg \max$ in Equation 3 is not differentiable. Using a differentiable surrogate loss function l (e.g. cross-entropy, variational ELBO), the gradient could be computed for the output layer of the uncertain model at the first step. However, different uncertainty estimation algorithms could be non-differentiable by design. For instance, in Gaussian process based algorithms (e.g. DUE [van Amersfoort *et al.*, 2021]), Monte Carlo sampling and other non-differentiable computation are usually employed to compute the final predictive uncertainty distribution. Therefore, to address such non-differentiable issues, we follow the exact implementation in the papers of DUE to study the feasibility of our proposed attack. For DUE [van Amersfoort *et al.*, 2021], the reparametrization trick is implemented as in [van Amersfoort *et al.*, 2021] to make the generation of the adversarial perturbation differentiable. As for the Gaussian process in SNGP [Liu *et al.*, 2020], we follow the training implementation presented in [Liu *et al.*, 2020] to guarantee the differentiability of the uncertain model. Regarding Deep Ensemble and DUQ, both of them are deterministic and differentiable, hence our attack could be launched directly.

5 Experiments

We evaluate the efficacy of our proposed out-domain uncertainty attack by assessing to which extent, the victim model could be deceived to make high-confident predictions for perturbed out-domain data. From a technical point of view, the victim models are trained to learn the in-domain data distribution using the in-domain training set $\mathcal{D}_{\text{train}}$ (as defined in Section 3), but are tested on perturbed out-domain data.

5.1 Dataset and Experimental Setup

Dataset. Following the experimental design presented in [Lakshminarayanan *et al.*, 2016; Van Amersfoort *et al.*, 2020; van Amersfoort *et al.*, 2021] we use MNIST [LeCun *et al.*, 1998] vs. NotMNIST¹, and CIFAR-10 [Krizhevsky and Hinton, 2009] vs. SVHN [Netzer *et al.*, 2011] to test the efficacy of our proposed attack on corrupting the out-domain uncertainty estimation algorithms in DNNs. MNIST dataset contains 0-9 handwritten digit images, and NotMNIST dataset contains images for letters A-J (10 classes) taken from different fonts. CIFAR10 is a dataset containing images of common objects of 10 kinds in the real-world (airplane, automobile, dogs, e.t.c.), whereas SVHN consists of colored digit images taken from real-world scenes. Based on the description above, it is clear that MNIST vs. NotMNIST and CIFAR10 vs. SVHN represent drastically different data domains.

Experimental Setup.² In many out-domain uncertainty estimation studies [Lakshminarayanan *et al.*, 2016; Van Amersfoort *et al.*, 2020; van Amersfoort *et al.*, 2021], it is common to train uncertainty estimation models with in-domain data, and then evaluate the predictive confidence returned by the model over the out-domain test data. Therefore, without loss of generality, we also follow the experimental setup for out-domain uncertainty estimation in [Lakshminarayanan *et al.*, 2016; Van Amersfoort *et al.*, 2020; van Amersfoort *et al.*, 2021]. There are two sets of experiments in our work. In the first set of experiment, the uncertain models are trained using MNIST (in-domain), and tested on NotMNIST (out-domain). In the second set of experiments, CIFAR10 is used as the in-domain training data, and the resulted model is tested on out-domain SVHN data.

5.2 Baselines

Baseline DNN Architectures. To efficiently model the data distribution and avoid overfitting, we select one commonly used neural network architecture for two sets of experiments respectively. In particular, for MNIST vs. NotMNIST, LeNet-5 [LeCun *et al.*, 1998] is used as the baseline architecture to build uncertainty estimation models. As for CIFAR10 vs. SVHN, we use ResNet-18 [He *et al.*, 2016]. Note that the baseline architectures are not equivalent to the final uncertainty estimation models. For instance, in a Deep Ensemble model, there are multiple ensemble classifiers. For DUQ, DUE and SNGP, the baseline models only work as feature extractors, after which there could still be further computational modules, such as RBF kernels or Gaussian process.

¹dataset link: <http://yaroslavvb.blogspot.co.uk/2011/09/notmnist-dataset.html>

²Code will be released after acceptance of this paper.

Baseline Uncertainty Estimation Algorithms. To show the feasibility and efficacy of our proposed attack algorithm, 4 uncertainty estimation algorithms are included in our experiments, namely Deep Ensemble [Lakshminarayanan *et al.*, 2016], DUQ [Van Amersfoort *et al.*, 2020], DUE [van Amersfoort *et al.*, 2021] and SNGP [Liu *et al.*, 2020]. For all experiments, the key hyperparameters for our method and all baseline algorithms are tuned to achieve their best performance for a fair comparison. In summary, we have following trained models under different training schemes:

- **Deep Ensemble:** The vanilla deep ensemble for uncertainty estimation [Lakshminarayanan *et al.*, 2016].
- **Deep Ensemble Adv.:** The augmented version of deep ensemble using adversarially trained ensemble networks [Lakshminarayanan *et al.*, 2016].
- **DUQ:** Uncertainty estimation using RBF kernels [Van Amersfoort *et al.*, 2020].
- **DUE:** Uncertainty estimation using Gaussian process and variational ELBO [van Amersfoort *et al.*, 2021].
- **SNGP:** Spectral-normalized Gaussian process for uncertainty estimation [Liu *et al.*, 2020].

Note that there exists another augmented variant of Deep Ensemble. That is, the augmented deep ensemble model could be built upon a set of adversarially trained ensemble networks. Following the configuration in the original paper of Deep Ensemble [Lakshminarayanan *et al.*, 2016], we train the ensemble networks with fast gradient sign method (FGSM) instead of projected gradient descent (PGD). For a fair comparison, when evaluating deep ensemble models, we also modify our attack with iterations to a single-step attack. When performing adversarial training, the training radius ϵ_{tr} must be specified. In our experiments, two different training adversarial radius (ϵ_{tr}) are used. For MNIST vs. NotMNIST, adversarial training with $\epsilon_{tr} = 0.1$ and $\epsilon_{tr} = 0.2$ are performed, whereas for CIFAR10 vs. SVHN, the training adversarial radius ϵ_{tr} is set to 0.016 and 0.031, respectively.

5.3 Evaluation Results

Evaluating Efficacy. To begin with, we conduct two sets of experiments to evaluate the efficacy of our proposed adversarial attack. The numerical results on entropy and rejection rate are reported in Table 1. \mathcal{H}_{clean} and \mathcal{R}_{clean} are the entropy and rejection rate computed using the clean, unperturbed out-domain data. In contrast, \mathcal{H}_{adv} and \mathcal{R}_{adv} are evaluated using perturbed out-domain data. In both MNIST vs. NotMNIST and CIFAR10 vs. SVHN, the confidence threshold τ for computing \mathcal{R}_{clean} and \mathcal{R}_{adv} is 0.9. That is, the uncertain model only makes predictions when the confidence score is greater than 0.9. Moreover, to attack NotMNIST images, the adversarial radius ϵ is set to 0.1, and $\epsilon = 0.016$ for SVHN images.

From Table 1, we observe that our proposed attack could reduce both the entropy and rejection rate of different uncertain models significantly. This indicates that the uncertain models would be fooled by our attack to be over-confident about the wrong predictions they made for the out-domain data. For instance, regarding the first evaluation metric, the entropy on NotMNIST of SNGP is reduced from 1.22 to 0.23.

	\mathcal{H}_{clean}	\mathcal{H}_{adv}	\mathcal{R}_{clean}	\mathcal{R}_{adv}
D.E.	0.42	0.26	0.44	0.26
D.E.Adv. ($\epsilon_{tr} = 0.1$)	0.65	0.49	0.60	0.44
D.E.Adv. ($\epsilon_{tr} = 0.2$)	0.72	0.58	0.65	0.52
DUQ	1.29	1.02	0.98	0.78
DUE	0.83	0.09	0.69	0.01
SNGP.	1.22	0.23	0.80	0.17

(a) MNIST vs. NotMNIST: 10-iteration adversarial attack with $\epsilon = 0.1$, rejection rate computed with confidence threshold $\tau = 0.9$.

	\mathcal{H}_{clean}	\mathcal{H}_{adv}	\mathcal{R}_{clean}	\mathcal{R}_{adv}
D.E.	0.45	0.34	0.53	0.43
D.E.Adv. ($\epsilon_{tr} = 0.016$)	1.01	0.65	0.93	0.65
D.E.Adv. ($\epsilon_{tr} = 0.032$)	1.30	1.03	0.98	0.90
DUQ	1.32	0.56	1.00	0.72
DUE	1.26	1.05	0.95	0.81
SNGP.	1.04	0.04	0.90	0.02

(b) CIFAR10 vs. SVHN: 10-iteration adversarial attack with $\epsilon = 0.016$, rejection rate computed with confidence threshold $\tau = 0.9$.

Table 1: Evaluating the efficacy of proposed attack algorithm on MNIST vs. NotMNIST and CIFAR10 vs. SVHN.

As for the second evaluation metric, when making predictions for clean NotMNIST images, the DUQ model would reject 98% of the predictions due to the lack of confidence. However, when attacked by our algorithm, the rejection rate of DUQ would be reduced to 78%. Similar observations are made on all other baseline uncertain models in terms of both entropy and rejection rate. In other words, the uncertain models would be less uncertain and could fail to reject many predictions, which originally would have been rejected when there is no attack. As for CIFAR10 vs. SVHN, our attack still shows high efficacy in terms of increasing the uncertain models' confidence on the out-domain data according to Table 1b. For example, the DUQ model could achieve the rejection rate of 100% on clean SVHN images, but fails to reject 28% perturbed images. Moreover, when comparing Table 1a and Table 1b, we also notice that different uncertainty estimation algorithms show different robustness on different datasets. For instance, the uncertainty of DUE could be corrupted dramatically on MNIST vs. NotMNIST (the entropy is reduced from 0.83 to 0.09), but on CIFAR10 vs. SVHN, DUE shows a much stronger robustness (the entropy is reduced merely from 1.26 to 1.05). In contrast, DUQ is much more robust against our attack on MNIST vs. NotMNIST, but more vulnerable on CIFAR10 vs. SVHN.

To summarize, the consistent reduction of entropy and rejection rates on all uncertain models in Table 1 verifies that the state-of-the-art uncertainty estimation algorithms could be extremely vulnerable in the presence of adversarial attacks. Moreover, despite the vulnerability of all tested uncertainty algorithms, they also display different robustness against the attacks on different datasets.

Robustness Study. In addition to the efficacy evaluation of our proposed attack shown in Table 1b, we also conduct a set of robustness study to understand the relationship between the adversarial radius and attack effect. In particular, we changed the adversarial radius ϵ from 0.1 (correspond-

	D.E.	D.E. Adv. $\epsilon_{tr} = 0.1$	D.E. Adv. $\epsilon_{tr} = 0.2$	DUQ	DUE	SNGP
\mathcal{H}_{clean}	0.42	0.65	0.72	1.29	0.83	1.22
\mathcal{H}_{adv}	0.26	0.49	0.58	1.02	0.09	0.23
\mathcal{H}'_{adv}	0.22	0.45	0.53	0.80	0.07	0.14
\mathcal{H}''_{adv}	0.22	0.48	0.73	0.72	0.08	0.15
\mathcal{R}_{clean}	0.44	0.60	0.65	0.98	0.69	0.80
\mathcal{R}_{adv}	0.26	0.44	0.52	0.78	0.01	0.17
\mathcal{R}'_{adv}	0.22	0.40	0.47	0.48	0.00	0.12
\mathcal{R}''_{adv}	0.21	0.41	0.57	0.40	0.00	0.11

(a) MNIST vs. NotMNIST: 10-iteration adversarial attack with $\epsilon = 0.1$, $\epsilon = 0.2$ and $\epsilon = 0.3$, rejection rate computed with confidence threshold $\tau = 0.9$.

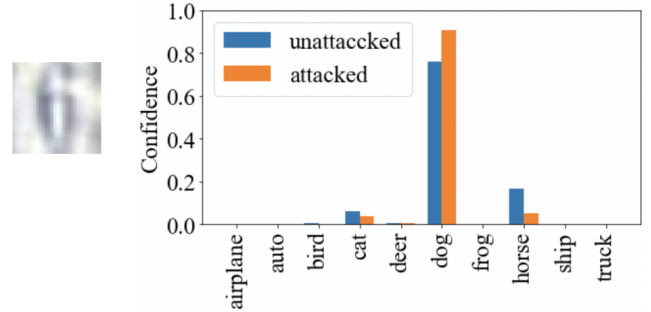
	D.E.	D.E. Adv. $\epsilon_{tr} = 0.016$	D.E. Adv. $\epsilon_{tr} = 0.031$	DUQ	DUE	SNGP
\mathcal{H}_{clean}	0.45	1.01	1.30	1.32	1.26	1.04
\mathcal{H}_{adv}	0.34	0.65	1.03	0.56	1.05	0.04
\mathcal{H}'_{adv}	0.51	0.49	0.86	0.32	0.98	0.03
\mathcal{H}''_{adv}	0.68	0.44	0.75	0.24	0.93	0.03
\mathcal{R}_{clean}	0.53	0.93	0.98	1.00	0.95	0.90
\mathcal{R}_{adv}	0.39	0.65	0.90	0.72	0.81	0.02
\mathcal{R}'_{adv}	0.59	0.48	0.80	0.13	0.80	0.02
\mathcal{R}''_{adv}	0.76	0.42	0.70	0.04	0.82	0.00

(b) CIFAR10 vs. SVHN: 10-iteration adversarial attack with $\epsilon = 0.016$, $\epsilon = 0.031$, $\epsilon = 0.063$ rejection rate computed with confidence threshold $\tau = 0.9$.

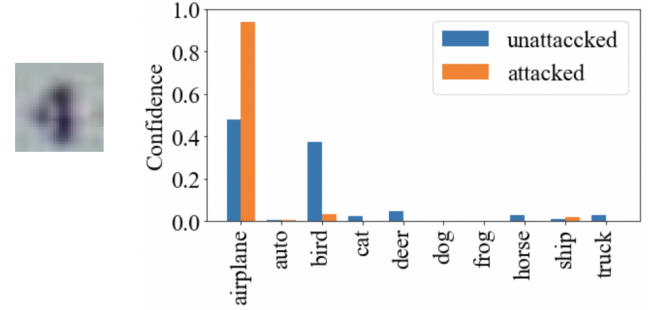
Table 2: Evaluating the efficacy of proposed attack algorithm on MNIST vs. NotMNIST and CIFAR10 vs. SVHN.

ing to \mathcal{H}_{adv} and \mathcal{R}_{adv}) to 0.2 (corresponding to \mathcal{H}'_{adv} and \mathcal{R}'_{adv}) and 0.3 (corresponding to \mathcal{H}''_{adv} and \mathcal{R}''_{adv}) for NotMNIST. Similarly, for SVHN, ϵ is changed from 0.016 to 0.031 and 0.063. As shown in Table 2a and in Table 2b, we observe that in general, with larger ϵ , the attack becomes more powerful and the uncertainty of the model would be further reduced (bold numbers). For instance, the entropy of DUQ on NotMNIST is consistently reduced from 1.02 to 0.80 and 0.72 with the increase of ϵ . However, an opposite trend is also observed, such as on SVHN, both entropy and rejection rate of the vanilla deep ensemble increase when ϵ becomes larger. We believe the optimal ϵ for attacking different dataset is at different ranges.

Visualization of Confidence Distribution. Due to the limited space, we randomly pick two of attacked SVHN images and show the shift of the confidence distribution after applying our attack. In Figure 2a, an adversarially trained deep ensemble model is attacked using our proposed algorithm, and in Figure 2b, a DUE model is attacked. Note that both uncertain models are trained using CIFAR10 as in-domain data, and are tested on SVHN images. As we can observe, both uncertain models are fooled by the adversarial examples to make high-confidence wrong predictions for the out-domain data. In Figure 2b, the DUE model is only around 50% confident about the prediction when there is no attack. However, when the image is perturbed, the confidence is raised to higher than 90%. Originally, the uncertain prediction on the unperturbed images could have been rejected by model, and further deferred the image to an expert, but under our attack, the model will be extremely confident about the wrong predictions, and



(a) Attack adversarially trained Deep Ensemble (D.E. Adv. $\epsilon_{tr} = 0.016$) with $\epsilon = 0.016$. The clean image of the digit '6' from SVHN dataset will be recognized by the model as a dog. After our attack, the confidence is higher than 90%.



(b) Attack DUE with $\epsilon = 0.031$. The clean image of the digit '4' from SVHN dataset will be recognized by the model as an airplane. After our attack, the confidence is higher than 90%.

Figure 2: The predictive confidence distribution over the in-domain categories.

no further critical analysis would be performed.

6 Conclusion

With the intention of raising the attention on the topic of the robustness of uncertainty estimation, we investigated and designed a prototype white-box adversarial attacks on *out-domain* uncertainty estimation. As shown in our experiments, state-of-the-art uncertainty estimation algorithms could be deceived easily by our proposed adversary to make predictions for the out-domain images with very high confidence. In terms of the applications with real-world consequences, the vulnerability of the uncertain models could open a fatal loophole for attackers, leading to catastrophic consequences. Therefore, despite the accuracy of the uncertainty estimation, it is also important to improve the robustness of various uncertainty estimation algorithms in DNNs.

Ethics Statement

Adversarial attacks could cause extremely high threat to existing uncertainty estimation in deep neural networks. Despite the impressive performance on uncertainty estimation on out-domain data, state-of-the-art algorithms still fail catastrophically under the attack of imperceptible perturbations. The existence of such “out-domain” adversarial examples exposes a serious vulnerability in current uncertainty-

based ML systems, such as autonomous driving, medical diagnosis and digital financial systems. In these applications with real-world consequences, the vulnerability of uncertainty systems could place our lives and security at risk.

Our work has the potential to inspire future studies on a new type of adversarial defense mechanism and the design of robust uncertainty estimation algorithms. Although in this work, we present a simple prototype adversarial attack on out-domain uncertainty estimation, the core intention of this work is to raise the attention of the research community to systematically investigate the robustness of the uncertainty estimation in DNNs. We believe our attack could be used to test future robust uncertainty estimation algorithms.

Acknowledgments

This research is supported in part by the National Science Foundation under Grant No. CHE-2105032, IIS-2008228, CNS-1845639, CNS-1831669. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

References

- [Alarab and Prakoonwit, 2021] Ismail Alarab and Simant Prakoonwit. Adversarial attack for uncertainty estimation: Identifying critical regions in neural networks. *arXiv preprint arXiv:2107.07618*, 2021.
- [Bradshaw *et al.*, 2017] John Bradshaw, Alexander G de G Matthews, and Zoubin Ghahramani. Adversarial examples, uncertainty, and transfer testing robustness in gaussian process hybrid deep networks. *arXiv preprint arXiv:1707.02476*, 2017.
- [Carlini and Wagner, 2017] Nicholas Carlini and David Wagner. Adversarial examples are not easily detected: Bypassing ten detection methods. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, pages 3–14, 2017.
- [Feng *et al.*, 2018] Di Feng, Lars Rosenbaum, and Klaus Dietmayer. Towards safe autonomous driving: Capture uncertainty in the deep neural network for lidar 3d vehicle detection. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 3266–3273. IEEE, 2018.
- [Gal and Ghahramani, 2016] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016.
- [Galil and El-Yaniv, 2021] Ido Galil and Ran El-Yaniv. Disrupting deep uncertainty estimation without harming accuracy. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [Goodfellow *et al.*, 2014] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [Han *et al.*, 2021] Lei Han, Xiao Dong, and Gianluca Demartini. Iterative human-in-the-loop discovery of unknown unknowns in image datasets. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 9, pages 72–83, 2021.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [Kannan *et al.*, 2018] Harini Kannan, Alexey Kurakin, and Ian Goodfellow. Adversarial logit pairing. *arXiv preprint arXiv:1803.06373*, 2018.
- [Krizhevsky and Hinton, 2009] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. *Technical report, University of Toronto*, 2009.
- [Lakshminarayanan *et al.*, 2016] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *arXiv preprint arXiv:1612.01474*, 2016.
- [LeCun *et al.*, 1998] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [Liu *et al.*, 2020] Jeremiah Zhe Liu, Zi Lin, Shreyas Padhy, Dustin Tran, Tania Bedrax-Weiss, and Balaji Lakshminarayanan. Simple and principled uncertainty estimation with deterministic deep learning via distance awareness. *arXiv preprint arXiv:2006.10108*, 2020.
- [Netzer *et al.*, 2011] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. *NeurIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [Reiter, 1981] Raymond Reiter. On closed world data bases. In *Readings in artificial intelligence*, pages 119–140. Elsevier, 1981.
- [Szegedy *et al.*, 2013] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [Van Amersfoort *et al.*, 2020] Joost Van Amersfoort, Lewis Smith, Yee Whye Teh, and Yarin Gal. Uncertainty estimation using a single deep deterministic neural network. In *International Conference on Machine Learning*, pages 9690–9700. PMLR, 2020.
- [van Amersfoort *et al.*, 2021] Joost van Amersfoort, Lewis Smith, Andrew Jesson, Oscar Key, and Yarin Gal. On feature collapse and deep kernel learning for single forward pass uncertainty. *arXiv preprint arXiv:2102.11409*, 2021.