
Improving Robustness against Real-World and Worst-Case Distribution Shifts through Decision Region Quantification

Leo Schwinn*^{1,2} Leon Bungert³ An Nguyen¹ René Raab¹ Falk Pulsmeier¹ Doina Precup^{2,4,5}
Björn Eskofier¹ Dario Zanca¹

Abstract

The reliability of neural networks is essential for their use in safety-critical applications. Existing approaches generally aim at improving the robustness of neural networks to either real-world distribution shifts (e.g., common corruptions and perturbations, spatial transformations, and natural adversarial examples) or worst-case distribution shifts (e.g., optimized adversarial examples). In this work, we propose the Decision Region Quantification (DRQ) algorithm to improve the robustness of any differentiable pre-trained model against both real-world and worst-case distribution shifts in the data. DRQ analyzes the robustness of local decision regions in the vicinity of a given data point to make more reliable predictions. We theoretically motivate the DRQ algorithm by showing that it effectively smooths spurious local extrema in the decision surface. Furthermore, we propose an implementation using targeted and untargeted adversarial attacks. An extensive empirical evaluation shows that DRQ increases the robustness of adversarially and non-adversarially trained models against real-world and worst-case distribution shifts on several computer vision benchmark datasets.

1. Introduction

Deep neural networks (DNNs) achieve remarkable results on a wide range of machine learning problems, including image classification (Krizhevsky et al., 2012), speech processing (van den Oord et al., 2016), and more. However, previous work has shown that distribution shifts of the data

can severely impact the classification accuracy of such models. Distribution shifts of the data can be separated into multiple categories, such as out-of-distribution data (Hendrycks & Gimpel, 2017), common corruptions (Hendrycks et al., 2021a), and adversarial examples (Madry et al., 2018).

Making deep learning models robust against distribution shifts in the data is a long-standing problem (Quiñonero-Candela et al., 2008). One area of research focuses on improving the robustness of models at *training-time*. Here, models are trained with specific procedures (e.g., data augmentations or adversarial training), which cannot be applied to pre-trained models (Madry et al., 2018; Yin et al., 2019; Geirhos et al., 2019; Hendrycks et al., 2020). These methods usually entail vast computational overhead and require access to the training data (Madry et al., 2018; Rebuffi et al., 2021). Another set of methods aims to improve the robustness at *test-time*. However, these approaches are usually semi-supervised, rely on assumptions on the distribution of the test data, and often require additional model training at test-time (Li et al., 2017; Shorten & Khoshgoftaar, 2019; Zhang et al., 2021). Moreover, most prior work focuses either on the robustness against real-world distribution shifts (e.g., common corruptions and perturbations, spatial transformations, and natural adversarial examples) or worst-case distribution shifts (e.g., optimized adversarial examples).

In this work, we propose the Decision Region Quantification (DRQ) algorithm that analyses the decision surface of a given model to improve its predictions. Unlike previous work, which mainly specialized in one threat model at a time, DRQ simultaneously improves the robustness of models to both real-world and worst-case distribution shifts. The algorithm is illustrated in Figure 1. Remarkably, the proposed approach does not require any further training data and can be directly combined with pre-trained models during test-time. Additionally, no batch statistics are used during test-time and a single test sample is sufficient for DRQ-based inference.

Our contributions can be summarized as follows: First, we theoretically motivate the proposed DRQ algorithm by demonstrating its ability to smooth small spurious local extrema in the decision surface. Additionally, we provide an

¹University of Erlangen-Nürnberg ²Mila - Quebec AI Institute
³Hausdorff Center for Mathematics, University of Bonn ⁴McGill University ⁵Google Deepmind. Correspondence to: Leo Schwinn <leo.schwinn@fau.de>.

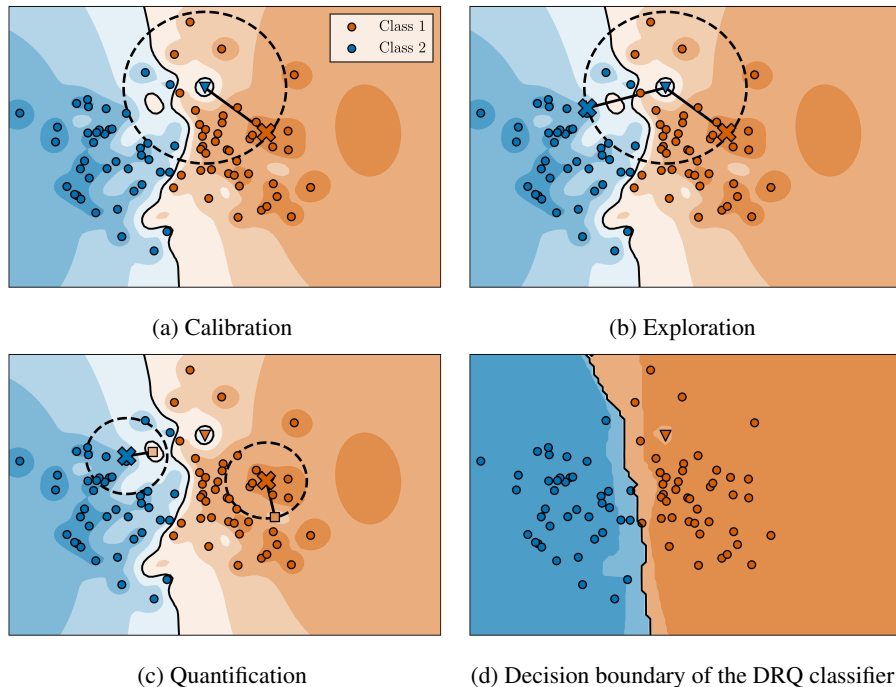


Figure 1: Illustration of the Decision Region Quantification (DRQ) algorithm. The decision boundary of an SVM classifier that is trained to separate the samples of two different two-dimensional Gaussian distributions is shown. Training samples from each class are shown as filled colored circles with a black border. One misclassified point from class 2 is shown as a triangle instead. DRQ is applied for this point as follows: (a) In the *calibration* step of the DRQ algorithm the nearest point in the feature space which is confidently classified as belonging to the other class is discovered (orange cross). The distance to this point is illustrated by the dashed circle. This distance is used as the exploration radius ε_p in the next step. (b) In the *exploration* step, for every class, the input with the highest confidence within the exploration radius is being found (blue and orange cross). (c) In the *quantification* step, for every cross discovered in the exploration step, we search for the lowest confident input within the quantification radius $\alpha \cdot \varepsilon_p$ (radius indicated by the dashed circle; resulting points indicated by the colored squares). The DRQ classifier then assigns the original point to the class that exhibits the highest robustness. Here, the prediction changes from class 2 to class 1. (d) The resulting decision boundary after applying the DRQ algorithm to the SVM classifier for every point in the input space.

implementation using targeted and untargeted adversarial attacks. Furthermore, in an extensive empirical study, we demonstrate that DRQ consistently improves the robustness of pre-trained models against real-world distribution shifts. Moreover, we show that DRQ can be combined with other pre-processing methods to enhance the robustness even further. Lastly, we show that DRQ considerably improves robustness against worst-case distribution shifts for all adversarially-trained models in our experiments.

2. Related Work

The robustness of neural networks to distribution shifts is a broad research area (Hendrycks et al., 2021a). Among the sub-fields of this research area are domain adaptation (Shimodaira, 2000; Wilson & Cook, 2020), out-of-distribution (OOD) detection (Hendrycks & Gimpel, 2017; Schwinn et al., 2021a), corruption and perturbation ro-

bustness (Hendrycks & Dietterich, 2019; Yin et al., 2019; Geirhos et al., 2019; Zhang et al., 2021), robustness to spatial transformations (Engstrom et al., 2019), adversarial robustness (Szegedy et al., 2014; Goodfellow et al., 2015; Madry et al., 2018; Bungert et al., 2021), and more. Here, we focus on robustness against common corruptions and perturbations, spatial transformations, natural adversarial examples, and optimized adversarial examples.

Most approaches that aim to improve the robustness of neural networks adapt the training phase of the model. This includes methods such as strong data augmentation (Hendrycks et al., 2020) and adversarial training (Goodfellow et al., 2015; Madry et al., 2018). Strong data augmentation has shown to be effective for improving the robustness of neural networks against common corruptions (e.g., noise, blur, digital weather corruptions) (Yin et al., 2019; Hendrycks et al., 2020; 2021a). However, data

augmentation has proven to be largely ineffective for increasing robustness against adversarial attacks and common corruptions at the same time. Similarly, adversarial training is mainly effective against adversarial attacks and only marginally improves the robustness of neural networks against common corruptions, while it considerably reduces accuracy on clean data (Raghunathan et al., 2020).

Another group of methods adjusts the inference process of models during test-time to improve their robustness. Data augmentations can be used during test-time to smooth the prediction of a model and make it more robust (Krizhevsky et al., 2012; Shorten & Khoshgoftaar, 2019). Here, the model prediction is averaged over multiple augmented versions of the input. Test-time augmentation methods only require a single test sample and no further re-training of the model. Other test-time adaptations need multiple test inputs to make distributional assumptions about the data. For instance, Li et al. (2017) and Schneider et al. (2020) calculate online batch-normalization statistics on a single batch of the test set to improve model robustness. Recently, Zhang et al. (2021) proposed to increase model robustness by enforcing invariance of the model to multiple augmented versions of the same input. This is done by adapting the model parameters to minimize the entropy of a model’s average output distribution for all augmentations of a given sample. Nevertheless, all these approaches are validated against real-world distribution shifts only and not against worst-case distribution shifts such as adversarial attacks.

Another related line of research are purification methods. These approaches can be used as a pre-processing step to remove data distribution shifts, such as Gaussian noise, thereby improving generalization. Denoising autoencoders (designed to remove noisy artifacts) are one of the first deep learning-based purification methods (Vincent et al., 2010). Based on a similar principle, neural networks can be trained to purify adversarial examples and thus improve the robustness of another pre-trained model (Samangouei et al., 2018; Yoon et al., 2021). However, while the classifier usually does not need to be re-trained to work with the purifier model, data is still needed to train the purifier for a given task. This introduces additional computational overhead and limits applicability when no training data is available.

To the best of our knowledge, previous test-time augmentations are designed either for real-world distribution shifts or worst-case distribution shifts. Moreover, prior approaches often require multiple test samples or additional training of the original or surrogate models.

3. Decision Region Quantification

In this work, we present the Decision Region Quantification (DRQ) algorithm that can be used with any pre-trained

neural network. DRQ does not need additional training of the classifier or a surrogate model and works with a single test sample. Furthermore, it is designed to simultaneously work for real-world and worst-case distribution shifts. This includes common corruptions and perturbations, spatial transformations, natural adversarial examples, and optimized adversarial examples.

3.1. Notation

Let $N : \mathbb{R}^d \rightarrow \mathbb{R}^C$ denote a differentiable classifier, where $C \in \mathbb{N}$ is the number of different classes. Let $F : \mathbb{R}^d \rightarrow \{1, \dots, C\}$, $F(x) := \arg \max(N(x))$ denote the function that directly maps the input to the respective class label. We additionally define $f : \mathbb{R}^d \rightarrow \mathbb{R}$, $f(x) := \max(\text{softmax}(N(x)))$ as the function that maps from the input to the confidence *w.r.t.* the predicted class and $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$, $f_i(x) = \text{softmax}(N(x))_i$ as the function that maps from the input to the confidence *w.r.t.* the class $i \in \{1, \dots, C\}$. Furthermore, we fix an ℓ_p norm $\|\cdot\|_p$ on \mathbb{R}^d and define the ball around $x \in \mathbb{R}^d$ with radius $\varepsilon > 0$ as $B_p(x; \varepsilon) := \{\tilde{x} \in \mathbb{R}^d : \|\tilde{x} - x\|_p \leq \varepsilon\}$.

Definition 3.1. A *decision region* is the subset of the input space containing all inputs $x \in \mathbb{R}^d$ which are classified as the same class i :

$$M_i := \{x : F(x) = i\} \text{ for } i \in \{1, \dots, C\}.$$

Every decision region can be regarded as the union of its disconnected components, i.e., $M_i = \cup_{j \in \mathcal{J}} M_i^j$, where $M_i^{j_1} \neq M_i^{j_2}, \forall j_1, j_2 \in \mathcal{J}$.

Definition 3.2. Let $x_k \in M_i$, such that $F(x_k) = i$. The *local decision region* associated with x_k is the unique connected component $M_i^{j_k}$ which encloses the respective sample, i.e., $x_k \in M_i^{j_k}$.

In Figure 1 the sample depicted as a triangle lies within the local decision region of class 2, which itself lies within a larger local decision region of class 1.

3.2. The DRQ Algorithm

The proposed DRQ algorithm aims to improve the robustness of a pre-trained classifier against real-world and worst-case distribution shifts. To this end, it is necessary to define robustness and how it can be quantified. An important property of a robust classification decision is that it does not change when the input is *slightly* perturbed (e.g., by adding Gaussian or adversarial noise) (Szegedy et al., 2014; Hendrycks et al., 2021a). Thus, we define a classification decision and the associated input sample as robust if no points in direct vicinity of the sample exhibit *considerably* lower confidence for the originally predicted class. For a given sample, DRQ quantifies the robustness of local decision re-

gion in its vicinity. Subsequently, the class label associated with the most robust local decision region is assigned.

Specifically, the DRQ classifier $F_{\text{DRQ}} : \mathbb{R}^d \rightarrow \{1, \dots, C\}$ takes the following steps to make a prediction for a data point $x \in \mathbb{R}^d$.

Calibration: In the calibration step, a vicinity S of x is defined. This can be achieved by setting $S := B_p(x; \varepsilon)$ for a suitable value of $\varepsilon > 0$. We propose to choose ε by searching for the distance of the closest sample which is confidently classified as a different class, i.e., we solve

$$\varepsilon_p := \min_{\tilde{x} \in \mathbb{R}^d} \{ \|x - \tilde{x}\|_p : F(\tilde{x}) \neq F(x), f(\tilde{x}) \geq c \}, \quad (1)$$

where $c \in [0, 1)$ is a confidence threshold value. The vicinity is then defined as $S := B_p(x; \varepsilon_p)$. In cases where prior knowledge about the robustness of a classifier is already available (e.g., adversarially trained models), the calibration step can be omitted and a predefined ε_p can be used instead (see Section 4.3).

Exploration: In the exploration step, local decision regions in the vicinity of the input $x \in \mathbb{R}^d$ are explored. For every class label $i \in \{1, \dots, C\}$, we search for a highest confidence sample \tilde{x}_i , such that $\tilde{x}_i \in B_p(x; \varepsilon_p)$ and $F(\tilde{x}_i) = i$, if it exists (otherwise we do not consider i as a candidate for the correct class). For this we compute

$$\tilde{x}_i \in \arg \max_{\tilde{x} \in B_p(x; \varepsilon_p)} \{ f(\tilde{x}) : F(\tilde{x}) = i \}. \quad (2)$$

Quantification: In the quantification step, the robustness of all local decision regions associated with \tilde{x}_i 's is evaluated to identify the most robust prediction. For every sample \tilde{x}_i , we quantify the robustness of its local decision region by searching for a sample $\hat{x} \in B_p(\tilde{x}_i; \alpha \varepsilon_p)$ for some $\alpha \in [0, 1]$ such that the confidence with respect to the i -th class, $f_i(\hat{x})$, is minimized. The final decision is then made by defining $F_{\text{DRQ}}(x)$ as the class label that exhibits the highest robustness:

$$F_{\text{DRQ}}(x) := \arg \max_{i \in \{1, \dots, C\}} \min_{\hat{x} \in B_p(\tilde{x}_i; \alpha \varepsilon_p)} f_i(\hat{x}). \quad (3)$$

A theoretical analysis of the potential of DRQ to smooth spurious extrema is provided in Appendix A and is outlined in the following. We derive lower and upper bounds for the confidence of DRQ-based predictions after quantification depending on α . Afterward, we show that DRQ effectively smooths spurious local extrema (i.e., misclassifications induced by distribution shifts such as noise) within the local decision region of the correct class. Moreover, we show that when a classifier is strongly convex around a local optimum, we can make precise statements about the impact

of DRQ on the confidence of the classifier depending on the sharpness of the optimum. The smoothing of spurious local extrema is exemplified in Figure 1. Here, the sample indicated by a triangle lies within a small region containing a local maximum of class 2 within the decision region of class 1. After applying DRQ to the classifier, the local extrema is removed and the sample is assigned to class 2. The extent of smoothing on the decision surface by the DRQ algorithm depends directly on the choice of the vicinity S (i.e., ε_p) and the quantification radius defined by α . Here, $\alpha \approx 0$ increases the confidence compared to the original classifier for all classes, whereas $\alpha \approx 1$ has the opposite effect and reduces the confidence. For intermediate choices of $0 < \alpha < 1$, we can expect that for spurious maxima, the confidence of the original class increases and the confidence of the spurious class decreases. Thus, in the experiments, we set $\alpha = 0.5$, as preliminary experiments confirmed this to be a good value for both real-world and worst-case distribution shifts.

3.3. DRQ Implementation

In the following, we provide a baseline implementation of the DRQ algorithm for differentiable classifiers.

Calibration: If we do not use prior information to skip the calibration step, we find the smallest possible ε_p in (1) with the Fast Minimum-Norm (FMN) attack, which can be used for several different p -norms (Pintor et al., 2021). Therefore, we adjust the attack to optimize for minimum-norm adversaries with predefined confidence instead of adversaries with arbitrary confidence. The norm of the final perturbation found by FMN for each sample is subsequently used as the respective exploration radius ε_p .

Exploration: The optimization problem (2) for finding the local decision region for every class i in the exploration step can be reformulated as a targeted adversarial attack:

$$\tilde{x}_i \in \arg \min_{\tilde{x} \in B_p(x; \varepsilon_p)} \mathcal{L}(N(\tilde{x}), y_i). \quad (4)$$

Here $y_i \in \mathbb{R}^C$ denotes the one-hot encoded vector of the i -th class and the loss function \mathcal{L} is the categorical cross-entropy loss by which problems (2) and (4) become equivalent. This process is repeated for all possible output classes $i \in 1, \dots, C$ or can be constrained to only the most likely predictions to reduce the computational overhead. We use a Projected Gradient Descent (PGD) based attack for this step (Madry et al., 2018).

Quantification: Similarly, we use an untargeted PGD-based adversarial attack

$$\hat{x}_i \in \arg \max_{\hat{x} \in B_p(\tilde{x}_i; \alpha \varepsilon_p)} \mathcal{L}(N(\hat{x}), y_j), \quad (5)$$

where $y_j \in \mathbb{R}^C$ is the one-hot encoded vector of the class

$j := F(\tilde{x}_i)$ and \mathcal{L} is the categorical cross-entropy loss, to solve the inner optimization problem in (3) for every class $i \in \{1, \dots, C\}$. We fix the starting point of every attack to make DRQ completely deterministic and simplify the robustness evaluation with adversarial attacks. Note that it is essential to perform the exploration step for the originally predicted class $j = F(x)$ and not quantify $f_j(x)$ directly. We observed that otherwise an adversary can manipulate the distance of the given sample x to the decision boundary of another class so as to change the outcome of the quantification step and mislead the DRQ-based classifier.

4. Experiments

4.1. Data

We evaluate DRQ for real-world and worst-case distribution shifts. These include common corruption and perturbations, natural adversarial examples, spatial transformations, image renditions, and optimized adversarial examples.

4.1.1. BENCHMARK DATASETS:

We use CIFAR10, CIFAR100 (Krizhevsky, 2009), and ImageNet (Deng et al., 2009) as benchmark datasets. Additionally, we include CIFAR10-C, CIFAR100-C, and ImageNet-C (Hendrycks & Dietterich, 2019), which are corrupted and perturbed versions of the original datasets that contain a wide variety of natural distribution shifts. For CIFAR10-C, CIFAR100-C, and ImageNet-C, we use the strongest corruption severity level of 5. We additionally use only 1000 images from every corruption to reduce the computational overhead. Furthermore, we include ImageNet-A, which consists of naturally occurring adversarial examples that can be classified correctly by humans but cause a considerable decrease in performance for machine learning models (Hendrycks et al., 2021b). Lastly, we include ImageNet-R (Hendrycks et al., 2021a), which contains various renditions of ImageNet classes (e.g., art, cartoon, or graffiti). For all ImageNet datasets, we apply DRQ only to the top-10 largest logits ($N(x)$) to reduce the computational overhead. Note that we do not perform any training in our experiments and only use the respective test datasets.

4.1.2. SPATIAL TRANSFORMATIONS:

In addition to common benchmarks from the literature, we test DRQ against random translations and rotations of the inputs. Here, we apply random translations of up to 4 pixels in the x and y coordinates for CIFAR10 and CIFAR100 (image size 32x32) and random translations of up to 20 pixels for ImageNet (image size 224x224). For all datasets, we apply random rotations between -20 and $+20$ degrees.

4.1.3. ADVERSARIAL EXAMPLES:

Lastly, we evaluate if DRQ can successfully improve the robustness of models against adversarial attacks. We use an ensemble of different attacks and report the worst-case performance against all attacks (Croce & Hein, 2020). Specifically, we consider an attack as successful if any of the ensemble attacks can induce a misclassification for the given input. Several different adversarial attacks have been proposed (Goodfellow et al., 2015; Madry et al., 2018; Croce & Hein, 2020; Schwinn et al., 2021b). We choose four complementary attacks from the autoattack ensemble for our evaluation, which have shown to lead to a reliable evaluation (Croce & Hein, 2020). This includes *FAB* with 100 iterations and 9 target classes, *square* with 5000 queries, and two versions of *APGD* with either the DLR or CE loss and 100 iterations.

Additionally, we designed two adaptive attacks to circumvent our defense. With these attacks, we aim to find adversarial examples that show high robustness in the quantification step of DRQ and simultaneously lead to misclassification. For both attacks, we use the EOT method (Athalye et al., 2018b) to find robust adversarial examples and use the PGD algorithm for optimization with 100 iterations (Madry et al., 2018). For both attacks, we calculate the gradient direction of the adversarial attack as the weighted average over multiple noisy gradients at every attack iteration. For the first attack, we calculate noisy gradients by sampling random noise in the ε -ball $B_p(x; \varepsilon)$ around sample x and injecting it to the current adversarial example. Here, we sample 10 noisy gradients at every attack iteration. Hence, the adversarial example found by the attack is made more robust against noise. For the second attack, we generate the noise with PGD-based adversarial attacks in the ε -ball with 7 iterations each. We use 7 iterations as this has proven to be sufficient for adversarial training (Madry et al., 2018). For the second attack, we sample only 4 noisy gradients at every attack iteration due to the computational complexity of the attack. We refer to the first attack as $\text{PGD}_{\text{noise}}$ and the second attack as $\text{PGD}_{\text{attack}}$. We do not early stop the adversarial attacks once the baseline model misclassifies a given input but rather construct high-confidence adversarial attacks. We additionally investigated end-to-end attacks on the whole DRQ pipeline in Appendix E. As the evaluation is empirical, we encourage other researchers to create adaptive attacks that circumvent our method to update our results accordingly.

4.2. Models

We combine DRQ with several prior methods that aim to make models robust against real-world or worst-case distribution shifts to assess the versatility of DRQ. For our evaluations we only use pre-trained models. We use models

trained to be robust against real-world distribution shifts. Here, we evaluate the model proposed by Hendrycks et al. (2021a) that is trained with a combination of Augmix and DeepAugment (referred to as No-AT) (Hendrycks et al., 2020). To investigate if DRQ can be combined with different adversarial training approaches, we additionally use models trained with ℓ_∞ or ℓ_2 -based adversarial training (referred to as ℓ_∞ -AT and ℓ_2 -AT) (Rice et al., 2020; Wong et al., 2020; Goyal et al., 2020). Models are either taken from the RobustBench library (Croce et al., 2020) or from the GitHub repositories of the authors’ (Hendrycks & Dietterich, 2019).

4.3. DRQ Hyperparameters

The DRQ algorithm has five different hyperparameters: the confidence threshold c for the calibration phase, the number of iterations for the calibration, search, and quantification step and the norm ℓ_p used to assess the robustness of a local decision region. In a preliminary experiment, we evaluated the sensitivity of the performance of DRQ to hyperparameter choices on the CIFAR10 dataset. We conducted all initial experiments on the adversarially trained model proposed in (Wong et al., 2020). We did all experiments on clean data and adversarial examples generated by a standard PGD attack with 100 iterations.

First, We explored the effect of the confidence threshold c on the performance of DRQ. For adversarially trained models, we have prior knowledge about their respective robustness since they were trained to be robust against a certain adversarial budget ε . Thus, we omitted the calibration phase for all adversarially trained models to reduce the computational overhead. For the model proposed in (Wong et al., 2020), the mean ε_p value to misclassify a given sample was approximately equal to twice the ε used for adversarial training the model and we used $\varepsilon_p = 2 \cdot \varepsilon$. This resulted in $\varepsilon_\infty = 16/255$ and $\varepsilon_2 = 1$ for adversarially trained CIFAR models and $\varepsilon_\infty = 8/255$ for adversarially trained ImageNet models. For the non-adversarially trained models, we set the confidence threshold to the confidence of the respective sample. Thereby, the confidence of the inputs found with DRQ in the exploration step is approximately equal to the original confidence of the model. An analysis of the number of iterations required for the different steps is given in Appendix B.1. We choose the ℓ_∞ and ℓ_2 norm to measure the robustness of the local decision regions that are commonly used to measure adversarial robustness in previous work (Croce et al., 2020). An important assumption of the DRQ algorithm is that the robustness of training data is higher than the robustness of out-of-distribution data and adversarial examples. Prior work indicates that models trained to be robust against perturbations in these norms have this property (Roth et al., 2020).

5. Results

In this section, we summarize the results of the experiments. We combine DRQ with previously proposed training-time methods from the literature that aim to increase the robustness of neural networks. We then evaluate the robustness for real-world and worst-case distribution shifts with and without DRQ. We specifically focus on training-time methods since current test-time methods have different prerequisites and are designed for either real-world or worst-case distribution shifts but not for both.

5.1. Common Corruptions, Natural Adversarial Examples, and Image Renditions

Table 1 compares the accuracy of standard inference with the DRQ method applied in the ℓ_∞ and ℓ_2 norm on the CIFAR10-C, CIFAR100-C, ImageNet-C, ImageNet-A, and ImageNet-R datasets. The DRQ method achieves the highest accuracy for every dataset and model. For all adversarially trained models, DRQ is most effective when the same norm is used for adversarial training and DRQ. For all non-adversarially trained models, DRQ is most effective in the ℓ_2 norm. Both DRQ approaches are more effective than standard inference in all but two cases. For the ImageNet model proposed in Hendrycks & Dietterich (2019) DRQ ℓ_∞ slightly decreases the accuracy compared to standard inference. Note that the accuracy for clean data differs only marginally between standard and DRQ-based inference in our experiments (see Appendix D).

DRQ can be used with any pre-trained differentiable classifier without further requirements. Nevertheless, we explored if DRQ can be combined with prior pre-processing approaches from the literature that need additional training to further increase the robustness. Therefore, we combined DRQ with a denoising autoencoder (Vincent et al., 2010) trained on the CIFAR10 dataset. We specifically implemented the DnCNN architecture proposed in (Zhang et al., 2017). Training details are included in Appendix C. The combination of the denoising autoencoder and DRQ increases the accuracy from 68.44% to 70.74% (see Table 2).

5.2. Spatial Transformations

Table 3 compares the accuracy of standard inference with the DRQ method applied in the ℓ_∞ and ℓ_2 norm on the CIFAR10, CIFAR100, and ImageNet datasets. DRQ improves robustness up to 17.57% for rotations and up to 2.3% for translations. The DRQ method achieves the highest accuracy for every dataset and model. In contrast to the evaluation on common corruptions, DRQ is not always most effective when the same norm is used for adversarial training and DRQ against spatial transformations. However, for

Table 1: Accuracy [%] of standard inference and DRQ inference in two different norms. The accuracy is shown for differently trained models on the CIFAR10-C, CIFAR100-C, ImageNet-C, and ImageNet-A datasets. The best accuracy is shown in **bold** and the second best is underlined. Note that we show the mean corruption accuracy instead of the mean corruption error for the common corruption datasets (indicated with -C) to simplify the comparison with the other results.

Models	Training	Standard	DRQ ℓ_∞	DRQ ℓ_2
CIFAR10-C				
Wong et al. (2020)	ℓ_∞ -AT	66.41	68.44	<u>67.32</u>
Gowal et al. (2020)	ℓ_∞ -AT	72.58	73.12	<u>73.03</u>
Rice et al. (2020)	ℓ_2 -AT	71.96	70.08	73.2
Hendrycks et al. (2021a)	No-AT	<u>81.33</u>	<u>81.61</u>	81.88
CIFAR100-C				
Rice et al. (2020)	ℓ_∞ -AT	35.09	36.12	<u>35.68</u>
ImageNet-C				
Wong et al. (2020)	ℓ_∞ -AT	22.82	23.25	<u>22.93</u>
Hendrycks et al. (2021a)	No-AT	<u>56.86</u>	56.48	57.48
ImageNet-A				
Wong et al. (2020)	ℓ_∞ -AT	1.6	2.1	<u>1.9</u>
Hendrycks et al. (2021a)	No-AT	3.76	<u>3.9</u>	4.5
ImageNet-R				
Wong et al. (2020)	ℓ_∞ -AT	31.4	33.3	<u>32.4</u>
Hendrycks et al. (2021a)	No-AT	<u>46.8</u>	45.2	47.0

Table 2: Mean corruption accuracy [%] of standard inference and DRQ in two different norms for the model proposed in (Wong et al., 2020) on the CIFAR10-C dataset with and without a denoising autoencoder. The best accuracy is shown in **bold** and the second best is underlined.

	Standard	DRQ ℓ_∞	DRQ ℓ_2
Model	66.41	68.44	<u>67.32</u>
Model + Denoiser	67.98	70.74	<u>69.92</u>

all adversarially trained models except for the CIFAR100 model proposed by Rice et al. (2020) both DRQ methods show higher robustness compared to standard inference. Furthermore, for adversarially trained models, using DRQ with the same norm as adversarial training always increases the robustness. As for common corruptions, DRQ is most effective in the ℓ_2 norm for all non-adversarially trained models,

5.3. Adversarial Attacks

Table 4 compares the accuracy of standard inference with the DRQ method applied in the ℓ_∞ and ℓ_2 norm on the CIFAR10 and CIFAR100 datasets. Against adversarial attacks, DRQ cannot be applied to the top-n predictions only, as this could be exploited by the adversary¹. Thus, we omitted ImageNet from this experiment due to the computational com-

¹The adversary could for example try to make the correct class the least likely prediction

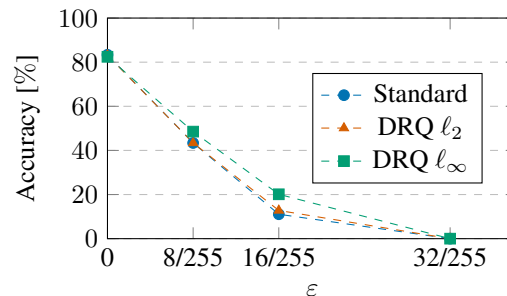


Figure 2: Worst-case accuracy [%] for the model proposed by Wong et al. (2020) on the CIFAR10 dataset. The accuracy is shown for standard inference, DRQ ℓ_∞ , and DRQ ℓ_2 . The accuracy converges to 0 for all methods for $\epsilon = 32/255$

plexity of performing DRQ on all 1000 target classes. DRQ improves the worst-case robustness up to 5.15%. Against adversarially trained models, DRQ is always most effective in the ℓ_∞ norm and always more effective than standard inference. For the non-adversarially trained model proposed in (Hendrycks & Dietterich, 2019), DRQ is not able to increase the worst-case robustness above 0%. While DRQ shows high robustness against some of the attacks (e.g., FAB: 84.49%) it is not able to defend against the adaptive attack. Out of the individual attacks, the two adaptive attacks achieved the highest success rate against DRQ-based inference in all cases (see Appendix E for more details).

Prior defenses against adversarial attacks were later shown to be ineffective against stronger adaptive attacks (Athalye

Table 3: Accuracy [%] of standard inference and DRQ inference in two different norms. The accuracy is shown for differently trained models for rotations and translation of the original input. The best accuracy is shown in **bold** and the second best is underlined. The values are arranged as follows (Standard / DRQ ℓ_∞ / DRQ ℓ_2).

Models	Training	Rotations	Translations
CIFAR10			
Hendrycks et al. (2021a)	No-AT	74.93 / 74.00 / 87.66	92.08 / 91.65 / 92.61
Wong et al. (2020)	ℓ_∞ -AT	43.74 / 61.31 / <u>52.72</u>	59.08 / 61.23 / <u>61.11</u>
Gowal et al. (2020)	ℓ_∞ -AT	51.93 / 64.56 / <u>58.01</u>	78.16 / 78.46 / <u>78.51</u>
Rice et al. (2020)	ℓ_2 -AT	43.49 / 60.03 / <u>56.20</u>	63.96 / 66.26 / <u>65.55</u>
CIFAR100			
Rice et al. (2020)	ℓ_∞ -AT	33.73 / 41.25 / <u>39.16</u>	<u>45.07</u> / 45.35 / 44.87
ImageNet			
Wong et al. (2020)	ℓ_∞ -AT	21.83 / 28.80 / <u>23.02</u>	42.51 / <u>42.80</u> / 42.97
Hendrycks et al. (2021a)	No-AT	53.59 / <u>54.03</u> / 56.73	<u>70.83</u> / 70.21 / 71.54

Table 4: Worst-case accuracy [%] of standard inference and DRQ inference in two different norms against an ensemble of adversarial attacks. The best accuracy is shown in **bold** and the second best is underlined.

Models	Training	Standard	DRQ ℓ_∞	DRQ ℓ_2
CIFAR10				
Wong et al. (2020)	ℓ_∞ -AT	43.37	48.52	<u>43.39</u>
Gowal et al. (2020)	ℓ_∞ -AT	61.51	64.41	<u>62.94</u>
Rice et al. (2020)	ℓ_2 -AT	66.77	75.16	<u>72.65</u>
Hendrycks et al. (2021a)	No-AT	0.00	0.00	0.00
CIFAR100				
Rice et al. (2020)	ℓ_∞ -AT	<u>19.01</u>	24.19	17.79

Table 5: Worst-case accuracy [%] of standard inference and DRQ ℓ_∞ against an ensemble of adversarial attacks. DRQ ℓ_∞ + PGD_{noise} denotes the DRQ method, where the adversarial attack in the exploration step is replaced with PGD_{Noise}.

Attacks	Clean	Worst-case
Standard	83.34	43.37
DRQ ℓ_∞	82.44	48.52
DRQ ℓ_∞ + PGD _{noise}	83.01	49.56

et al., 2018a). Adaptive attacks against DRQ need to find robust local decision regions that belong to the wrong class in the vicinity of the input. Such attacks could be used to find more robust regions in the exploration step of the algorithm. We additionally explored if DRQ-based inference can be made more robust by using the adaptive PGD_{noise} attack during the exploration step of the algorithm. The results are summarized in Table 5. We observe that replacing standard PGD in the exploration step with the adaptive variant increases the worst-case robustness and clean accuracy of DRQ-based inference.

Furthermore, we investigated if DRQ robustness against adversarial attacks largely stems from gradient obfuscation.

For this purpose, we conducted additional experiments on the CIFAR10 model proposed in (Wong et al., 2020) that are inspired by the findings of (Athalye et al., 2018a). Figure 2 demonstrates how the worst-case accuracy decreases with an increasing maximum perturbation norm ε for the adversarial attacks. For all inference approaches, the robustness decreases rapidly to 0% against values of $\varepsilon = 32/255$. We also observed an increase in the success rate of the attacks when we increased the number of attack iterations by a factor β . For $\beta = 2$, the robustness of DRQ ℓ_∞ decreases from 48.52% to 48.01% and it further reduces to 47.93% for $\beta = 4$. Additionally, we conducted a random noise attack, by constructing 1000 random perturbations within the ε -ball, but this only reduced the accuracy marginally against DRQ.

Lastly, we investigated the cosine similarity between DRQ direction $\gamma_{\text{DRQ}} = \hat{x}_{\text{DRQ}} - x$ and the adversarial perturbation direction $\gamma = x_{\text{adv}} - x$ for the model proposed by Wong et al. (2020). Here, we denote by \hat{x}_{DRQ} the most robust sample found in the quantification step (3) of DRQ. We specifically explored samples that were classified correctly by DRQ but wrongly by regular inference. The cosine similarity between the adversarial perturbation direction and the DRQ direction is negative for all attacks. This indicates that DRQ removes the perturbation generated by the adversarial

attacks, thereby improving classification performance. It also supports the assumption that in-distribution data exhibits higher robustness than distribution-shifted data. More details are included in Appendix F.

6. Conclusion and Outlook

In this paper, we propose a test-time algorithm, Decision Region Quantification (DRQ), designed to improve the robustness of any pre-trained differentiable classifier against a wide variety of distribution shifts. We theoretically motivate DRQ and show its effectiveness in an extensive empirical study which includes various models and datasets. In contrast to prior work, the proposed algorithm can handle common corruptions, spatial transformations, natural adversarial examples, image renditions, and adversarial examples simultaneously. DRQ proves to be most effective for models that already exhibit adversarial robustness in our experiments. Furthermore, DRQ still entails considerable computational overhead compared to standard inference and more effective implementations need to be explored. Additionally, while DRQ has been shown to be effective in computer vision tasks in our experiments, its performance in other application areas such as speech processing and natural language processing has yet to be evaluated.

Author Contributions

L.S. conceived the proposed method, conducted all experiments and wrote the initial draft of the paper. D.P., B.E., and D.Z. provided feedback on the experiments and structure of the paper. L.B. and L.S. developed the theorem and L.B. formalized it. All authors analyzed the results and contributed to the final manuscript.

Acknowledgements

Leo Schwinn gratefully acknowledges support by a fellowship within the IFI programme of the German Academic Exchange Service (DAAD). An Nguyen gratefully acknowledges personal funding from Siemens Healthineers. Bjoern M. Eskofier gratefully acknowledges the support of the German Research Foundation (DFG) within the framework of the Heisenberg professorship programme (grant number ES 434/8-1).

References

Athalye, A., Carlini, N., and Wagner, D. A. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *ICML*, volume 80, pp. 274–283. PMLR, 2018a.

Athalye, A., Engstrom, L., Ilyas, A., and Kwok, K. Syn-

thesizing robust adversarial examples. In Dy, J. G. and Krause, A. (eds.), *ICML*, volume 80, pp. 284–293. PMLR, 2018b.

Bungert, L., Raab, R., Roith, T., Schwinn, L., and Tenbrinck, D. CLIP: cheap lipschitz training of neural networks. In *SSVM*, volume 12679, pp. 307–319. Springer, 2021.

Croce, F. and Hein, M. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *ICML*, volume 119, pp. 2206–2216. PMLR, 2020.

Croce, F., Andriushchenko, M., Sehwag, V., Flammarion, N., Chiang, M., Mittal, P., and Hein, M. Robustbench: a standardized adversarial robustness benchmark. *CoRR*, abs/2010.09670, 2020.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *CVPR*, pp. 248–255, 2009.

Engstrom, L., Tran, B., Tsipras, D., Schmidt, L., and Madry, A. Exploring the landscape of spatial robustness. In *ICML*, volume 97, pp. 1802–1811. PMLR, 2019.

Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F. A., and Brendel, W. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In *ICLR*, 2019.

Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. In *ICLR*, 2015.

Gowal, S., Qin, C., Uesato, J., Mann, T. A., and Kohli, P. Uncovering the limits of adversarial training against norm-bounded adversarial examples. *CoRR*, abs/2010.03593, 2020.

Hendrycks, D. and Dietterich, T. G. Benchmarking neural network robustness to common corruptions and perturbations. In *ICLR*, 2019.

Hendrycks, D. and Gimpel, K. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *ICLR*, 2017.

Hendrycks, D., Mu, N., Cubuk, E. D., Zoph, B., Gilmer, J., and Lakshminarayanan, B. Augmix: A simple data processing method to improve robustness and uncertainty. In *ICLR*, 2020.

Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *ICCV*, pp. 8340–8349, 2021a.

- Hendrycks, D., Zhao, K., Basart, S., Steinhardt, J., and Song, D. Natural adversarial examples. *CVPR*, 2021b.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- Krizhevsky, A. Learning multiple layers of features from tiny images. Technical report, 2009.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012.
- Li, Y., Wang, N., Shi, J., Liu, J., and Hou, X. Revisiting batch normalization for practical domain adaptation. In *ICLR*, 2017.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018.
- Pintor, M., Roli, F., Brendel, W., and Biggio, B. Fast minimum-norm adversarial attacks through adaptive norm constraints. In *NeurIPS*, 2021.
- Quiñonero-Candela, J., Sugiyama, M., Schwaighofer, A., and Lawrence, N. *Dataset shift in machine learning*. MIT Press, 2008.
- Raghunathan, A., Xie, S. M., Yang, F., Duchi, J. C., and Liang, P. Understanding and mitigating the tradeoff between robustness and accuracy. In *ICML*, volume 119, pp. 7909–7919. PMLR, 2020.
- Rebuffi, S., Goyal, S., Calian, D. A., Stimberg, F., Wiles, O., and Mann, T. A. Fixing data augmentation to improve adversarial robustness. *CoRR*, abs/2103.01946, 2021.
- Rice, L., Wong, E., and Kolter, J. Z. Overfitting in adversarially robust deep learning. In *ICML*, volume 119, pp. 8093–8104. PMLR, 2020.
- Roth, K., Kilcher, Y., and Hofmann, T. Adversarial training is a form of data-dependent operator norm regularization. In *NeurIPS*, 2020.
- Samangouei, P., Kabkab, M., and Chellappa, R. Defense-gan: Protecting classifiers against adversarial attacks using generative models. In *ICLR*, 2018.
- Schneider, S., Rusak, E., Eck, L., Bringmann, O., Brendel, W., and Bethge, M. Improving robustness against common corruptions by covariate shift adaptation. In *NeurIPS*, 2020.
- Schwinn, L., Nguyen, A., Raab, R., Bungert, L., Tenbrinck, D., Zanca, D., Burger, M., and Eskofier, B. Identifying untrustworthy predictions in neural networks by geometric gradient analysis. In *UAI*, volume 161, pp. 854–864. PMLR, 2021a.
- Schwinn, L., Raab, R., Nguyen, A., Zanca, D., and Eskofier, B. M. Exploring misclassifications of robust neural networks to enhance adversarial attacks. *CoRR*, abs/2105.10304, 2021b.
- Shimodaira, H. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of statistical planning and inference*, 90(2):227–244, 2000.
- Shorten, C. and Khoshgoftaar, T. M. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6:1–48, 2019.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I. J., and Fergus, R. Intriguing properties of neural networks. In *ICLR*, 2014.
- van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A. W., and Kavukcuoglu, K. Wavenet: A generative model for raw audio. In *ISCA Speech Synthesis Workshop*, pp. 125, 2016.
- Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., and Manzagol, P. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research*, 11: 3371–3408, 2010.
- Wilson, G. and Cook, D. J. A survey of unsupervised deep domain adaptation. *ACM Transactions on Intelligent Systems and Technology*, 11(5):1–46, 2020.
- Wong, E., Rice, L., and Kolter, J. Z. Fast is better than free: Revisiting adversarial training. In *ICLR*, 2020.
- Yin, D., Lopes, R. G., Shlens, J., Cubuk, E. D., and Gilmer, J. A fourier perspective on model robustness in computer vision. In *NeurIPS*, 2019.
- Yoon, J., Hwang, S. J., and Lee, J. Adversarial purification with score-based generative models. In Meila, M. and Zhang, T. (eds.), *ICML*, volume 139, pp. 12062–12072. PMLR, 2021.
- Zhang, K., Zuo, W., Chen, Y., Meng, D., and Zhang, L. Beyond a gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017.
- Zhang, M., Levine, S., and Finn, C. MEMO: test time robustness via adaptation and augmentation. In *NeurIPS, Workshop on Distribution Shifts*, 2021.

A. Theoretical Analysis of DRQ for Binary Classification

In this section we perform a theoretical analysis of DRQ in the setting of binary classification, as it was illustrated in Figure 1. The general case of more than two classes is treated analogously by a one-versus-all approach. In the binary situation the DRQ algorithm can be reformulated as follows: We consider a differentiable binary classifier $u : \mathbb{R}^d \rightarrow [0, 1]$ for classifying data into the classes 0 and 1 with the label decision $\text{round}(u(x))$. Here $\text{round}(t)$ for $t \in [0, 1]$ is the rounding operation, defined as $\text{round}(t) = 0$ if $t < 0.5$ and $\text{round}(t) = 1$ if $t \geq 0.5$. We, furthermore, fix a calibration value of $\varepsilon > 0$ and a parameter $\alpha \in [0, 1]$. We define the exploration and quantification steps associated with an input $x \in \mathbb{R}^d$ as

$$\begin{aligned} \tilde{x}_0 &\in \arg \min_{\tilde{x} \in B_p(x; \varepsilon)} u(\tilde{x}), & \tilde{x}_1 &\in \arg \max_{\tilde{x} \in B_p(x; \varepsilon)} u(\tilde{x}), \\ \hat{x}_0^\alpha &\in \arg \max_{\hat{x} \in B_p(\tilde{x}_0; \alpha \varepsilon)} u(\hat{x}), & \hat{x}_1^\alpha &\in \arg \min_{\hat{x} \in B_p(\tilde{x}_1; \alpha \varepsilon)} u(\hat{x}). \end{aligned}$$

The label decision is now taken based on which of the two points \hat{x}_0^α and \hat{x}_1^α realizes the smaller value in

$$\min \{u(\hat{x}_0^\alpha), 1 - u(\hat{x}_1^\alpha)\}.$$

In other words, the DRQ classifier with parameter $\alpha \in [0, 1]$ acts as follows:

$$u_{\text{DRQ}}^\alpha(x) = \begin{cases} u(\hat{x}_0^\alpha), & \text{if } u(\hat{x}_0^\alpha) \leq 1 - u(\hat{x}_1^\alpha), \\ u(\hat{x}_1^\alpha), & \text{if } u(\hat{x}_0^\alpha) > 1 - u(\hat{x}_1^\alpha), \end{cases} \quad F_{\text{DRQ}}^\alpha(x) = \text{round}(u_{\text{DRQ}}^\alpha(x)). \quad (6)$$

We can prove the following statement regarding the behavior of DRQ with respect to the parameter $\alpha \in [0, 1]$.

Proposition A.1. *If $0 \leq \underline{\alpha} \leq \bar{\alpha} \leq 1$ then for all $x \in \mathbb{R}^d$ it holds*

$$u(\hat{x}_0^{\underline{\alpha}}) \leq u(\hat{x}_0^{\bar{\alpha}}) \quad \text{and} \quad u(\hat{x}_1^{\underline{\alpha}}) \geq u(\hat{x}_1^{\bar{\alpha}}). \quad (7)$$

Furthermore, we have the following two estimates for the special cases $\underline{\alpha} = 0$ and $\bar{\alpha} = 1$:

$$u(\hat{x}_0^0) \leq u(x) \leq u(\hat{x}_0^1) \quad \text{and} \quad u(\hat{x}_1^0) \geq u(x) \geq u(\hat{x}_1^1). \quad (8)$$

Proof. Since $\underline{\alpha} \leq \bar{\alpha}$ it holds

$$u(\hat{x}_0^{\underline{\alpha}}) = \max_{B_p(\tilde{x}_0; \underline{\alpha} \varepsilon)} u \leq \max_{B_p(\tilde{x}_0; \bar{\alpha} \varepsilon)} u = u(\hat{x}_0^{\bar{\alpha}}).$$

Analogously, it holds

$$u(\hat{x}_1^{\underline{\alpha}}) = \min_{B_p(\tilde{x}_1; \underline{\alpha} \varepsilon)} u \geq \min_{B_p(\tilde{x}_1; \bar{\alpha} \varepsilon)} u = u(\hat{x}_1^{\bar{\alpha}}).$$

This proves the inequalities (7). Regarding the special cases one notes that, since $B_p(\tilde{x}_i, 0\varepsilon) = \{\tilde{x}_i\}$, it obviously holds $\hat{x}_i^0 = \tilde{x}_i$ for $i \in \{0, 1\}$. Hence, it holds

$$u(\hat{x}_0) = u(\tilde{x}_0) = \min_{B_p(x; \varepsilon)} u \leq u(x) \quad \text{and} \quad u(\hat{x}_1) = u(\tilde{x}_1) = \max_{B_p(x; \varepsilon)} u \geq u(x),$$

and analogously

$$u(\hat{x}_0^1) = \max_{B_p(\tilde{x}_0, \varepsilon)} u \geq u(x) \quad \text{and} \quad u(\hat{x}_1^1) = \min_{B_p(\tilde{x}_1, \varepsilon)} u \leq u(x),$$

where we utilized that, since $\tilde{x}_i \in B_p(x; \varepsilon)$, it also holds $x \in B_p(\tilde{x}_i, \varepsilon)$ for $i \in \{0, 1\}$. This proves (8). \square

Inequality (7) in Proposition A.1 shows that with increasing values of α , the confidences of the DRQ algorithm decrease for both classes. Furthermore, for $\alpha = 0$ DRQ increases confidences compared to the original binary classifier u , whereas for $\alpha = 1$ confidences are decreased. To prevent DRQ from being over- or underconfident, it therefore seems natural to choose $\alpha = 0.5$, which is also confirmed by our empirical results.

We now theoretically investigate the potential of DRQ to smooth out spurious misclassifications and start with the following definition.

Definition A.2. Let $u : \mathbb{R}^d \rightarrow \mathbb{R}$ be a function and $\varepsilon > 0$. A point $x \in \mathbb{R}^d$ is called ε -wide local minimum if

$$u(x) \leq u(\tilde{x}) \quad \forall \tilde{x} \in B_p(x; \varepsilon) \setminus \{x\}.$$

Analogously, x is called ε -wide local maximum if

$$u(x) \geq u(\tilde{x}) \quad \forall \tilde{x} \in B_p(x; \varepsilon) \setminus \{x\}.$$

An ε -wide local minimum or maximum is called *strict* if these inequalities are strict.

Note that every $\bar{\varepsilon}$ -wide local extremum is also an $\underline{\varepsilon}$ -wide local extremum for $0 < \underline{\varepsilon} < \bar{\varepsilon}$. Furthermore, u can have several other local extrema in the ball $B_p(x; \varepsilon)$ which does not change the fact that x is an ε -wide local extremum. For simplicity we restrict the following discussion to local minima of a classifier $u : \mathbb{R}^d \rightarrow [0, 1]$.

The first observation is a trivial one:

Proposition A.3. Let $x \in \mathbb{R}^d$ be a 2ε -wide local minimum of the classifier $u : \mathbb{R}^d \rightarrow [0, 1]$ and let $\alpha \in [0, 1]$. Then it holds

$$u(\hat{x}_0^\alpha) \geq u(x) \quad \text{and} \quad u(\hat{x}_1^\alpha) \geq u(x) \quad (9)$$

and therefore

$$u_{\text{DRQ}}^\alpha(x) \geq u(x). \quad (10)$$

Proof. Since $\tilde{x}_i \in B_p(x; \varepsilon)$ and therefore $\hat{x}_i^\alpha \in B_p(\tilde{x}_i; \alpha\varepsilon) \subset B_p(x; 2\varepsilon)$ it trivially holds

$$u(\hat{x}_0^\alpha) = \max_{B_p(\tilde{x}_0; \alpha\varepsilon)} u \geq u(x), \quad u(\hat{x}_1^\alpha) = \min_{B_p(\tilde{x}_1; \alpha\varepsilon)} u \geq u(x),$$

which implies $u_{\text{DRQ}}^\alpha(x) \geq u(x)$. \square

For a more precise analysis, we have to quantify the sharpness of a local minimum. If $x \in \mathbb{R}^d$ is a ε -wide local minimum of u , we measure its sharpness by the non-negative and non-decreasing function

$$g_\varepsilon : [0, \infty) \rightarrow [0, 1], \quad g_\varepsilon(\alpha) := \max_{B_p(x; \alpha\varepsilon)} u - \min_{B_p(x; \varepsilon)} u. \quad (11)$$

Note that $g_\varepsilon(0) = 0$ and $g_\varepsilon(1)$ describes the oscillation of u on $B_p(x; \varepsilon)$. We now prove that, if the classifier u is strongly convex around a local minimum, then DRQ pushes the classification of x from class 0 to class 1 based on the sharpness of the local minimum, measured through (11) and the constant of strict convexity.

Proposition A.4. Let $x \in \mathbb{R}^d$ be a 2ε -wide local minimum of the classifier $u : \mathbb{R}^d \rightarrow [0, 1]$ and let $\alpha \in [0, 1]$. Assume that u is μ -strongly convex on $B_p(x; 2\varepsilon)$, meaning that there exists a constant $\mu > 0$ such that

$$u(\hat{x}) \geq u(\tilde{x}) + \langle \nabla u(\tilde{x}), \hat{x} - \tilde{x} \rangle + \mu \|\hat{x} - \tilde{x}\|_p^2, \quad \forall \hat{x}, \tilde{x} \in B_p(x; 2\varepsilon). \quad (12)$$

Then it holds

$$u_{\text{DRQ}}^\alpha(x) \geq \min \{u(x) + g_\varepsilon(\alpha), u(x) + \mu(1 - \alpha)^2 \varepsilon^2\}. \quad (13)$$

Proof. Since u is μ -strongly convex, x is in fact a strict ε -wide local minimum, meaning $\arg \min_{B_p(x; \varepsilon)} u = \{x\}$. Hence, we get $\tilde{x}_0 = x$ and by definition of g_ε it therefore holds

$$u(\hat{x}_0^\alpha) = \max_{B_p(\tilde{x}_0; \alpha\varepsilon)} u = \max_{B_p(x; \alpha\varepsilon)} u = u(x) + g_\varepsilon(\alpha). \quad (14)$$

Using the strong convexity and the fact that $\nabla u(x) = 0$ —since x is a local minimum—it also holds

$$u(\hat{x}_1^\alpha) \geq u(x) + \underbrace{\langle \nabla u(x), \hat{x}_1^\alpha - x \rangle}_{=0} + \mu \|\hat{x}_1^\alpha - x\|_p^2 = u(x) + \mu \|\hat{x}_1^\alpha - x\|_p^2.$$

Now by definition it holds $\tilde{x}_1 \in \arg \max_{B_p(x; \varepsilon)} u$ and $\hat{x}_1^\alpha \in \arg \min_{B_p(\tilde{x}_1, \alpha\varepsilon)} u$. The strong convexity of u implies that \tilde{x}_1 lies on the boundary of $B_p(x; \varepsilon)$. Hence, \hat{x}_1^α has distance at least $(1 - \alpha)\varepsilon$ from x which can be seen as follows:

$$\|\hat{x}_1^\alpha - x\|_p \geq \|\tilde{x}_1 - x\|_p - \|\hat{x}_1^\alpha - \tilde{x}_1\|_p \geq \varepsilon - \alpha\varepsilon = (1 - \alpha)\varepsilon.$$

Using this in the estimate above we find

$$u(\hat{x}_1^\alpha) \geq u(x) + \mu(1 - \alpha)^2\varepsilon^2 \tag{15}$$

and, combining (14) and (15), the DRQ classifier (6) satisfies (13) as claimed. □

Again, the bound in (13) shows that $\alpha = 0.5$ is a balanced choice: For $\alpha = 0$ the confidence of class 0 remains $u(x)$ where as the one for class 1 increases by $\mu\varepsilon^2$. In contrast for $\alpha = 1$ the confidence of class 0 decreases whereas the one for class 1 satisfies the bound $u(\hat{x}_1^1) \geq u(x)$. Together with the converse inequality from Proposition A.1 this implies $u(\hat{x}_1^1) = u(x)$, meaning that the confidence of class 1 remains constant. For in-between choices of α one can therefore expect that for spurious minima the confidence of class 0 decreases and the one for class 1 increases, which is the desired effect of DRQ.

The choice $\alpha = 0.5$ is further motivated by the following corollary.

Corollary A.5. *Under the conditions of Proposition A.4 it holds*

$$u_{\text{DRQ}}^\alpha(x) \geq \min \{ u(x) + \mu\alpha^2\varepsilon^2, u(x) + \mu(1 - \alpha)^2\varepsilon^2, \} \tag{16}$$

which for $\alpha = 0.5$ yields

$$u_{\text{DRQ}}^\alpha(x) \geq u(x) + \frac{\mu}{4}\varepsilon^2. \tag{17}$$

Proof. Using the strong convexity one can estimate the sharpness $g_\varepsilon(\alpha)$ from below. Let $\hat{x} \in \arg \max_{B_p(x; \alpha\varepsilon)} u$. Then it holds

$$g_\varepsilon(\alpha) = u(\hat{x}) - u(x) \geq \mu \|\hat{x} - x\|_p^2.$$

Since u is strongly convex, \hat{x} lies on the boundary of $B_p(x; \alpha\varepsilon)$ and hence

$$g_\varepsilon(\alpha) \geq \mu\alpha^2\varepsilon^2.$$

This then yields the estimates. □

B. DRQ Hyperparameters

B.1. Exploration and Quantification Iterations

Figure 3 shows the accuracy of DRQ for clean and adversarial data. Subfigure (a) shows the accuracy for a fixed number of quantification iterations and increasing exploration iterations. In contrast, the exploration iterations are fixed in subfigure (b) while the quantification iterations are increased.

B.2. Confidence Threshold

Figure 4 illustrates how the calibration distance ε_p is dependent on the confidence threshold c . The higher the confidence threshold, the harder it is for the FMN attack to find an adversarial example, and hence the larger the mean calibration distance becomes.

C. Denoising Autoencoder

For the experiment described in Section 5.1, we trained a denoising autoencoder based on the architecture proposed in (Zhang et al., 2017) on the training set of the CIFAR10 dataset. We used 17 layers and set the number of output channels of the first layer to 64. We used a batch size of 128 and an Adam optimizer with a learning rate of 0.001 Kingma & Ba (2015). We injected Gaussian noise to every image during training with a standard deviation of 0.1. We used the mean squared error loss and trained the model for 40 epochs, which was sufficient for convergence.

Decision Region Quantization

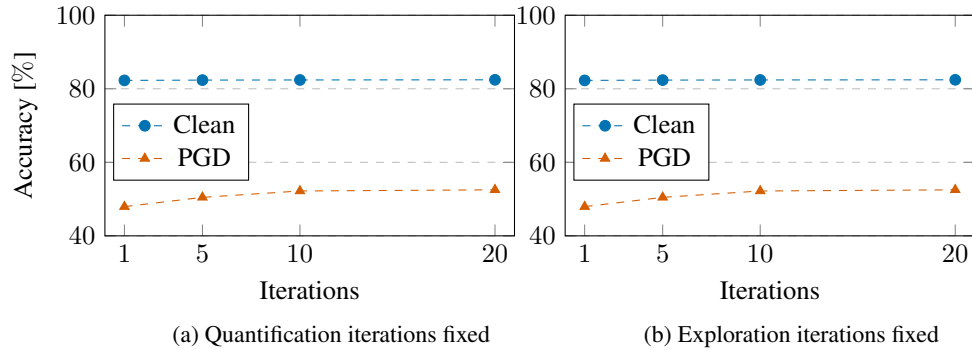


Figure 3: Accuracy for clean images and PGD attacked images for the model proposed by Wong et al. (2020) on the CIFAR10 dataset with the DRQ ℓ_∞ method. The graph shows the accuracy for: (a) A fixed number of 20 quantification step iterations and variable exploration iterations and (b) A fixed number of 20 exploration step iterations and variable quantification iterations. The performance increases monotonically with the amount of iterations till it converges at about 20 iterations for both steps.

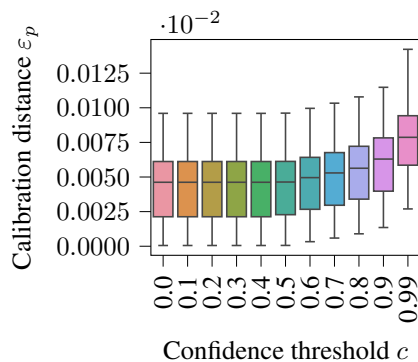


Figure 4: Calibration distance ε_p for different confidence thresholds c for the non-adversarially trained CIFAR10 model proposed in (Hendrycks & Dietterich, 2019). The calibration distance ε_p remains largely constant for values $c \leq 0.5$ and rapidly increases for $c > 0.9$.

D. DRQ Additional Experiments

D.1. Normally Trained Models

We additionally evaluated DRQ on networks trained without additional data augmentation methods or adversarial training. We observed no considerable differences in robustness or clean accuracy for any of the datasets and models. A more in-depth analysis revealed that the ε_p determined during the calibration step was orders of magnitude smaller than the one obtained for the adversarially trained models. We hypothesize that networks must exhibit at least some robustness to distribution shifts in order for DRQ to affect the robustness.

D.2. Accuracy on Clean Data

In Table 6, we investigate the difference in accuracy on clean data for different models and datasets for the two DRQ approaches and standard inference. We observe that DRQ has no considerable effect on the clean accuracy of the different models and will explore this behavior in future work.

E. Attack Success Rate

The success rate of the individual attacks is exemplified in Table 7 for the model proposed in (Wong et al., 2020). Furthermore, we attack the whole DRQ pipeline by including DRQ in the forward and backward pass of an PGD-based adversarial attack.

Decision Region Quantification

Table 6: Accuracy [%] of standard inference and DRQ inference in two different norms. The accuracy is shown for differently trained models on the CIFAR10, and ImageNet datasets. The best accuracy is shown in **bold** and the second best is underlined.

Models	Training	Standard	DRQ ℓ_∞	DRQ ℓ_2
CIFAR10				
Wong et al. (2020)	ℓ_∞ -AT	83.34	82.44	<u>83.11</u>
Gowal et al. (2020)	ℓ_∞ -AT	<u>89.48</u>	89.22	89.62
Rice et al. (2020)	ℓ_2 -AT	<u>88.67</u>	88.63	88.84
Hendrycks et al. (2021a)	No-AT	<u>95.08</u>	<u>95.07</u>	95.28
ImageNet				
Wong et al. (2020)	ℓ_∞ -AT	53.8	53.8	53.9
Hendrycks et al. (2021a)	No-AT	75.8	75.8	75.8

To this end, we substitute the sign function of the gradient in the different steps with the gradient to make the whole pipeline differentiable (this does not affect the performance of DRQ). We further skipped the calibration step in these experiments and used a predefined ε_p . For the same amount of attack iterations, this attack achieved a marginally better success rate than our adaptive attack (+0.11% success rate for the model proposed by (Wong et al., 2020)). However, the proposed adaptive attacks were considerably more effective with the same amount of computational complexity. We explored the behavior of the end-to-end attack and observed a vanishing/exploding gradient problem due to the size of the computational graph. This phenomenon was also observed by Athalye et al. (2018a), when attacking models that include an iterative process. We further explored using the score-based black-box square attack on the whole pipeline but this was not more effective.

Table 7: Accuracy [%] of standard inference and DRQ inference in two different norms against different adversarial attacks for the CIFAR10 model proposed by (Wong et al., 2020). The worst-case accuracy of the combined attacks is shown in the rightmost column.

Attacks	Clean	FAB	Square	APGD ℓ_2	APGD ℓ_∞	APGD ℓ_1	APGD ℓ_0	APGD ℓ_2	Worst-case
Standard	83.34	44.35	51.22	61.38	45.83	46.77	53.59	47.05	43.37
DRQ ℓ_2	83.17	70.19	70.0	65.36	49.93	48.68	54.05	51.64	43.39
DRQ ℓ_∞	81.16	71.43	72.28	72.02	58.97	52.62	53.73	62.27	48.52

F. DRQ Exploration Directions

Figure 5 shows, the cosine similarity between DRQ direction $\gamma_{\text{DRQ}} = \hat{x}_{\text{DRQ}} - x$ and the adversarial perturbation direction $\gamma = x_{\text{adv}} - x$ for the model proposed by Wong et al. (2020), as described in Section 5.3. The mean cosine similarity of all attacks is closer or equal to zero. The lowest mean cosine similarity is displayed by the adaptive PGD Noise attack. However, for this attack DRQ was not able to correct the prediction as often (+5.85% accuracy with DRQ) as for example for the other attacks (e.g., Square: +21.06% accuracy with DRQ). We argue that the perturbations found by the adaptive PGD Noise attack are more robust. Thus, there exist less perturbation directions that can correct the prediction and DRQ is forced to find the opposing direction to the adversarial perturbation γ , such that $-\gamma_{\text{DRQ}} \approx \gamma$. This leads to a lower cosine similarity for the PGD Noise attack in 5.

G. Exploration and Quantification Step

We conducted additional experiments to assess if the adversarial optimization in the exploration and quantification steps of DRQ are necessary to increase robustness. To this end, we replaced the exploration step with random noise augmentations, which we describe as *random exploration*. We also replaced the quantification step with random noise augmentations. Instead of searching for the sample that shows the lowest robustness, we inject 20 different random noise augmentations to every input and select the augmented sample with the lowest robustness. We refer to this procedure as *random quantification*. Table 8 summarizes the results of replacing either the exploration or the quantification step or both with their random versions. The experiment was conducted for the model proposed in (Wong et al., 2020) on the CIFAR10-C dataset and against the adversarial attack ensemble on CIFAR10. Replacing either the exploration or quantification step reduces the

Decision Region Quantification

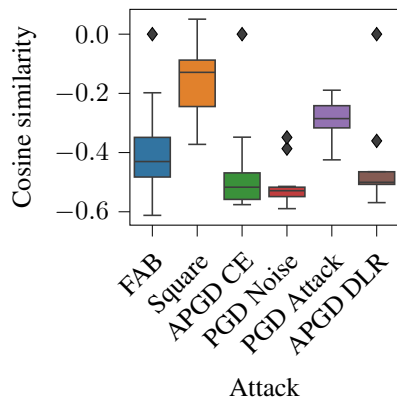


Figure 5: Cosine similarity between the DRQ direction $\gamma_{\text{DRQ}} = \hat{x}_{\text{DRQ}} - x$ and the adversarial perturbation direction $\gamma = x_{\text{adv}} - x$ for the CIFAR10 model proposed in (Wong et al., 2020). Negative values indicate that DRQ moves the adversarial example closer to the original input.

robustness in all cases. This highlights the importance of accurately estimating the robustness of the local decision regions in both steps of the DRQ algorithm.

Table 8: Ablation study on the accuracy differences [%] between DRQ and standard prediction for different region exploration and region quantification approaches for the model proposed in (Wong et al., 2020). Here, either the exploration or quantification step or both are replaced by a random approximation.

Data	CIFAR10-C		CIFAR10 worst-case	
	Exploration - Random	Exploration - ℓ_∞	Exploration - Random	Exploration- ℓ_∞
Quantification - Random	-0.1	-0.5	-0.75	-0.27
Quantification - ℓ_∞	-0.25	+2.03	-0.35	+5.85