

Defense Against Adversarial Images using Web-Scale Nearest-Neighbor Search

Abhimanyu Dubey^{*1,2}, Laurens van der Maaten², Zeki Yalniz², Yixuan Li², and Dhruv Mahajan²

¹Massachusetts Institute of Technology

²Facebook AI

Abstract

A plethora of recent work has shown that convolutional networks are not robust to adversarial images: images that are created by perturbing a sample from the data distribution as to maximize the loss on the perturbed example. In this work, we hypothesize that adversarial perturbations move the image away from the image manifold in the sense that there exists no physical process that could have produced the adversarial image. This hypothesis suggests that a successful defense mechanism against adversarial images should aim to project the images back onto the image manifold. We study such defense mechanisms, which approximate the projection onto the unknown image manifold by a nearest-neighbor search against a web-scale image database containing tens of billions of images. Empirical evaluations of this defense strategy on ImageNet suggest that it is very effective in attack settings in which the adversary does not have access to the image database. We also propose two novel attack methods to break nearest-neighbor defenses, and demonstrate conditions under which nearest-neighbor defense fails. We perform a series of ablation experiments, which suggest that there is a trade-off between robustness and accuracy in our defenses, that a large image database (with hundreds of millions of images) is crucial to get good performance, and that careful construction the image database is important to be robust against attacks tailored to circumvent our defenses.

1. Introduction

A range of recent studies has demonstrated that many modern machine-learning models are not robust to *adversarial examples*: examples that are intentionally designed to be misclassified by the models, whilst being nearly indistinguishable from regular examples in terms of some distance measure. Whilst adversarial examples have been con-

structed against speech recognition [3] and text classification [5] systems, most recent work on creating adversarial examples has focused on computer vision [2, 9, 17, 19, 21, 32], in which adversarial images are often perceptually indistinguishable from real images. Such adversarial images have successfully fooled systems for image classification [32], object detection [38], and semantic segmentation [6]. In practice, adversarial images are constructed by maximizing the loss of the machine-learning model (such as a convolutional network) with respect to the input image, starting from a “clean” image. This maximization creates an adversarial perturbation of the original image; the perturbation is generally constrained or regularized to have a small ℓ_p -norm in order for the adversarial image to be perceptually (nearly) indistinguishable from the original.

Because of the way they are constructed, many adversarial images are different from natural images in that there exists no physical process by which the images could have been generated. Hence, if we view the set of all possible natural images¹ as samples from a manifold that is embedded in image space, many adversarial perturbations may be considered as transformations that take a sample from the image manifold and move it away from that manifold. This hypothesis suggests an obvious approach for implementing *defenses* that aim to increase the robustness of machine-learning models against “off-manifold” adversarial images [24, 39]: *viz.*, projecting the adversarial images onto the image manifold before using them as input into the model.

As the true image manifold is unknown, in this paper, we develop defenses that approximate the image manifold using a massive database of tens of billions of web images. Specifically, we approximate the projection of an adversarial example onto the image manifold by the finding nearest neighbors in the image database. Next, we classify the “projection” of the adversarial example, *i.e.*, the identified nearest neighbor(s), rather than the adversarial example itself. Using modern techniques for distributed approximate nearest-neighbor search to make this strategy practical, we

^{*}This work was done while Abhimanyu Dubey was at Facebook AI.

¹For simplicity, we ignore synthetic images such as drawings.

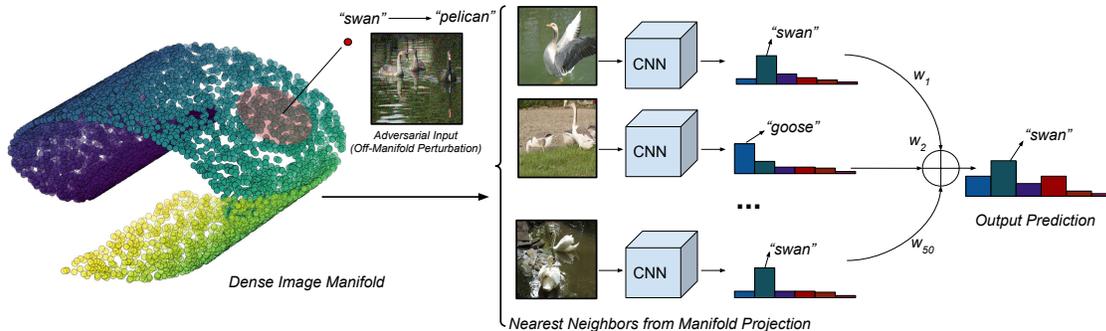


Figure 1. Illustration of our defense procedure for improving adversarial robustness in image classification. We first “project” the image on to the image manifold by finding the nearest neighbors in the image database, followed by a weighted combination of the predictions of this nearest neighbor set to produce our final prediction.

demonstrate the potential of our approach in ImageNet classification experiments. Our contributions are:

1. We demonstrate the feasibility of web-scale nearest-neighbor search as a defense mechanism against a variety of adversarial attacks in both *gray-box* and *black-box* attack settings, on an image database of an unprecedented scale (~ 50 billion images). We achieve robustness comparable to prior state-of-the-art techniques in *gray-box* and *black-box* attack settings in which the adversary is unaware of the defense strategy.
2. To analyze the performance of our defenses in *white-box* settings in which the adversary has full knowledge of the defense technique used, we develop two novel attack strategies designed to break our nearest-neighbor defenses. Our experiments with these attacks show that our defenses break in pure white-box settings, but remain effective in attack settings in which the adversary has access to a comparatively small image database and the defense uses a web-scale image database, even when architecture and model parameters are available to the adversary.

We also conduct a range of ablation studies, which show that: (1) nearest-neighbor predictions based on earlier layers in a convolutional network are more robust to adversarial attacks and (2) the way in which the image database for nearest-neighbor search is constructed substantially influences the robustness of the resulting defense.

2. Related Work

After the initial discovery of adversarial examples [32], several adversarial attacks have been proposed that can change model predictions by altering the image using a perturbation with small ℓ_2 or ℓ_∞ norm [2, 9, 17, 19, 21]. In particular, [19] proposed a general formulation of the fast gradient-sign method based on projected gradient descent (PGD), which is currently considered the strongest attack.

A variety of defense techniques have been studied that aim to increase adversarial robustness². *Adversarial training* [9, 12, 16, 17] refers to techniques that train the network with adversarial examples added to the training set. Defensive distillation [25, 22] tries to increase robustness to adversarial attacks by training models using model distillation. Input-transformation defenses try to remove adversarial perturbations from input images via JPEG compression, total variation minimization, or image quilting [4, 10]. Certifiable defense approaches [29, 27] aim to guarantee robustness under particular attack settings. Other studies have used out-of-distribution detection approaches to detect adversarial examples [18]. Akin to our approach, PixelDefend [30] and Defense-GAN [28] project adversarial images back onto the image manifold, but they do so using parametric density models rather than a non-parametric one.

Our work is most closely related to the nearest-neighbor defenses of [24, 39]. [39] augments the convolutional network with an off-the-shelf image retrieval system to mitigate the adverse effect of “off-manifold” adversarial examples, and uses local mixup to increase robustness to “on-manifold” adversarial examples. In particular, inputs are projected onto the feature-space convex hull formed by the retrieved neighbors using trainable projection weights; the feature-producing convolutional network and the projection weights are trained jointly. In contrast to [39], our approach does not involve alternative training procedures and we do not treat on-manifold adversarial images separately [8].

3. Problem Setup

We consider multi-class image classification of images $\mathbf{x} \in [0, 1]^{H \times W}$ into one of C classes. We assume we are given a labeled training set with N examples, $\mathcal{D} = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_N, y_N)\}$ with labels $y \in \mathbb{Z}_C$. Training a classification model amounts to selecting a hypothesis $h(\mathbf{x}) \rightarrow \mathbb{Z}_C$ from some hypothesis set \mathcal{H} . The hypothe-

²See <https://www.robust-ml.org/defenses/> for details.

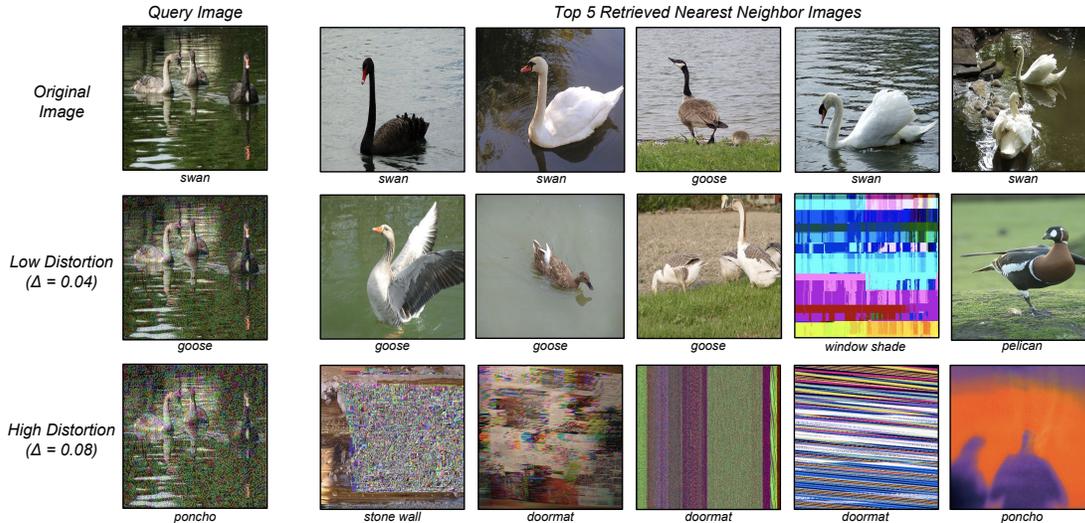


Figure 2. Visualization of an image and its five nearest neighbors in the *YFCC-100M* database (from left to right) based on `conv_5_1` features for a clean image (top), an image with a small adversarial perturbation ($\Delta = 0.04$; center), and an image with a large adversarial perturbation ($\Delta = 0.08$; bottom). Adversarial images generated using PGD with a ResNet-50 trained on ImageNet.

sis set \mathcal{H} is the set of all possible parameter values for a convolutional network architecture (such as a ResNet), and the hypothesis $h(\mathbf{x})$ is selected using empirical risk minimization: specifically, we minimize the sum of a loss function $L(\mathbf{x}_n, y_n; h)$ over all examples in \mathcal{D} (we omit h where it is obvious from the context). Throughout the paper, we choose $L(\cdot, \cdot; \cdot)$ to be the multi-class logistic loss.

3.1. Attack Model

Given the selected hypothesis (*i.e.*, the model) $h \in \mathcal{H}$, the adversary aims to find an adversarial version \mathbf{x}^* of a real example \mathbf{x} for which: (1) \mathbf{x}^* is similar to \mathbf{x} under some distance measure and (2) the loss $L(h(\mathbf{x}^*), y)$ is large, *i.e.*, the example \mathbf{x}^* is likely to be misclassified. In this paper, we measure similarity between \mathbf{x}^* and \mathbf{x} by the normalized ℓ_2 distance³, given by $\Delta(\mathbf{x}, \mathbf{x}^*) = \frac{\|\mathbf{x} - \mathbf{x}^*\|_2}{\|\mathbf{x}\|_2}$. Hence, the adversary’s goal is to find for some similarity threshold ϵ :

$$\mathbf{x}^* = \arg \max_{\mathbf{x}': \Delta(\mathbf{x}, \mathbf{x}') \leq \epsilon} L(\mathbf{x}', y; h).$$

Adversarial attacks can be separated into three categories: (1) *white-box attacks*, where the adversary has access to both the model h and the defense mechanism; (2) *black-box attacks*, where the adversary has no access to h nor the defense mechanism; and (3) *gray-box attacks* in which the adversary has no direct access to h but has partial information of the components that went into the construction of h , such as the training data \mathcal{D} , the hypothesis set \mathcal{H} , or a superset of the hypothesis set \mathcal{H} . While robustness against *white-box* adversarial attacks is desirable since it is the strongest

notion of security [23], in real-world settings, we are often interested in robustness against *gray-box* attacks because it is rare for an adversary to have complete information (*cf.* *white-box*) or no information whatsoever (*cf.* *black-box*) on the model it is attacking [13].

3.2. Adversarial Attack Methods

The **Iterative Fast Gradient Sign Method (I-FGSM)** [17] generates adversarial examples by iteratively applying the following update for $m = \{1, \dots, M\}$ steps:

$$\mathbf{x}^{(m)} = \mathbf{x}^{(m-1)} + \epsilon \cdot \text{sign} \left(\nabla_{\mathbf{x}^{(m-1)}} L(\mathbf{x}^{(m-1)}, y) \right),$$

where $\mathbf{x}_{\text{IFGSM}}^* = \mathbf{x}^{(M)}$, and $\mathbf{x}^{(0)} = \mathbf{x}$.

When the model is available to the attacker (*white-box* setting), the attack can be run using the true gradient $\nabla_{\mathbf{x}} L(h(\mathbf{x}), y)$, however, in *gray-box* and *black-box* settings, the attacker has access to a surrogate gradient $\nabla_{\mathbf{x}} L(h'(\mathbf{x}), y)$, which in practice, has been shown to be effective as well. The **Projected Gradient Descent (PGD)** [19] attack generalizes the I-FGSM attack by: (1) clipping the gradients to project them on the constraints formed by the similarity threshold and (2) including random restarts in the optimization process. Throughout the paper, we employ the PGD attack in our experiments because recent benchmark competitions suggest it is currently the strongest attack method.

In the appendix, we also show results with **Fast Gradient Sign Method (FGSM)** [9] and **Carlini-Wagner’s ℓ_p (CW-Lp)** [2] attack methods. For all the attack methods, we use the implementation of [10] and enforce that the im-

³Other choices for measuring similarity include the ℓ_∞ metric [37].

age remains within $[0, 1]^{H \times W}$ by clipping pixel values to lie between 0 and 1.

4. Adversarial Defenses via Nearest Neighbors

The underlying assumption of our defense is that adversarial perturbations move the input image away from the image manifold. The goal of our defense is to project the images back onto the image manifold before classifying them. As the true image manifold is unknown, we use a sample approximation comprising a database of billions of natural images. When constructing this database, the images may be selected in a weakly-supervised fashion to match the target task, for instance, by including only images that are associated with labels or hashtags that are relevant to that task [20]. To “project” an image on the image manifold, we identify its K nearest neighbors from the database by measuring Euclidean distances in some feature space. Modern implementations of approximate nearest neighbor search allow us to do this in milliseconds even when the database contains billions of images [14]. Next, we classify the “projected” adversarial example by classifying its nearest neighbors using our classification model and combining the resulting predictions. In practice, we pre-compute the classifications for all images in the image database and store them in a key-value map [7] to make prediction efficient.

We combine predictions by taking a weighted average of the *softmax probability vector* of all the nearest neighbors⁴. The final class prediction is the arg max of this average vector. We study three strategies for weighting the importance of each of the K predictions in the overall average:

Uniform weighting (UW) assigns the same weight ($w = 1/K$) to each of the predictions in the average.

We also experimented with two confidence-based weighting schemes that take into account the “confidence” that the classification model has in its prediction for a particular neighbor. This is important because, empirically, we observe that “spurious” neighbors exist that do not correspond to any of the classes under consideration, as displayed in Figure 2 (center row, fourth retrieved image). The entropy of the softmax distribution for such neighbors is very high, suggesting we should reduce their contribution to the overall average. We study two measures for compute the weight, w , associated with a neighbor: (1) an entropy-based measure, CBW-E(ntropy); and (2) a measure for diversity among the top-scoring classes, CBW-D(iversity).

CBW-E measures the entropy gap between a class prediction and the entropy of a uniform prediction. Hence, for softmax vector \mathbf{s} over C classes ($\forall c \in \{1, \dots, C\} : s_c \in$

$[0, 1]$ and $\sum_{c \in \{1, \dots, C\}} s_c = 1$), the weight w is given by:

$$w = \left| \log C + \sum_{c=1}^C s_c \log s_c \right|.$$

CBW-D computes w as a function of the difference between the maximum value of the softmax distribution and the next top M values. Specifically, let $\hat{\mathbf{s}}$ be the sorted (in descending order) version of the softmax vector \mathbf{s} . The weight w is defined as:

$$w = \sum_{m=2}^{M+1} (\hat{s}_1 - \hat{s}_m)^P.$$

We tuned M and P using cross-validation in preliminary experiments, and set $M = 20$ and $P = 3$ in all experiments that we present in the paper.

5. Experiments: Gray and Black-Box Settings

To evaluate the effectiveness of our defense strategy, we performed a series of image-classification experiments on the ImageNet dataset. Following [16], we assume an adversary that uses the state-of-the-art PGD adversarial attack method (see Section 3.2) with 10 iterations. In the appendix, we also present results obtained using other attack methods.

5.1. Experimental Setup

To perform image classification, we use ResNet-18 and ResNet-50 models [11] that were trained on the ImageNet training set. We consider two different attack settings: (1) a *gray-box attack* setting in which the model used to generate the adversarial images is the same as the image-classification model, *viz.* the ResNet-50; and (2) a *black-box attack* setting in which the adversarial images are generated using the ResNet-18 model and the prediction model is ResNet-50 (following [10]). We experiment with a number of different implementations of the nearest-neighbor search defense strategy by varying: (1) the image database that is queried by the defense and (2) the features that are used as basis for the nearest-neighbor search.

Image database. We experiment with three different web-scale image databases as the basis for our nearest-neighbor defense.

- *IG-N-** refers to a database of N public images with associated hashtags that is collected from a social media website, where \star can take two different values. Specifically, *IG-N-All* comprises images that were selected at random. Following [20], *IG-N-Targeted* contains exclusively images that were tagged with at least one 1, 500 hashtags that match one of the 1, 000 classes in the ImageNet-1K benchmark. The 1, 500 hashtags

⁴In preliminary experiments, we also tried averaging “hard” rather than “soft” predictions but we did not find that to work better in practice.

Image database	Clean			Gray box			Black box		
	UW	CBW-E	CBW-D	UW	CBW-E	CBW-D	UW	CBW-E	CBW-D
IG-50B-All (conv_5_1-RMAC)	0.632	0.644	0.676	0.395	0.411	0.427	0.448	0.459	0.491
IG-1B-Targeted (conv_5_1)	0.659	0.664	0.681	0.415	0.429	0.462	0.568	0.574	0.587
IN-1.3M (conv_5_1)	0.472	0.469	0.471	0.285	0.286	0.286	0.311	0.312	0.312

Table 1. ImageNet classification accuracies of a ResNet-50 using our nearest-neighbor defense with three different weighting strategies (UW, CBW-E, and CBW-D) on PGD adversarial ImageNet images with a normalized ℓ_2 distance of 0.06. Nearest-neighbor searches were performed on three image databases (rows) with $K = 50$. Accuracies using KNN defense on clean images are included for reference.

were obtained by canonicalizing all synonyms corresponding to the synsets, which is why the dataset contains more hashtags than classes. Our largest database contains $N = 50$ billion images.

- *YFCC-100M* is a publicly available dataset of 100 million Flickr images with associated meta-data [33].
- *IN-1.3M* refers to the training split of the publicly ImageNet-1K dataset of approximately 1.28M images.

Features. We constructed feature representations for the images in each of these image databases were constructed by: (1) computing pre-ReLU activations from the conv_2_3, conv_3_4, conv_4_6, or conv_5_1 layer of a ResNet-50 trained on ImageNet-1K and (2) reducing these feature representations to 256 dimensions using a spatial average pooling followed by PCA. For our largest database of 50 billion images, we used a feature representation that requires less storage: following [20], we use conv_5_1-RMAC features that were obtained by using conv_5_1 features from a ResNet-50 model followed by R-MAC pooling [34], bit quantization, and dimensionality reduction (see the appendix for details).

5.2. Results

Table 4 presents the classification accuracy of a ResNet-50 using our defense strategy on PGD adversarial ImageNet images with a normalized ℓ_2 dissimilarity of 0.06; the table presents results in both the gray-box and the black-box settings. The table presents results for both the uniform weighting (UW) and the confidence-based weighting (CBW-E and CBW-D) strategies, performing nearest-neighbor search in three different image databases using a value of $K = 50$. The results presented in the table demonstrate the potential of large-scale nearest neighbors as a defense strategy: our best model achieves a top-1 accuracy of 46.2% in the gray-box and of 58.7% in the black-box setting. The results also show that CBW-D weighting consistently outperforms the other weighting schemes, and that the use of web-scale IG- N -* databases with billions of images leads to substantially more effective defenses than using the ImageNet training set.

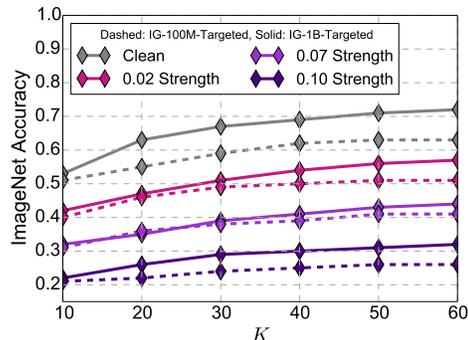


Figure 3. Classification accuracy of ResNet-50 using our CBW-D defense on PGD adversarial ImageNet images, as a function of the normalized ℓ_2 norm of the adversarial perturbation. Defenses are implemented via nearest-neighbor search using conv_5_1 features on the IG-1B-Targeted (solid lines) and IG-100M-Targeted (dashed lines). Results are for the black-box setting.

Motivated by its strong performance in our initial experiments, we use the CBW-D strategy for all the ablation studies and analyses that we perform below.

How does the choice of K influence the effectiveness of the defense? The robustness⁵ of nearest-neighbor algorithms depends critically on the number of nearest neighbors, K [36]. Motivated by this observation, we empirically analyze the effect of K on the effectiveness of our defense. Figure 3 presents the classification accuracy of a ResNet-50 with CBW-D using conv_5_1 features as a function K ; results are presented on two different databases, *viz.*, IG-100M-Targeted (dashed lines) and IG-1B-Targeted (solid lines). The results in the figure reveal increasing K appears to have a positive effect on classification accuracy for both databases, although the accuracy appears to saturate at $K = 50$. Therefore, we use $K = 50$ in the remainder of our experiments. The figure also shows that defenses based on the larger image database (1B images) consistently outperform those based on the smaller database (100M images).

⁵Recent work also proposed a nearest-neighbors algorithm that is claimed to be robust under adversarial perturbations [36]. We do not study that algorithm here because it does not scale to web-scale image datasets.

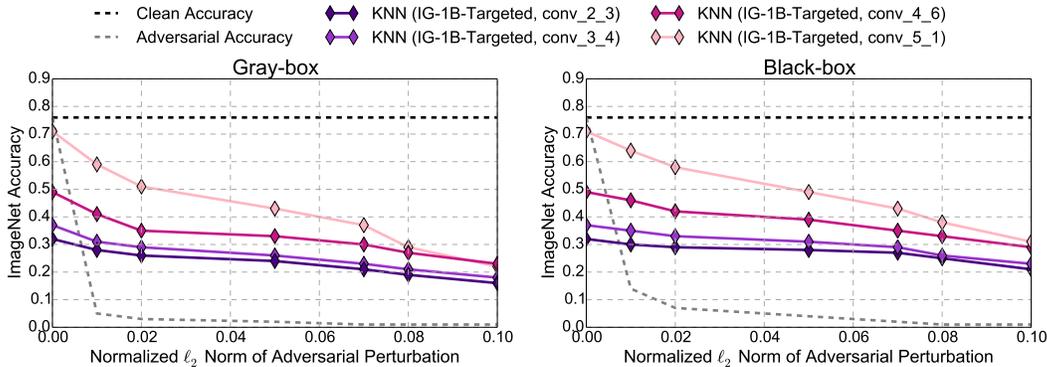


Figure 4. Classification accuracy of ResNet-50 using the CBW-D defense on PGD adversarial ImageNet images, as a function of the normalized ℓ_2 norm of the adversarial perturbation. Defenses use four different feature representations of the images in the IG-1B-Targeted image database. Results are presented for the gray-box (left) and black-box (right) settings.

How does the choice of features influence the effectiveness of the defense? Figure 4 presents classification accuracies obtained using CBW-D defenses based on four different feature representations of the images in the IG-1B-Targeted database. The results presented in the figure show that using “later” features for the nearest-neighbor search generally provide better accuracies, but that “earlier” features are more robust (*i.e.*, deteriorate less) in the perturbation norm regime we investigated. Earlier features presumably provide better robustness because they are less affected by the adversarial perturbation in the image, which induces only a small perturbation in those features. The downside of using early features is that they are more susceptible to non-semantic variations between the images, which is why the use of later features leads to higher classification accuracies. Motivated by these results, we use `conv_5_1` features in the remainder of our experiments.

How does nearest-neighbor index size influence the effectiveness of the defense? Next, we measured the effectiveness of our defense strategy as a function of the size of the image database used for nearest-neighbor searches. Figure 5 presents the effect of the image database size (N) on the classification accuracy for different attack strengths; experiments were performed on the IG- N -Targeted database. In line with earlier work [20, 31], the results suggest a log-linear relation between accuracy and database size: each time the database size is doubled, the classification accuracy increases by a fixed number of percentage points. This result appears to be consistent across a range of magnitudes of the adversarial perturbation.

How does selection of images in the index influence the effectiveness of the defense? In Figure 5, we also investigate how important it is that the images in the database used for the CBW-D defense are semantically related to the (ad-

versarial) images being classified, by comparing defenses based on the IG- N -All and IG- N -Targeted databases for varying index sizes. The results reveal that there is a positive effect of “engineering” the defense image database to match the task at hand: defenses based on IG- N -Targeted outperforms those based on IG- N -All by a margin of 1%–4% at the same database size.

5.3. Comparison with State-of-the-Art Defenses

In Table 2, we compare the effectiveness of our nearest-neighbor defense with that of other state-of-the-art defense strategies. Specifically, the table displays the classification accuracy on adversarial ImageNet images produced with PGD, using a normalized ℓ_2 distance of 0.06. The results in the table show that, despite its simplicity, our defense strategy is at least as effective as alternative approaches (including approaches that require re-training of the network). To the best of our knowledge, our defense strategy even outperforms the current state-of-the-art in the gray-box setting. In the black-box setting, image quilting [10] performs slightly better than our nearest-neighbor defense. Interestingly, image quilting is also a nearest-neighbor approach but it operates at the image patch level rather than at the image level.

6. Experiments: White-Box Setting

So far, we have benchmarked our adversarial defense technique against PGD attacks that are unaware of our defense strategies, *i.e.*, we have considered gray-box and black-box attack settings. However, in real-world attack settings, we may expect adversaries to be aware of defense mechanism and tailor their attacks to circumvent the defenses. In this section, we study such a *white-box* setting by designing adversarial attacks that try to circumvent nearest-neighbor defenses, and we study the effectiveness of these novel attacks against our defense strategy.

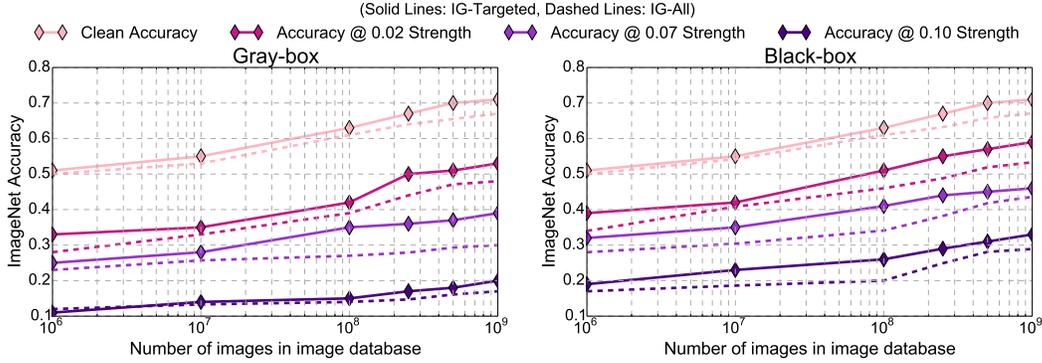


Figure 5. Classification accuracy of ResNet-50 using the CBW-D defense on PGD adversarial ImageNet images, using the IG- N -Targeted database (solid lines) and IG- N -All database (dashed lines) with different values of N . Results are presented in the gray-box (left) and black-box (right) settings.

Defense	Clean	Gray box	Black box
No defense	0.761	0.038	0.046
Crop ensemble [10]	0.652	0.456	0.512
TV Minimization [10]	0.635	0.338	0.597
Image quilting [10]	0.414	0.379	0.618
Ensemble training [35]	–	–	0.051
ALP [16]	0.557	0.279	0.348
RA-CNN [39]*	0.609	0.259	–
<i>Our Results</i>			
IG-50B-All (conv_5_1-RMAC)	0.676	0.427	0.491
IG-1B-Targeted (conv_5_1)	0.681	0.462	0.587
YFCC-100M (conv_5_1)	0.613	0.309	0.395
IN-1.3M (conv_5_1)	0.462	0.235	0.292

Table 2. ImageNet classification accuracies of ResNet-50 models using state-of-the-art defense strategies against the PGD attack, using a normalized ℓ_2 distance of 0.06. * RA-CNN [39] experiments were performed using a ResNet-18 model.

6.1. Defense-Aware Attacks

We develop two defense-aware attacks in which the adversary uses nearest-neighbor search on a web-scale image database to simulate the defense. Whilst this “attack database” may be identical to the “defense database”, it is likely that the databases are not exactly the same in practice (if only, because the defender can randomize its defense database). We develop two defense-aware attacks:

Nearest-neighbor prediction attack (PGD-PR). Given an attack database $\hat{\mathcal{D}}_A$ and a function $g(\mathbf{x})$ that the attacker uses to compute a feature representation for example \mathbf{x} , we first compute the corresponding set of K nearest neighbors, $\hat{\mathcal{D}}_{A,K}(g(\mathbf{x}))$. Subsequently, we add an extra loss term that maximizes the loss between those K nearest neighbors and the model prediction when constructing the adversarial sample \mathbf{x}^* . Specifically, we perform PGD-like updates:

$$\mathbf{x}^* = \mathbf{x} + \varepsilon \cdot \text{sign} \left[\nabla_{\mathbf{x}} L(h(\mathbf{x}), y) + \gamma \sum_{\mathbf{x}' \in \hat{\mathcal{D}}_{A,K}(\mathbf{x})} \nabla_{\mathbf{x}'} L(h(\mathbf{x}'), y) \right].$$

Herein, the hyperparameter γ trade off the loss suffered on the sample itself with the loss suffered on its nearest neighbors. We set $\gamma = 0.05$ and $K = 50$ in our paper. We perform

an iterative PGD-like update using this objective function, where we fix the neighbor set beforehand.

Nearest-neighbor feature-space attack (PGD-FS). In contrast to the previous attack, this attack targets the feature space used for nearest-neighbor retrieval. Specifically, it attacks the feature extractor $g(\mathbf{x})$ directly by producing the adversarial sample \mathbf{x}^* via PGD-like updates:

$$\mathbf{x}^* = \mathbf{x} + \varepsilon \cdot \text{sign} \left[\sum_{\mathbf{x}' \in \hat{\mathcal{D}}_{A,K}(\mathbf{x})} \nabla_{\mathbf{x}'} \|g(\mathbf{x}') - g(\mathbf{x})\|_2^2 \right].$$

6.2. Experiments

We perform experiments in which we evaluate the effectiveness of the PGD, the PGD-PR, and the PGD-FS attack. In all attacks, we follow [16] and set the number of attack iterations to 10, we set $K = 50$, and use conv_5_1 features to implement $g(\cdot)$.

We assume a pure white-box setting in which the adversary uses the defense image database (IG-1B-Targeted) as its attack database, and has access to both the ResNet-50 classification model and the feature-generating function $g(\cdot)$. Figure 6(a) shows the accuracy of a ResNet-50 using the CBW-D defense as a function of the attack strength for all three attacks. The results show that in this pure white-box setting, CBW-D only provides limited robustness against defense-aware attacks.

We perform ablation experiments for the PGD-FS attack to investigate the effect of: (1) the size of the attack image database and (2) the amount of overlap between the attack and defense databases. These ablation experiments emulate a gray-box attack scenario in which the attacker only has partial access to the image database used by the defender.

Effect of index size. We construct the attack databases of varying size by randomly selecting a subset from the 500M images in the defense index, and measure the accuracy of our models under PGD-FS attacks using these attack databases. Figure 6(b) displays accuracy as a function

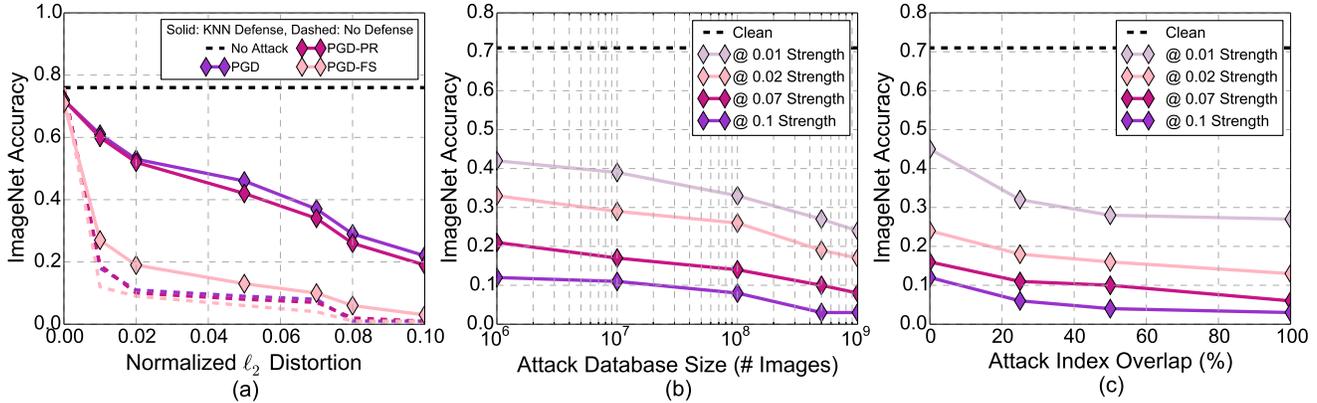


Figure 6. (a) Comparison of regular PGD, PGD-PR and PGD-FS on ResNet-50 in the *white-box* setting, i.e., attacker has access to the same database as the KNN defense (in this case, it is IG-1B-Targeted), (b) Variation of classification performance under KNN-PGD attack on ResNet-50 with IG-500M-Targeted defense database as the attacker database size increases, (c) Variation of classification performance under KNN-PGD attack on ResNet-50 using IG-500M-Targeted as defense database, as the overlap of images between the attacker database and the defense database increases. Note that in (c) the attack database size is always 500M.

of attack database size, showing that the performance of our defense degrades as the size of attack database increases.

Effect of attack and defense index overlap. We fix the size of the defense and attack databases to 500M images, but we vary the percentage of images that is present in both databases. Figure 6(c) shows the accuracy of our models under PGD-FS attacks as a function of the intersection percentage. We observe that the accuracy of our defense degrades rapidly as the overlap between the attack and the defense database grows. However, our nearest-neighbor defense is effective when in the overlap between both databases is limited.

The experimental results show that nearest-neighbor defense strategies can be effective even when the adversary is aware of the defense strategy and adapts its attack accordingly, provided that the defender can perform some kind of “data obfuscation”. Such security-by-obfuscation is insecure in true white-box attack scenarios, but in real-world attack settings, it may be practical because it is difficult for the adversary to obtain the same set of hundreds of millions of images that the defender has access to.

7. Discussion and Future Work

In this study, we have explored the feasibility of using defense strategies web-scale nearest-neighbor search to provide image-classification systems robustness against adversarial attacks. Our experiments show that, whilst such defenses are not effective in pure white-box attack scenarios, they do perform competitively in realistic gray-box settings in which the adversary knows the defense strategy used but cannot access the exact web-scale image database.

Qualitative analyses of the nearest neighbors of adversarial images show that the model predictions for these neigh-

bors images are generally closely related to the true label of the adversarial image, even if the prediction is incorrect. On ImageNet, the prediction errors are often “fine-grained” [1]: for example, changing the label *golden retriever* into *labrador*. As a result, we believe that studies such as ours tend to overestimate the success rate of the adversary: if, for example, an attacker aims to perturb an objectionable image so that it can pass as benign, it often cannot get away with fine-grained changes to the model prediction. Randomly choosing a target class has been explored previously [1] and is a more realistic attack setting, since this attack setting is more likely to perturb predictions to a class that is semantically unrelated.

Our results provide the following avenues for future research. Specifically, they suggest that varying the depth of the network that constructs features for the nearest-neighbor search trades off adversarial robustness for accuracy on clean images: “early” features are more robust to attacks but work less well on clean images. This suggests future work should explore combinations of either features or nearest neighbors [24] at different depths. Future work should also investigate approaches that decrease the similarity between the feature-producing model and the image-classification model used in our approach, as well as train these networks using adversarial-training approaches such as adversarial logit pairing [16]. Other directions for future work include developing better strategies for selecting images to use in the image database, e.g., using supervision from hashtags or text queries associated with the images.

Acknowledgements. We thank Matthijs Douze, Jeff Johnson, Viswanath Sivakumar, Hervé Jegou, and Jake Zhao for many helpful discussions and code support, and Anish Athalye and Shibani Santurkar for their comments on an earlier draft of this paper.

References

- [1] A. Athalye, L. Engstrom, A. Ilyas, and K. Kwok. Synthesizing robust adversarial examples. *arXiv preprint arXiv:1707.07397*, 2017. **8**
- [2] N. Carlini and D. Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017. **1, 2, 3**
- [3] N. Carlini and D. Wagner. Audio adversarial examples: Targeted attacks on speech-to-text. *arXiv preprint arXiv:1801.01944*, 2018. **1**
- [4] G. K. Dziugaite, Z. Ghahramani, and D. M. Roy. A study of the effect of JPG compression on adversarial images. *arXiv preprint arXiv:1608.00853*, 2016. **2**
- [5] J. Ebrahimi, A. Rao, D. Lowd, and D. Dou. Hotflip: White-box adversarial examples for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 31–36, 2018. **1**
- [6] V. Fischer, M. C. Kumar, J. H. Metzen, and T. Brox. Adversarial examples for semantic image segmentation. *arXiv preprint arXiv:1703.01101*, 2017. **1**
- [7] S. Ghemawat and J. Dean. Leveldb. URL: <https://github.com/google/leveldb,%20http://leveldb.org>, 2011. **4**
- [8] J. Gilmer, L. Metz, F. F. Faghri, S. S. Schoenholz, R. Maithra, M. Wattenberg, and I. J. Goodfellow. Adversarial spheres. *arXiv preprint arXiv:1801.02774*, 2018. **2**
- [9] I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. **1, 2, 3, 11**
- [10] C. Guo, M. Rana, M. Cisse, and L. van der Maaten. Countering adversarial images using input transformations. *arXiv preprint arXiv:1711.00117*, 2017. **2, 3, 4, 6, 7, 11**
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. **4, 11**
- [12] R. Huang, B. Xu, D. Schuurmans, and C. Szepesvári. Learning with a strong adversary. *arXiv preprint arXiv:1511.03034*, 2015. **2**
- [13] A. Ilyas, L. Engstrom, A. Athalye, and J. Lin. Black-box adversarial attacks with limited queries and information. *arXiv preprint arXiv:1804.08598*, 2018. **3**
- [14] J. Johnson, M. Douze, and H. Jégou. Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734*, 2017. **4, 11**
- [15] E. Jones, T. Oliphant, P. Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. [Online; accessed ;today]. **11**
- [16] H. Kannan, A. Kurakin, and I. Goodfellow. Adversarial logit pairing. *arXiv preprint arXiv:1803.06373*, 2018. **2, 4, 7, 8**
- [17] A. Kurakin, I. Goodfellow, and S. Bengio. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*, 2016. **1, 2, 3, 11**
- [18] K. Lee, K. Lee, H. Lee, and J. Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in Neural Information Processing Systems*, 2018. **2**
- [19] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017. **1, 2, 3, 11**
- [20] D. Mahajan, R. Girshick, V. Ramanathan, K. He, M. Paluri, Y. Li, A. Bharambe, and L. van der Maaten. Exploring the limits of weakly supervised pretraining. *arXiv preprint arXiv:1805.00932*, 2018. **4, 5, 6**
- [21] S.-M. Moosavi-Dezfooli, A. Fawzi, and P. Frossard. Deep-fool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2574–2582, 2016. **1, 2**
- [22] N. Papernot and P. McDaniel. Extending defensive distillation. *arXiv preprint arXiv:1705.05264*, 2017. **2**
- [23] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, 2017. **3**
- [24] N. Papernot and P. D. McDaniel. Deep k-nearest neighbors: Towards confident, interpretable and robust deep learning. *arXiv preprint arXiv:1803.04765*, 2018. **1, 2, 8**
- [25] N. Papernot, P. D. M. McDaniel, X. Wu, S. Jha, and A. Swami. Distillation as a defense to adversarial perturbations against deep neural networks. *arXiv preprint arXiv:1511.04508*, 2015. **2**
- [26] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017. **11**
- [27] A. Raghunathan, J. Steinhardt, and P. Liang. Certified defenses against adversarial examples. *arXiv preprint arXiv:1801.09344*, 2018. **2**
- [28] P. Samangouei, M. Kabkab, and R. Chellappa. Defense-gan: Protecting classifiers against adversarial attacks using generative models. 2018. **2**
- [29] A. Sinha, H. Namkoong, and J. Duchi. Certifiable distributional robustness with principled adversarial training. *arXiv preprint arXiv:1710.10571*, 2017. **2**
- [30] Y. Song, T. Kim, S. Nowozin, S. Ermon, and N. Kushman. Pixeldefend: Leveraging generative models to understand and defend against adversarial examples. 2017. **2**
- [31] C. Sun, A. Shrivastava, S. Singh, and A. Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *Proc. ICCV*, 2017. **6**
- [32] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013. **1, 2**
- [33] B. Thomee, D. A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth, and L. Li. The new data and new challenges in multimedia research. *arXiv preprint arXiv:1503.01817*, 2015. **5**
- [34] G. Toliás, R. Sicre, and H. Jégou. Particular object retrieval with integral max-pooling of cnn activations. *arXiv preprint arXiv:1511.05879*, 2015. **5**
- [35] F. Tramèr, A. Kurakin, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel. Ensemble adversarial training: Attacks and defenses. *arXiv preprint arXiv:1705.07204*, 2017. **7**

- [36] Y. Wang, S. Jha, and K. Chaudhuri. Analyzing the robustness of nearest neighbors to adversarial examples. *arXiv preprint arXiv:1706.03922*, 2017. [5](#)
- [37] D. Warde-Farley and I. Goodfellow. Adversarial perturbations of deep neural networks. *Advanced Structured Prediction*, T. Hazan, G. Papandreou, and D. Tarlow, Eds, 2016. [3](#)
- [38] C. Xie, J. Wang, Z. Zhang, Y. Zhou, L. Xie, and A. L. Yuille. Adversarial examples for semantic segmentation and object detection. *CoRR*, abs/1703.08603, 2017. [1](#)
- [39] J. Zhao and K. Cho. Retrieval-augmented convolutional neural networks for improved robustness against adversarial examples. *arXiv preprint arXiv:1802.09502*, 2018. [1](#), [2](#), [7](#)