

SafetyNet: Detecting and Rejecting Adversarial Examples Robustly

Jiajun Lu, Theerasit Issaranon, David Forsyth
University of Illinois at Urbana Champaign
{jlu23, issaran1, daf}@illinois.edu

Abstract

We describe a method to produce a network where current methods such as DeepFool have great difficulty producing adversarial samples. Our construction suggests some insights into how deep networks work. We provide a reasonable analysis that our construction is difficult to defeat, and show experimentally that our method is hard to defeat using several standard networks and datasets. We use our method to produce a system that can reliably detect whether an image is a picture of a real scene or not. Our system applies to images captured with depth maps (RGBD images) and checks if a pair of image and depth map is consistent. It relies on the relative difficulty of producing naturalistic depth maps for images in post processing. We demonstrate that our system is robust to adversarial examples built from currently known attacking approaches.

1. Introduction

Adversarial examples are images with tiny, imperceptible perturbations that fool a classifier into predicting the wrong labels with high confidence. \mathbf{x} denotes the input to some classifier, which is a *natural* example and has label l . A variety of constructions [9, 14, 19, 24] can generate an *adversarial* example $\mathbf{a}(\mathbf{x})$ to make the classifier label it $m \neq l$. This is interesting, because $\|\mathbf{a}(\mathbf{x}) - \mathbf{x}\|_2$ is so small that we would expect $\mathbf{a}(\mathbf{x})$ to be labelled l .

Adversarial examples are a persistent problem of classification neural networks, and of many other classification schemes. Adversarial examples are easy to construct [29, 21, 3], and there are even universal adversarial perturbations [18]. Adversarial examples are important for practical reasons, because one can construct physical adversarial examples, suggesting that neural networks in current status are unusable in some image classification applications (*e.g.* imagine a small physical modification that could reliably get a stop sign classified as a go faster sign [24]). Adversarial examples are important for conceptual reasons too, because an explanation of why adversarial examples are easy to construct could cast some light on the inner life




Figure 1: SafetyNet consists of a conventional classifier (in our experiments, either VGG19 or ResNet) with an RBF-SVM that uses discrete codes computed from late stage ReLUs to detect adversarial examples. We show that (a) SafetyNet detects adversarial examples reliably, even if they are produced by methods not represented in the detectors’ training set and (b) it is very difficult to produce examples that are both misclassified and slip past SafetyNet’s detector.

of neural networks. The absence of theory means it is hard to defend against adversarial examples (for example, distillation was proposed as a defense [25], but was later shown to not work [2]).

Adversarial example constructions (*e.g.*, line search along the gradient [9]; LBFGS on an appropriate cost [29]; DeepFool [19]) all rely on the gradient of the network, but it is known that using the gradient of another similar network is sufficient [24], so concealing the gradient does not work as a defense for current networks. An important puzzle is that networks that generalize very well remain susceptible to adversarial examples [29]. Another important puzzle is that examples that are adversarial for one network tend to be adversarial for another as well [29, 15, 26]. Some network architectures appear to be robust to adversarial examples [13], which still need more empirical verification. At least some adversarial attacks appear to apply to many distinct networks [18].

We denote the probability distribution of examples by $P(X)$. At least in the case of vision, $P(X)$ has support on some complicated subset of the input space, which is known

as the “manifold” of “real images”. Nguyen *et al.* show how to construct examples that appear to be noise, but are confidently classified as objects [22]. This construction yields $\mathbf{a}(\mathbf{x})$ lies outside the support of $P(X)$, so the classifier’s labeling is unreliable because it has not seen such examples. However, most adversarial examples “look like” images to humans, such as figure 5 in [29], so they are likely to lie within the support of $P(X)$.

One way to build a network that is robust to adversarial examples is to train networks with enhanced training data (adding adversarial samples [17]); this approach faces difficulties, because the dimension of the images and features in networks means an unreasonable quantity of training data is required. Alternatively, we can build a network that detects and rejects an adversarial sample. Metzen *et al.* show that, by attaching a detection subnetwork that observes the state of the original classification network, one can tell whether it has been presented with an adversarial example or not [16]. However, because the gradients of their detection subnetwork are quite well behaved, the joint system can be attacked easily (in our experiments; [16] does not provide any analysis for attacking on both detectors and classifiers). We also show their detection subnetwork is also easily fooled by adversarial samples produced by attacking methods which are not used in detector training process.

Our method focuses on codes produced by quantizing individual ReLUs in particular layers of the classification network (“patterns of activation”), and proceed from the hypothesis:

Hypothesis 1 *Adversarial attacks work by producing different patterns of activation in late stage ReLUs to those produced by natural examples.*

These patterns lie outside the family for which the softmax layer would be reliable. This hypothesis suggests that: (a) the presence of an adversarial example can be detected (as in Metzen *et al.* [16]); (b) such detectors can be made very difficult to defeat (unlike Metzen *et al.* [16]; section 3); (c) such detectors should be good at generalization for different adversarial attacks (unlike Metzen *et al.* [16]); (d) transfer attacks work because an example that generates unfamiliar patterns in one network tends to generate unfamiliar patterns in other networks too; (e) transfer attacks could be defended as well (section 3).

Contributions: Section 2 describes our SafetyNet architecture, which consists of the original classifier network and a detector that rejects adversarial examples. A **type I attack** on SafetyNet consists of a standard adversarial example crafted to be (a) similar to a natural image; (b) misclassified by the original network. A **type II attack** consists of an example that is crafted to be (a) similar to a natural image; (b) misclassified; and (c) not rejected by SafetyNet.

We show that SafetyNet is robust to both types of attacks and generalize well. Concealing the gradients is highly effective for SafetyNet, and it produces a black box that is strongly resistant to the best attacks we have been able to construct. This is in sharp contrast to all other known methods [24, 16].

In section 3, we demonstrate SceneProof, a robust and reasonably effective proof that an image is an image of a real scene (a “real” image; contrast a “fake” image, which is not an image of a real scene). We identify images of real scenes by checking a match between the image and a depth map, which is hard to manipulate. We show that SceneProof is (a) accurate and (b) strongly resistant to attacks that try to get manipulated scenes identified as authentic scenes.

In section 4, we propose a model that explains why our approach works, and it also demonstrates that SafetyNet is difficult to attack in principle.

2. SafetyNet: Spotting Adversarial Examples

SafetyNet consists of the original classifier, and an adversary detector which looks at the internal state of the later layers in the original classifier, as in Figure 1. If the adversary detector declares that an example is adversarial, then the sample is rejected.

2.1. Detecting Adversarial Examples

The adversary detector needs to be hard to attack. We force an attacker to solve a hard discrete optimization problem. For a layer of ReLUs at a high level in the classification network, we quantize each ReLU at some set of thresholds to generate a discrete code (binarized code in the case of one threshold). Our hypothesis 1 suggests that different code patterns appear for natural examples and adversarial examples. We use an adversary detector that compares a code produced at test time with a collection of examples, meaning that an attacker must make the network produce a code that is acceptable to the detector (which is hard; section 3). The adversary detector in SafetyNet uses an RBF-SVM on binary or ternary codes (activation patterns) to find adversarial examples.

We denote a code by \mathbf{c} . The RBF-SVM classifies by

$$f(\mathbf{c}) = \sum_i^N \alpha_i y_i \exp(-\|\mathbf{c} - \mathbf{c}_i\|^2 / 2\sigma^2) + b \quad (1)$$

In this objective function, when σ is small, the detector produces essentially no gradient unless the attacking code \mathbf{c} is very close to a positive example \mathbf{c}_i . Our quantization process makes the detector more robust and the gradients even harder to get. Experiments show that this form of gradient obfuscation is quite robust, and that confusing the detector is very difficult without access to the RBF-SVM, and still

difficult even when access is possible. Experiments in section 3 and theory in section 4 confirms that the optimization problem is hard.

2.2. Attacking Methods

We use the following standard and strong attacks [2], with various choice of hyper-parameters, to test the robustness of the systems. Each attack searches for a nearby $\mathbf{a}(\mathbf{x})$ which changes the class of the example and does not create visual artifacts. We use these methods to produce both type I attack (fool the classifier) and type II attack (fool the classifier *and* sneak past the detector).

Fast Sign method: Goodfellow et al [9] described this simple method. The applied perturbation is the direction in image space which yields the highest increase of the linearized cost under l_∞ norm. It uses a hyper-parameter ϵ to govern the distance between adversarial and original image.

Iterative methods: Kurakin et al. [14] introduced an iteration version of the fast sign method, by applying it several times with a smaller step size α and clipping all pixels after each iteration to ensure that results stay in the ϵ neighborhood of the original image. We apply two versions of this method, one where the neighborhood is in L_∞ norm and another where it is in L_2 norm.

DeepFool method: Moosavi-Dezfooli et al. [19] introduced the DeepFool adversary, which is able to choose which class an example is switched to. DeepFool iteratively perturbs an image x_0^{adv} , linearizes the classifier around x_n^{adv} and finds the closest class boundary. The minimal step according to the l_p distance from x_n^{adv} to traverse this class boundary is determined and the resulting point is used as x_{n+1}^{adv} . The algorithm stops once x_{n+1}^{adv} changes the class of the actual classifier. We use a powerful L_2 version of DeepFool.

Transfer method: Papernot et al. [24] described a way to attack a black-box network. They generated adversarial samples using another accessible network, which performs the same task, and used these adversarial samples to attack the black-box network. This strategy has been notably reliable.

2.3. Type I Attacks Are Detected

Accuracy: Our SafetyNet can detect adversarial samples with high accuracy on CIFAR-10 [12] and ImageNet-1000 [4]. For classification networks, we used a 32-layer ResNet [10] for CIFAR-10 and a VGG19 network [28] for ImageNet-1000. Figure 2 shows the detection accuracy of our Binarized RBF-SVM detector on the x5 layer of ResNet for Cifar10 and on the fc7 layer of VGG19 trained for ImageNet-1000. Adversarial samples are generated by **Iterative-L2**, **Iterative-Linf**, **DeepFool-L2** and **FastSign** methods. Figure 2 compares our RBF-SVM detection results with the detector subnetwork results of [16]. The RoC

for our detector for Cifar-10 and ImageNet-1000 appears in Figure 3.

Our results show: When our detector is tested on the same adversary as it is trained on, its performance is similar to the detector subnetwork, even though our detector works on quantized activation patterns while the detector subnetwork works on original continuous activation patterns. DeepFool is a strong attack. Increasing the number of categories in the problem makes it easier for DeepFool to produce an undetected adversarial example, likely because it becomes easier to exploit local classification errors without producing strange ReLU activations. If DeepFool is required to produce a label outside the top-5 for the original example, the attack is much weaker.

Generalization across attacks: Generally, a detector cannot know at training time what attacks will occur at test time. We test generalization across attacks by training a detector on one class of attack, then testing with other classes of attack. Figure 2 shows that our RBF-SVM generalizes across attacks more reliably than a detector subnetwork. We believe this is because the representation presented to the RBF-SVM has been aggressively summarized (by quantization), so that the classifier is not distracted by subtle but irrelevant features. Another example that questions the generalization ability of neural networks is the JPEG example in Table 7. Networks trained on normal quality JPEG images fails significantly on low quality JPEG test images.

2.4. Rejecting by Classification Confidence

Our experiments demonstrate that there is a trade-off between classification confidence and detection easiness for adversarial examples. Adversarial examples with high confidence in wrong classification labels tend to have more abnormal activation patterns, so they are easier to be detected by detectors. While adversarial examples with low classification confidence in wrong labels are harder to be detected. For example, attacks like DeepFool add small and just enough perturbations to change the classification label, so these adversarial examples are sometimes hard to detect. However, these adversarial examples could not assign high classification confidence to the wrong label. If they perform more iterations and increase the wrong class classification confidence, our detector could detect them much easier.

Experiments also show that Type II attacks on our quantized SVM detector together with the classifier produce adversarial examples with low confidence. All these experiments mean that we can use classification confidence as a detection criteria, and it could help us increase the detector's detection ability and decrease the potential to be attacked by Type II attacks.

The classification confidence in our experiments is measured by the ratio of the example's second highest classification confidence to the highest classification confidence. For




Figure 2: SafetyNet accurately detects adversarial attacks. To facilitate comparison, we follow the conventions of [16], plotting the success of the adversary (i.e. its ability to fool the classifier; leftward is better) on the horizontal axis and the accuracy of the detector on the vertical axis (higher is better). We show results for binary (SVM) and ternary (M-SVM) codes, and for a variety of attacks (columns are tagged with the attack: L - iterative L_2 attack; ∞ - iterative L_∞ attack; D - DeepFool attack; F - fast sign attack; trailing digits indicate distinct hyperparameter settings). **A**: Results for the detection subnetwork on CIFAR-10 from [16]. **B**: Results for SafetyNet on CIFAR-10, where the detector was trained and tested on adversarial samples generated by the same attacking method (same setting as **A**). **C**: Results for SafetyNet *and* the detection subnetwork (cnn) of [16] on CIFAR-10, where the detector was trained on L_∞ attack and tested on other attacking methods; SafetyNet generalizes better than detection subnetwork to different adversarial attacking methods. **D**: Results for SafetyNet on ImageNet-1000, where the detector was trained and tested on the same adversarial method. The classifier is evaluated with top-5 accuracy (**E** is evaluated with top-1 accuracy, note difference in x axis); using top-5 accuracy significantly advantages the adversary detector, because forcing an adversarial example to move out of top-5 requires larger changes. **F**: Results for SafetyNet on ImageNet-1000 (top-5), where the detector was trained on L_∞ attack and tested on other attacking methods; SafetyNet has relatively small loss of detection accuracy (compared to **E**). We cannot compare to the detection subnetwork of [16], because they do not provide results for ImageNet-1000.




Figure 3: ROC curve for our adversary detector on various adversaries. **Left**: CIFAR-10; **center**: ImageNet-1000, top-1; **right**: ImageNet-1000, top-5. Deepfool-5 is a variant of deepfool that is required to force the adversarial example out of the original example’s top 5 classes. Deepfool is a strong adversarial attack, and seems to benefit from being able to choose attacking target class from multiple classes.

Method	Non Attack			Type I Attack			Type II Attack		
	F→T	T→F	T→T reject	F→T	T→F	T→T reject	F→T	T→F	T→T reject
Non Attack Data	9.7%	0%	9.4%	N/A	N/A	N/A	N/A	N/A	N/A
Unfamiliar Data Average	17.3%	0%	0%	N/A	N/A	N/A	N/A	N/A	N/A
Gradient Descent Attack	N/A	N/A	N/A	9.9%	5.0%	6.1%	16.3%	3.7%	6.2%
Transfer Attack Average	N/A	N/A	N/A	4.6%	9.4%	33.6%	7.9%	9.8%	26.6%

Table 1: Summary of our fc7 RBF-SVM detector’s reaction on various non attack data and Type I, Type II attacks (smaller is better). **F→T** means the rate at which false label images are classified as true and the detector does not spot, same for **T→F**. **T→T reject** means the rate at which true label images are classified as true, however, they are rejected by the detector. This number only matters for non attack data because attacks are likely to distort activation patterns even when the label keeps same. As expected, Type I attacks are less successful than Type II attacks. This is because a Type I attack does not explicitly try to fool the detector.




Figure 4: We show figures for successful Type I attacks (fool the classifier) on the original classifier network, and successful Type II attacks (fool both the classifier and detector) on our SafetyNet. Attackers are only allowed to manipulate the depth. Our SafetyNet is very difficult to attack and attacks changing label from False to True is harder. Successful attacks on our SafetyNet requires the original inputs hard to classify and the attacks also need to manipulate the images more.

example, if an image has 60% probability to be a dog and 15% probability to be a cat, our classification confidence is 0.25. We reject examples with classification confidence ratio bigger than a threshold, which means the classifier is unsure about the classification.

The classification confidence rejection results for non attack images and various Type II attack adversarial examples are included in Table 2 for Cifar-10 and Table 3 for ImageNet-1000. Both tables show that rejecting by classification confidence rejects few non attack images while

hugely increase the rejection of Type II attack adversarial examples. The benefits of rejecting by classification confidence is also demonstrated in the Type II attacks section.

2.5. Type II Attacks fail

A type II attack involves a search for an adversarial example that will be (a) mislabelled and (b) not detected. We perform the gradient descent based Type II attacks for Cifar-10 and ImageNet-1000 with SVM detector, and compare to detection subnetwork [16]. Metzen et al. [16] only investigated type I attacks and has not investigated type II attacks on the detection subnetwork. Because the gradients of detection subnetwork are better formed, it should be easier to attack with Type II gradient descent attacks.

In our experiments for Cifar-10 and ImageNet-1000, we use different gradient descent based Type II attacks (L0, L2, Fast, DeepFool and top-5 DeepFool) to attack the detector and classifier at the same time. In the later SceneProof part, gradient descent based Type II attacks on SceneProof dataset use L2 LBFGS method.

The summary for Type II attacks on Cifar-10 could be found in Table 4. The numbers reported in the table are the percentages of adversarial examples that are both misclassified and undetected (lower is better). Without classification confidence rejection, quantized SVM detector and detection subnetwork perform similar under Type II attacks for L0, L2 and Fast methods, and quantized SVM detector performs significantly better under DeepFool Type II attacks. With classification confidence rejection, quantized SVM detector is very hard to attack and performs better than detection subnetwork on almost all attacking methods. The classification confidence rejection increases at maximum 7% false rejection on non attack images. The detailed percentages of Type II attacks on Cifar-10 could be found in Table 10.

The summary for Type II attacks on ImageNet-1000 could be found in Table 5. The table arrangement is same to Table 4, and DeepFool5 is top-5 DeepFool attack. Quantized SVM detector consistently performs better than detection subnetwork for various attacking methods and for

	Statistics	Non Attack	L0 (II)	L2 (II)	Fast (II)	DeepFool (II)
m-SVM Det	Mean-confident	95.45%	73.95%	69.36%	74.73%	73.71%
	Mean-ratio	0.05	0.29	0.36	0.31	0.36
	Rejection-rate	7.22%	43.58%	53.96%	45.46%	63.22%
Subnet Det	Mean-confident	95.45%	95.71%	96.68%	79.21%	73.72%
	Mean-ratio	0.05	0.03	0.04	0.25	0.36
	Rejection-rate	7.22%	3.98%	5.50%	37.73%	63.22%

Table 2: CIFAR-10 classification confidence rejection results on non attack images, and various gradient descent based Type II attack adversarial examples. **Mean-confident** is the mean of classification confidence for the label with highest probability. **Mean-ratio** is the mean of the ratio of the second highest predicted label confidence to the highest predicted label confidence. **Rejection-rate** is the rate that examples are rejected because the ratio is higher than the threshold. The ratio for Cifar-10 is 0.25, which means the first predicted label confidence must be four times higher than the second one. For non attack data, the classification confidence rejection only rejects small amount of examples; for quantized SVM detector, it rejects majority of Type II attack adversarial examples; for detection subnetwork, the rejection is not as efficient as quantized SVM detector, because getting high classification confidence while fooling detection subnetwork is easier (compared to quantized SVM detector).

	Statistics	Non Attack	L0 (II)	L2 (II)	Fast (II)	DeepFool (II)	DeepFool5 (II)
m-SVM Det	Mean-confident	81.55%	76.80%	41.25%	40.64%	43.93%	37.83%
	Mean-ratio	0.15	0.17	0.43	0.49	0.77	0.51
	Rejection-rate	10.98%	14.26%	43.89%	49.55%	95.51%	51.90%
Subnet Det	Mean-confident	81.55%	67.53%	67.13%	36.65%	43.93%	37.82%
	Mean-ratio	0.15	0.28	0.30	0.51	0.77	0.51
	Rejection-rate	10.98%	25.21%	28.55%	51.80%	95.51%	51.84%

Table 3: ImageNet-1000 classification confidence rejection results on non attack images, and various gradient descent based Type II attack adversarial examples. The table arrangement is same to Table 2, and DeepFool5 is top-5 DeepFool. The rejection ratio threshold is 0.5. For non attack data, the classification confidence rejection only rejects small amount of examples; for quantized SVM detector and detection subnetwork, they reject majority of Type II attack adversarial examples.

both with classification confidence rejection and without. It’s very difficult to perform Type II attacks on quantized SVM detector with rejection. The classification confidence rejection increases at maximum 10% false rejection on non attack images. The detailed percentages of Type II attacks on ImageNet-1000 could be found in Table 11.

3. SceneProof

We would like Alice to be able to prove to Bob that her photo is real without the intervention of a team of experts, and we’d like Bob to have high confidence in the proof. This proof needs to operate at large scales (i.e. anyone could produce a proof while taking a picture), and automatically.

Current best methods to identify fake images require careful analysis of vanishing points [8], illumination angles [6], and shadows [11] (reviews in [8, 7]). Such analyses are difficult to conduct at large scales or automatically. We construct a proof by capturing an RGBD image. The proof of realness is achieved by a classifier that checks both image and depth and determines whether they are consistent. Such a system works if (a) the classifier is acceptably

accurate (i.e. it can determine whether the pair is real or not accurately); (b) it can detect a variety of adversarial manipulations of depth or image or both (i.e. type I attacks fail); and (c) type II attacks generally fail. We achieve this by using the SafetyNet architecture.

We are mainly concerned with attacks that make “fake” images to be labeled “real”. Natural attacks on our system are: produce a depth map for an RGB image using some regression method to obtain an RGBD image (regression); manipulate RGBD image by inserting new objects; take an RGBD image labeled “fake” and manipulate it to be labeled “real” (type I adversarial); take an RGBD image labeled “fake” and manipulate it to be labeled “real” in a way that fools SafetyNet’s adversary detector (type II adversarial). There is a wide range of available regression/adversarial attacks, and our system needs to be robust to various methods which might be used to prepare the regression/adversarial attack.

Real test data is easily obtained. We use the raw Kinect captures of LivingRoom and Bedroom from NYU v2 dataset [20]. However, fake data requires care. To evalu-

Method	L0 (II)	L2 (II)	Fast (II)	DeepFool (II)
m-SVM Det	19.73	18.70	6.86	22.01
m-SVM Det - R	9.86	7.32	3.41	8.32
Subnet Det	20.73	12.30	1.89	96.24
Subnet Det - R	19.69	11.57	1.19	35.39

Table 4: Percentages of CIFAR-10 Type II attack adversarial examples that are both misclassified and undetected, lower is better. - R means classification confidence rejection is used (rejection ratio is 0.25), otherwise only the detector is on duty. Without classification confidence rejection, quantized SVM detector and detection subnetwork perform similar under Type II attacks for L0, L2 and Fast methods, and quantized SVM detector performs significantly better under DeepFool Type II attacks. With classification confidence rejection, quantized SVM detector is hard to attack and performs better than detection subnetwork on almost all attacking methods.

Method	L0 (II)	L2 (II)	Fast (II)	DeepFool (II)	DeepFool5 (II)
m-SVM Det	25.15	26.40	12.97	45.26	30.08
m-SVM Det - R	23.19	15.05	8.26	2.32	15.52
Subnet Det	70.52	36.43	21.25	100.00	42.24
Subnet Det - R	52.56	26.66	12.16	4.49	21.99

Table 5: Percentages of IMAGENET-1000 Type II attack adversarial examples that are misclassified and undetected, lower is better. - R means classification confidence rejection is used (rejection ratio is 0.5), otherwise only the detector is on duty. Quantized SVM detector consistently performs better than detection subnetwork for various attacking methods and for both with classification confidence rejection and without. It is difficult to perform Type II attacks on quantized SVM detector with rejection.

ate generalization over different attacks, we omit some “regression” methods from the training data and use them only in test. “Regression” methods used in both train and test are: random swaps of depth and image planes; single image predicted depth [5]; rectangle cropped region insertion and random shifted or scaled misaligned depth and image. “Regression” methods used *only* in test are: all zero depth values; nearest neighbor down-sample and up-sampled images and depths; low quality JPEG compressed images and depths; Middlebury stereo RGBD dataset [27] and Sintel RGBD dataset [1](which should be classified “fake” because they are renderings). Refer to Figure 4 for dataset and attacks.

Type I attacks on SafetyNet fail: Type I attacks on SceneProof using a familiar adversary (i.e. one used to train the detector) fail. We report results for two detectors A (applied to fc7 of VGG19) and B (applied to fc6 of VGG19) in Table 6. Type I attacks on SceneProof using an unfamiliar adversary (i.e. one not used to train the detector) generally fail. We report results for two detectors A (applied to fc7 of VGG19) and B (applied to fc6 of VGG19) in Table 7.

A type II attack must both fool the classifier *and* sneak past the detector. We distinguish between two conditions. In non-blackbox case, the internals of the SafetyNet system is *accessible* to the attacker. Alternatively, the network may be a black box, with internal states and gradients concealed. In this case, attackers must probe with inputs and gather

outputs, or build another approximate network as in [24].

Type II attacks on accessible SafetyNet fail: a type II attack involves a search for an adversarial example that will be (a) mislabelled and (b) not detected. This search is made difficult by the quantization procedure and by the narrow basis functions in the RBF-SVM, so we smooth the quantization operation and the RBF-SVM kernel operation. Smoothing is essential to make the search tractable, but can significantly misapproximate SafetyNet (which is what makes attacks hard). Our smoothing attack uses a sigmoid function with parameter λ to simulate the quantization process. We also help the search process by increasing the size of the RBF parameter σ to form smoother gradients. Even after smoothing the objective function, attacks tend to fail, likely because it is hard to make an effective tradeoff between easy search and approximation. Table 8 includes Type I and Type II, blackbox and non-blackbox attacking results on SceneProof dataset. Our SafetyNet is the most robust architecture to various attacks.

Type II attacks on black box SafetyNet fail: Assume the state of SafetyNet is concealed. We follow [23, 18] by building attacks on various alternative networks, then transferring these network’s adversarial samples. These attacks fail for our SafetyNet, refer to Table 8. In contrast to SafetyNet, the detector subnetwork of [16] is generally susceptible to type II attacks in both blackbox and non-blackbox settings. This is because of quantization process and detec-

Test example type	Classifier Acc	B	A	B \wedge A	A \vee B	A \vee B, T
Natural RGBD, False	91.8%	15.2%	17.1%	14.3%	18.8%	19.6%
Natural RGBD, True	97.7%	10.1%	11.6%	9.2%	12.7%	10.8%
Adversarial RGBD, False	33.1%	89.1%	88.6%	87.3%	90.4%	88.9%
Adversarial RGBD, True	15.3%	81.3%	81.0%	79.1%	83.3%	83.7%

Table 6: This table shows detailed information about behavior of classifiers and detectors, summarized in Table 1. In this table, results are calculated on the whole SceneProof dataset testing split. Type I attacks on SceneProof using a familiar adversary (i.e. one used to train the detector) fail. We report results for two detectors: A is applied to fc7 of VGG19; B is applied to fc6 of VGG19. Each detector column gives the rate at which the detector labels examples as adversarial. A \vee B means the example is reported adversarial if either detector says it is adversarial. A \wedge B means the example is reported adversarial if both detectors say it is adversarial. A \vee B, T gives the fraction of examples that were detected as adversarial among these that were labelled True by the classifier. We break out results by type and ground truth label. For example, for natural RGBD images that do not represent real scenes (first row), the classifier labels 91.8% correctly; of the 8.2% incorrectly labelled true, the A \vee B adversary detector rejects 19.6% as adversarial examples (last column).

Test example type	Classifier Acc	B	A	B \wedge A	A \vee B	A \vee B, T
zero D channel	76.5%	6.5%	25.6%	6.1%	26.0%	82.0%
down-up sampled	75.2%	54.9%	60.6%	51.3%	63.4%	87.6%
low quality JPEG	36.4%	80.1%	79.2%	77.2%	82.2%	81.8%
Sintel RGBD [1]	27.6%	45.3%	51.7%	39.7%	57.2%	61.4%
Middlebury RGBD [27]	24.0%	39.7%	40.3%	33.4%	46.6%	47.8%

Table 7: This table shows detailed information about behavior of classifiers and detectors, summarized in Table 1. The table arrangement is same to Table 6. Type I attacks on SceneProof using an unfamiliar adversary (i.e. one not used to train the detector) generally fail. All these examples should be labelled false, OR rejected as adversarial. The column for each detector reports the rate at which the detector identifies examples as adversarial. For example, in the first row, 76.5% of zero D channel RGBD images are correctly labelled as false by the classifier; of those labelled “true”, 82.0% are rejected as adversarial (last column). This means that a total of 4.2% of zero D channel RGBD images pass through SafetyNet with “true” labels.

tion subnetwork’s classification boundary problem [18].

4. Theory: Bars and P-domains

We construct one possible explanation for adversarial examples that successfully explains (a) the phenomenology and (b) why SafetyNet works. We have a network with N layers of ReLU’s, and study $y_i^{(k)}(\mathbf{x})$, the values at the output of the k ’th layer of ReLUs. This is a piecewise linear function of \mathbf{x} . Such functions break up the input space into *cells*, at whose boundaries the piecewise linear function changes (i.e. is only C^0). Now assume that for some $y_i^{(k)}(\mathbf{x})$ there exist *p-domains* (union of cells) \mathcal{D} in the input space such that: (a) there are no or few examples in the p-domain; (b) the measure of \mathcal{D} under $P(X)$ is small; (c) $|y_i^{(k)}(\mathbf{x})|$ is large inside \mathcal{D} and small outside \mathcal{D} . We will always use the term “p-domain” to refer to domains with these properties. We think that the total measure of all p-domains under $P(X)$ is small.

By construction, ReLU networks can represent such p-domains. We construct a p-domain using a basis function with small support. $R(\mathbf{u})$ denote a ReLU applied to \mathbf{u} . We

have *basic bar function* ϕ .

$$\phi(\mathbf{x}; i, s, \epsilon) = \frac{1}{\epsilon} \begin{pmatrix} R((x_i - s) + \epsilon) - \\ 2R((x_i - s)) + \\ R((x_i - s) - \epsilon) \end{pmatrix}$$

where ϕ has support when $|x_i - s| < \epsilon$ and has peak value 1. For an index set \mathcal{I} with cardinality $\#\mathcal{I}$ and vectors \mathbf{s}, ϵ , we write *bar function* b as

$$b(\mathbf{x}; \mathcal{I}, \mathbf{s}, \epsilon) = R\left(\sum_{i \in \mathcal{I}} \phi(\mathbf{x}; i, s_i, \epsilon_i) - \#\mathcal{I} + 1\right)$$

where b has support when $\|\mathbf{x}_{\mathcal{I}} - \mathbf{s}_{\mathcal{I}}\|_1 < \epsilon_{\mathcal{I}}$. Figure 5 illustrates these functions. It is clear that a CNN can encode bars and weighted sums of bars, and that for at least $k \geq 2$ every $y_i^{(k)}$ could in principle be a bar function. Appropriate choices of \mathbf{s}, ϵ and \mathcal{I} choose the location and support of the bar and so can produce bars which have low measure under $P(X)$. Now the functions presented to the softmax layer are a linear combination of the $y_i^{(N)}(\mathbf{x})$. This means that with choice of weight and parameters, a bar can appear at this level, and create a p-domain.

Method	Ori		Subnet Det			Det A			Det ABC		
	F→T	T→F	F→T	T→F	T→T reject	F→T	T→F	T→T reject	F→T	T→F	T→T reject
Non Attack Data	16.3%	0.6%	8.4%	0%	10.2%	9.7%	0%	9.4%	8.4%	0%	9.9%
Gradient Descent (I)	32.8%	55.3%	13.4%	9.5%	6.0%	9.9%	5.0%	6.1%	8.4%	0.3%	6.3%
VGG FastSign TF (I)	30.6%	2.8%	14.9%	2.2%	54.1%	7.5%	2.5%	44.1%	6.6%	1.9%	47.2%
ResNet GradDesc TF (I)	28.9%	36.7%	15.3%	22.4%	33.2%	3.6%	13.4%	29.1%	2.7%	11.9%	30.3%
ResNet FastSign TF (I)	22.2%	29.1%	7.6%	15.1%	29.8%	2.8%	12.2%	27.5%	2.2%	11.6%	27.8%
Type I Average	28.6%	30.9%	12.8%	12.3%	30.8%	6.0%	8.3%	26.7%	5.0%	6.4%	27.9%
Gradient Descent (II)	32.8%	55.3%	26.3%	21.9%	11.9%	16.3%	3.7%	6.2%	13.2%	2.6%	9.6%
VGG Finetune TF (II)	20%	3.1%	17.1%	0%	43.5%	17.2%	0%	45.6%	17.2%	0%	48.4%
VGG Subnet Det TF (II)	16.3%	0.6%	13.7%	0%	15.6%	10.3%	0%	12.5%	9.1%	0%	13.1%
ResNet Finetune TF (II)	15.6%	40.3%	8.5%	31.3%	29.3%	1.3%	27.2%	20.6%	0.3%	25%	21.0%
ResNet Subnet Det TF (II)	23.8%	29.7%	17.6%	19.3%	29.8%	2.8%	12.2%	27.5%	2.2%	11.6%	27.5%
Type II Average	21.7%	25.8%	16.6%	14.5%	26.0%	9.6%	8.6%	22.5%	8.4%	7.84%	23.9%

Table 8: Type I and Type II attacks, non-blackbox and blackbox attacks on SceneProof all fail. This table is gather by attacking a randomly selected subset of 3200 images from the whole SceneProof dataset test split (contains 80K images). The table compares a VGG19 network (Ori) with the detection subnetwork of [16] (Subnet), and two variants of SafetyNet (Det A, where we have an RBF-SVM on fc7; and Det ABC, where we have an RBF-SVM on each of fc7, fc6 and pool5, and declare an adversary when any detector responds). T→F shows the rate at which true label classified as false and not detected and F→T shows false label classified as true and not detected (i.e. lower is better). T→T reject shows the rate at which true samples are classified as true, but rejected by detector. *This rate only matters for non attack data, and does not matter for all attacks because attacks are likely to distort the activation patterns even if the classification label has not been changed.* There is no manipulation for the non attack data, which represents unforced errors by the classifier; note that each of the adversary detectors catches a high percentage of the false positives committed by the classifier and rejects them as adversarial. We group attacks by type I attack (I) and type II attack (II). The gradient descent shows the performance of an attack by gradient descent method (type I or type II) on an accessible network. Even when the network is accessible, attacks tend to be unsuccessful. TF represents blackbox transfer attacks where adversarial samples are obtained from another network (VGG - a VGG19 model; ResNet - a ResNet model). The VGG19 (ResNet) FastSign TF gives results for a type I attack by transferring FastSign adversarials from a VGG19 (ResNet) model. VGG (ResNet) Finetune TF finetunes a VGG19 (ResNet) network with adversarial examples labelled false, and generate adversarials; VGG (ResNet) Subnet Det TF uses a VGG19 (ResNet) network with the detection subnetwork. The results show that original classifier network is easy to be attacked by all attacking methods. Subnet can detect type I attacks, but is not robust to transfer attacks and is vulnerable to type II attacks. Our SafetyNet is robust to Type I and Type II attacks, as well as gradient descent and transfer attacks. This is because of quantization process and our SafetyNet works like an example matcher, while detection subnetwork suffers from classification boundary problem [18].




Figure 5: Simple example bar functions on the x, y plane, where black is 0 and white is 1. **Left:** $\phi(\mathbf{x}; 1, 0, 1)$ (i.e. a bar in x , independent of y); **center:** $\phi(\mathbf{x}; 2, 0, 1)$; and **right:** $b(\mathbf{x}; \{1, 2\}, 0, 1)$.

We expect such p-domains to have several important properties. **Adversarial fertility:** P-domains can be used to make adversarial examples by choosing a point in a p-domain close to \mathbf{x} . Because there are no or few examples in the p-domain, the loss may not cause the classifier to control the maximum value attained by $y_i^{(k)}(\mathbf{x})$ in

this p-domain; and the large range of values inside the p-domain can be used to change the values in layers up-stream of k , by moving the example around the p-domain. **Generalization-neutral:** The requirement that p-domains have small measure in $P(X)$ means that both train and test examples are highly unlikely to lie in p-domains. A system with p-domains could generalize well without being immune to adversarial examples. Some subset of p-domains are likely **findable by LBFGS**. Consider the gradient of $y_i^{(N)}(\mathbf{x})$ with respect to \mathbf{x} in two cells separated by a boundary, where some ReLU changes state, weight decay encourages a relatively small change in gradient over these boundaries. If cells neighboring a p-domain have no or few examples in them, we can expect that the gradient change within cell is small too and a second order approximation of $y_i^{(N)}(\mathbf{x})$ could be reliable. We also expect cells to be small, so search and entering a p-domain are possible and requires crossing multiple cell boundaries, which

means many changes in ReLU activation. This argument suggests p-domains present **odd patterns of ReLU activation**, particularly in p-domains where some of the $y_i^{(k)}(\mathbf{x})$ are large in the absence of examples.

Why p-domains could exist: As Zhang *et al.* point out, the number of training examples available to a typical modern network is small compared to the relative capacity of deep networks [30]. For example, excellent training error is obtainable for randomly chosen image labels [30]. We expect that $y_i^{(N)}(\mathbf{x})$ will have a number of cells that is exponential in the dimension of \mathbf{x} , ensuring that the vast majority of cells lack any example. However, the weight decay term is not sufficient to ensure that $y_i^{(N)}$ is zero in these cells. Overshoot by stochastic gradient descent, caused by poor scaling in the loss, is the likely reason that $y_i^{(N)}(\mathbf{x})$ has support in these cells. Szegedy *et al.* demonstrate that, in practice, ReLU layers can have large norm as linear operators, despite weight decay (see [29], sec. 4.3), so large values in p-domains are plausible. This large norm is likely to be the result of overshoot. Recall that the value of $y_i^{(N)}(\mathbf{x})$ is determined by the *product* of numerous weights, so in some locations in \mathbf{x} , the value of $y_i^{(N)}$ could be large, which is a result of multiple layer norms interacting poorly.

An alternative to attacking by search using smoothed RBF gradients is as follows. One might pass an example through the main classifier, determine what code it had, then seek an adversarial example that produces that code (and so must fool the RBF-SVM). We sketch a proof that the optimization problem is extremely difficult. Choose some threshold $t > 0$. We use $b_t(u)$ for the function that binarizes its argument with t . Assume we have at least one unit $y_i^{(k)}$ that encodes a weighted sum of bar functions. We wish to create an adversarial example $\mathbf{a}(\mathbf{x}^*)$ that (a) meets criteria for being adversarial and (b) ensures that $b_t(y_i^{(k)}(\mathbf{a}))$ takes a prescribed value (either one or zero). The feasible set for this constraint can be disconnected (*e.g.* a sum of the bump functions of Figure 5 (right)), and so need not be convex, implying that the optimization problem is intractable. As a simple example, the following constraint set is disconnected for $\epsilon < 1/2$

$$\{\mathbf{x} \mid b_t(b(\mathbf{x}; 1, \mathbf{0}, \epsilon) + b(\mathbf{x}; 1, \mathbf{1}, \epsilon)) = 1\}.$$

5. Discussion

We have described a method to produce a classifier that identifies and rejects adversarial examples. Our SafetyNet is able to reject adversarial examples that come from attacking methods not seen in training data. We have shown that it is hard to produce an example that (a) is mislabeled and (b) is not detected as adversarial by SafetyNet. We have sketched one possible reason that SafetyNet works, and is hard to attack. Many interesting problems are opened by

our work, and we provides lots of insights into the mechanism that neural network works.

SaferNet: There might be some better architecture than our SafetyNet, whose objective function is harder to optimize. The ideal case would be an architecture that forces the attacker to solve a hard discrete optimization problem which does not naturally admit smoothing.

Neural network pruning: Our work suggests that networks behave poorly for input space regions where no data has been seen. We speculate that this behavior could be discouraged by a post-training pruning process, which removes neurons, paths or activation patterns not touched by training data.

Explicit management of overshoot during training: we have explained adversarial examples using p-domains, which is the result of poor damping of weights during training. We speculate that constructing adversarial examples during training, by identifying locations where this damping problem occurs and exploiting structural insights into network behavior, could control the adversarial sample problem (rather than just using adversarial examples as training data).

6. Supporting Materials

6.1. SceneProof Dataset

Our SceneProof dataset is processed from NYU Depth v2 raw captures, Sintel Synthetic RGBD dataset and Middlebury Stereo dataset. The dataset is split into part I and part II. Part I contains NYU natural image & depth pairs, along with manipulated unnatural scenes (swap depth, insert region, predicted depth, scale & shift depth), refer to Figure 6. It is used to train our classifier and work as test data part I. Part II contains unnatural scenes manipulated by other methods (set depth channel to zero, down sample and then up-sample both RGBD channels, aggressively compress the JPG RGBD images), and image & depth pairs from synthetic dataset and stereo dataset, refer to Figure 7. Part II is used as test data part II to test the generalization ability of our SceneProof network, and check the reactions of our detectors to unseen unnatural inputs. A good detector need to tend to reject unfamiliar data type, which does not exist in training data, because it is hard for classifier to do right classifications on unseen data types. In real application scenarios, it needs to be a human computer hybrid system where computer provides suspicious cases and human makes final decisions. Table 9 includes the dataset constitution, and we plan to release the dataset for academia usages.

6.2. Type II Attacks on Cifar-10 and ImageNet-1000

In this section, we include the detailed percentages of Type II attacks on Cifar-10 could be found in Table 10, and

the detailed percentages of Type II attacks on ImageNet-1000 could be found in Table 11.

	Training	Testing I	Testing II
Natural Scene	141780	57542	N/A
Swap Depth	33927	16094	N/A
Insert Region	30426	13741	N/A
Predicted Depth	53904	17026	N/A
Scale&Shift Depth	23523	10681	N/A
zeroD channel	N/A	N/A	1449
down-up sampled	N/A	N/A	1449
low quality JPG	N/A	N/A	1449
Sintel RGBD	N/A	N/A	54
Middlebury RGBD	N/A	N/A	30
Total	283560	115084	4431

Table 9: Number of image & depth pairs for each data type in each dataset split. Natural Scene has true label and the other data types have false labels.

References

- [1] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In A. Fitzgibbon et al. (Eds.), editor, *European Conf. on Computer Vision (ECCV)*, Part IV, LNCS 7577, pages 611–625. Springer-Verlag, Oct. 2012. 7, 8
- [2] N. Carlini and D. Wagner. Defensive distillation is not robust to adversarial examples. 1, 3
- [3] N. Carlini and D. Wagner. Towards evaluating the robustness of neural networks. *arXiv preprint arXiv:1608.04644*, 2016. 1
- [4] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. IEEE, 2009. 3
- [5] D. Eigen, C. Puhrsch, and R. Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in neural information processing systems*, pages 2366–2374, 2014. 7
- [6] H. Farid. Exposing photo manipulation with inconsistent reflections. *ACM Trans. Graph.*, 31(1):4, 2012. 6
- [7] H. Farid. Photo forensics. *MIT Press*, 2016. 6
- [8] H. Farid. How to detect faked photos. *American Scientist*, 2017. 6
- [9] I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 1, 3
- [10] 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. 3
- [11] E. Kee, J. F. O’Brien, and H. Farid. Exposing photo manipulation with inconsistent shadows. *ACM Transactions on Graphics (ToG)*, 32(3):28, 2013. 6
- [12] A. Krizhevsky. Learning multiple layers of features from tiny images. 2009. 3
- [13] D. Krotov and J. J. Hopfield. Dense associative memory is robust to adversarial inputs. *arXiv preprint arXiv:1701.00939*, 2017. 1
- [14] A. Kurakin, I. Goodfellow, and S. Bengio. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533*, 2016. 1, 3
- [15] Y. Liu, X. Chen, C. Liu, and D. Song. Delving into transferable adversarial examples and black-box attacks. *arXiv preprint arXiv:1611.02770*, 2016. 1
- [16] J. H. Metzen, T. Genewein, V. Fischer, and B. Bischoff. On detecting adversarial perturbations. *arXiv preprint arXiv:1702.04267*, 2017. 2, 3, 4, 5, 7, 9
- [17] T. Miyato, S.-i. Maeda, M. Koyama, K. Nakae, and S. Ishii. Distributional smoothing with virtual adversarial training. *arXiv preprint arXiv:1507.00677*, 2015. 2
- [18] S.-M. Moosavi-Dezfooli, A. Fawzi, O. Fawzi, and P. Frossard. Universal adversarial perturbations. *arXiv preprint arXiv:1610.08401*, 2016. 1, 7, 8, 9
- [19] S.-M. Moosavi-Dezfooli, A. Fawzi, and P. Frossard. Deepfool: 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, 3
- [20] P. K. Nathan Silberman, Derek Hoiem and R. Fergus. Indoor segmentation and support inference from rgbd images. In *ECCV*, 2012. 6
- [21] A. Nguyen, J. Yosinski, and J. Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 427–436, 2015. 1
- [22] A. Nguyen, J. Yosinski, and J. Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *CVPR*, 2015. 2
- [23] N. Papernot, P. McDaniel, and I. Goodfellow. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *arXiv preprint arXiv:1605.07277*, 2016. 7
- [24] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami. Practical black-box attacks against deep learning systems using adversarial examples. *arXiv preprint arXiv:1602.02697*, 2016. 1, 2, 3, 7
- [25] N. Papernot, P. McDaniel, X. Wu, S. Jha, and A. Swami. Distillation as a defense to adversarial perturbations against deep neural networks. In *Security and Privacy (SP), 2016 IEEE Symposium on*, pages 582–597. IEEE, 2016. 1
- [26] N. Papernot, P. D. McDaniel, and I. J. Goodfellow. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *arXiv preprint arXiv:1605.07277*, 2016. 1
- [27] D. Scharstein, H. Hirschmüller, Y. Kitajima, G. Krathwohl, N. Nešić, X. Wang, and P. Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In *German Conference on Pattern Recognition*, pages 31–42. Springer, 2014. 7, 8

Cifar-10		L0 (II)		L2 (II)		Fast (II)		DeepFool (II)	
		undet	det	undet	det	undet	det	undet	det
m-SVM Det	=	37.95	22.58	51.16	19.23	33.45	41.75	1.03	2.71
	≠	19.73	19.75	18.70	10.90	6.86	17.95	22.01	74.23
m-SVM Det - R	=	31.79	28.74	42.23	28.17	30.87	44.32	0.41	3.34
	≠	9.86	29.62	7.32	22.29	3.41	21.39	8.32	87.94
Subnet Det	=	16.91	21.64	28.57	32.01	8.06	66.57	3.76	0.00
	≠	20.73	40.72	12.30	27.13	1.89	23.48	96.24	0.00
Subnet Det - R	=	16.25	22.30	28.02	32.56	7.53	67.10	1.15	2.61
	≠	19.69	41.76	11.57	27.85	1.19	24.18	35.39	60.85

Table 10: Percentage details of Table 4 with correct classification (=) and undetected as adversarials (undet), correct classification and detected as adversarials (det), misclassification (\neq) and undetected as adversarials, misclassification and detected as adversarials. Table 4 comes from misclassification and undetected as adversarials (left down corner). *For all Type II attacks, correct classification and detected as adversarials percentage does not matter, because attacks tend to distort activation patterns even when the labels have not been changed.*

ImageNet-1000		L0 (II)		L2 (II)		Fast (II)		DeepFool (II)		DeepFool5 (II)	
		undet	det	undet	det	undet	det	undet	det	undet	det
m-SVM Det	=	0.00	0.00	3.12	1.58	55.21	7.04	0.00	0.00	0.00	0.00
	≠	25.15	74.84	26.40	68.90	12.97	24.78	45.26	54.74	30.08	69.92
m-SVM Det - R	=	0.00	0.00	2.43	2.27	53.06	9.19	0.00	0.00	0.00	0.00
	≠	23.19	76.80	15.05	80.24	8.26	29.48	2.32	97.67	15.52	84.48
m-SVM Det	=	17.67	4.13	33.13	20.69	21.96	13.28	0.00	0.00	0.00	0.00
	≠	70.52	7.68	36.43	9.74	21.25	43.52	100.00	0.00	42.24	57.76
m-SVM Det - R	=	16.03	5.77	30.86	22.97	19.63	15.61	0.00	0.00	0.00	0.00
	≠	52.56	25.64	26.66	19.51	12.16	52.61	4.49	95.51	21.99	78.01

Table 11: Percentage details of Table 5 with correct classification (=) and undetected as adversarials (undet), correct classification and detected as adversarials (det), misclassification (\neq) and undetected as adversarials, misclassification and detected as adversarials. Table 5 comes from misclassification and undetected as adversarials (left down corner). *For all Type II attacks, correct classification and detected as adversarials percentage does not matter, because attacks tend to distort activation patterns even when the labels have not been changed.*

- [28] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. 3
- [29] 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, 10
- [30] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. Understanding deep learning requires rethinking generalization. In *ICLR 2016*, 2016. 10




Figure 6: SceneProof dataset part I. Natural Scene has true label, and others have false labels.




Figure 7: SceneProof dataset part II. All have false labels.