



Trace and Detect Adversarial Attacks on CNNs Using Feature Response Maps

Mohammadreza Amirian^{1,2(✉)}, Friedhelm Schwenker², and Thilo Stadelmann¹

¹ ZHAW Datalab & School of Engineering, Winterthur, Switzerland
`{amir, stdm}@zhaw.ch`

² Institute of Neural Information Processing, Ulm University, Ulm, Germany
`friedhelm.schwenker@uni-ulm.de`

Abstract. The existence of adversarial attacks on convolutional neural networks (CNN) questions the fitness of such models for serious applications. The attacks manipulate an input image such that misclassification is evoked while still looking normal to a human observer—they are thus not easily detectable. In a different context, backpropagated activations of CNN hidden layers—“feature responses” to a given input—have been helpful to visualize for a human “debugger” what the CNN “looks at” while computing its output. In this work, we propose a novel detection method for adversarial examples to prevent attacks. We do so by tracking adversarial perturbations in feature responses, allowing for automatic detection using average local spatial entropy. The method does not alter the original network architecture and is fully human-interpretable. Experiments confirm the validity of our approach for state-of-the-art attacks on large-scale models trained on ImageNet.

Keywords: Model interpretability · Feature visualization · Diagnostic

1 Introduction

The success of deep neural nets for pattern recognition [35] has been a main driver behind the recent surge of interest in AI. A substantial part of this success is due to the Convolutional Neural Net (CNN) [5, 20] and its descendants, applied to image recognition tasks. Respective methods have reached the application level in business and industry [38] and lead to a wide variety of deployed models for critical applications like automated driving [2] or biometrics [46].

However, concerns regarding the reliability of deep neural networks have been raised through the discovery of so-called adversarial examples [41]. These inputs are specifically generated to “fool” [28] a classifier by visually appearing as some class (to humans), but being misclassified by the network with high confidence through the addition of barely visible perturbations (see Fig. 1). The perturbations are achieved by an optimization process on the input: the network weights are fixed, and the input pixels are optimized for the dual criterion of (a) classifying the input differently than the true class, and (b) minimizing the

changes to the input. A growing body of literature confirms the impact of this discovery on practical applications of neural nets [1]. It raises questions on how—and in what respect different from humans—they achieve their performance, and threatens serious deployments with the possibility of tailor-made adversarial attacks.

For instance, Su et al. [40] report on successfully attacking neural networks by modifying a single pixel. The attack works without having access to the internal structure nor the gradients in the network under attack. Moosavi-Dezfooli et al. [27] furthermore show the existence of universal adversarial perturbations that can be added to any image to fool a specific model, whereas transferability of perturbations from one model to another is for example shown by Xu et al. [44]. The impact of similar attacks extends beyond classification [26], is transferable to other modalities than images [6], and also works on models distinct from neural networks [31]. Finally, adversarial attacks have been shown to work reliably even after perturbed images have been printed and captured again via a mobile phone camera [18]. Apparently, such research touches a weak spot.

On the other hand, there is a recent interest in the interpretability of AI agents and in particular machine learning models [30, 42]. It goes hand in hand with societal developments like the new European legislation on data protection that is impacting any organization using algorithms on personal data [13]. While neural networks are publicly perceived as “black boxes” with respect to how they arrive at their conclusions [15], several methods have been developed recently to allow insight into the representation and decision surface of a trained model, improving interpretability. Prime candidates amongst these methods are feature visualization approaches that make the operations in hidden layers of a CNN visible [29, 37, 45]. They can thus serve a human engineer as a diagnostic tool in support of reasoning over success and failure of a model on the task at hand.

In this paper, we propose to use a specific form of CNN feature visualization, namely feature response maps, to not only *trace* the effect of adversarial inputs on algorithmic decisions throughout the CNN; we subsequently also use it as input to a novel automated *detection* approach, based on statistical analysis of the feature responses using average of image local spatial entropy. The goal is to decide if a model is currently under attack by the given input. Our approach has the advantage over existing methods of not changing the network architecture, i.e., not affecting classification accuracy; and of being interpretable both to humans and machines, an intriguing property also for future work on the method. Experiments on the validation set of ImageNet [34] with VGG19 networks [36] shows the validity of our approach for detecting various state-of-the-art attacks.

Below, Sect. 2 reviews related work in contrast to our approach. Section 3 presents the background on adversarial attacks and feature response estimation before Sect. 4 introduces our approach in detail. Section 5 reports on experimental evaluations, and Sect. 6 concludes with an outlook to future work.

2 Related Work

Work on adversarial examples for neural networks is a very active research field. Potential attacks and defenses are published at a high rate and have been surveyed recently by Akhtar and Mian [1]. Amongst potential defenses, directly comparable to our approach are those that focus on the sole detection of a possible attack and not on additionally recovering correct classification.

On one hand, several detection approaches exist that exploit specific abnormal behavioral traces that adversarial examples leave while passing through a neural network: Liang et al. [22] consider the artificial perturbations as noise in the *input* and attempt to detect it by quantizing and smoothing image filters. A similar concept underlies the SqueezeNet approach by Xu et al. [43], that compares the network’s *output* on the raw and filtered input, and raises a flag if detecting a large difference between both. Feinman et al. [9] observe the network’s output confidence as estimated by dropout in the forward pass [11], and Lu et al’s SafetyNet [23] looks for abnormal patterns in the ReLU activations of *higher layers*. In contrast, our method performs detection based on statistics of activation patterns in the complete *representation learning* part of the network as observed in feature response maps, whereas Li and Li [21] directly observe convolutional filter statistics there.

On the other hand, a second class of detection approaches trains sophisticated classifiers for directly sorting out malformed inputs: Meng and Chen’s MagNet [24] learns the manifold of friendly images, rejects far away ones as hostile and modifies close outliers to be attracted to the manifold before feeding them back to the network under attack. Grosse et al. [14] enhance the output of an attacked classifier by an additional class and retrain the model to directly classify adversarial examples as such. Metzen et al. [25] have a similar goal but target it via an additional subnetwork. In contrast, our method uses a simple threshold-based detector and pushes all decision power to the human-interpretable feature extraction via the feature response maps.

Finally, as shown in [1], different and mutually exclusive explanations for the existence of adversarial examples and the nature of neural network decision boundaries exist in the literature. Because our method enables a human investigator to trace attacks visually, it can be helpful in this debate in the future.

3 Background

We briefly present adversarial attacks and feature response estimation in general before assembling both parts into our detection approach in the next Section.

3.1 Adversarial Attacks

The main idea of adversarial attacks is to find a small perturbation for a given image that changes the decision of the Convolutional Neural Network. Pioneering work [41] demonstrated that negligible and visually insignificant perturbations

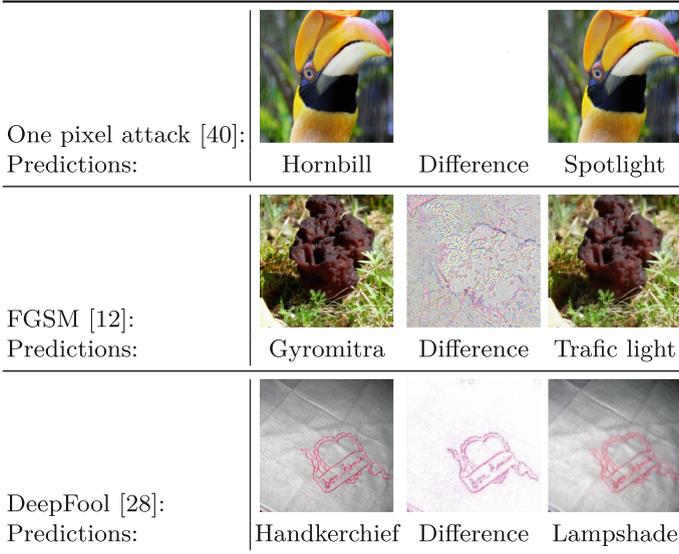


Fig. 1. Examples of different state-of-the-art adversarial attacks on a VGG19 model: original image and label (left), perturbation (middle) and mislabeled adversarial example (right). In the middle column difference of zero is encoded white and maximum difference is black because of visual enhancement.

could lead to considerable deviations in the networks' output. The problem of finding a perturbation $\boldsymbol{\eta}$ for a normalized clean image $\mathbf{I} \in \mathbb{R}^m$, where m is the image width \times height, is stated as follows [41]:

$$\min_{\boldsymbol{\eta}} \|\boldsymbol{\eta}\|_2 \quad \text{s.t.} \quad \mathcal{C}(\mathbf{I} + \boldsymbol{\eta}) \neq \ell; \quad \mathbf{I} + \boldsymbol{\eta} \in [0, 1]^m \quad (1)$$

where $\mathcal{C}(\cdot)$ presents the classifier and ℓ is the ground truth label. Szegedy et al. [41] proposed to solve the optimization problem in Eq. 1 for an arbitrary label ℓ' that differs from the ground truth to find the perturbation. However, box-constrained Limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) [10] is alternatively used to find perturbations satisfying Eq. 1 to improve computational efficiency. Optimization based on the L-BFGS algorithm for finding adversarial attacks are computational inefficient compared with gradient-based methods. Therefore, we use a couple of gradient-based attacks, a one-pixel attack, and boundary attack to compute adversarial examples (see Fig. 1).

Fast Gradient Sign Method (FGSM) [12] is a method suggested to compute adversarial perturbations based on the gradient $\nabla_{\mathbf{I}} J(\boldsymbol{\theta}, \mathbf{I}, \ell)$ of the cost function with respect to the original image pixel values:

$$\boldsymbol{\eta} = \epsilon \text{sign}(\nabla_{\mathbf{I}} J(\boldsymbol{\theta}, \mathbf{I}, \ell)) \quad (2)$$

where $\boldsymbol{\theta}$ represents the network parameters and ϵ is a constant factor that constrains the max-norm l_{∞} of the additive perturbation $\boldsymbol{\eta}$. The ground truth label

is presented by ℓ in Eq. 2. The sign function in Eq. 2 computes the elementwise sign of the gradient of the loss function with respect to the input image. Optimizing the perturbation in Eq. 2 in a single step is called Fast Gradient Sign Method (FGSM) in the literature. This method is a white box attack, i.e. the algorithm for finding the adversarial example requires the information of weights and gradients of the network.

Gradient attack is a simple and straightforward realization of finding adversarial perturbations in the FoolBox toolbox [33]. It optimizes pixel values of an original image to minimize the ground truth label confidence in a single step.

One pixel attack [40] is a semi-black box approach to compute adversarial examples using differential evolution [39]. The algorithm is not white box since it does not need the gradient information of the classifier; however, it is not fully black box as it needs the class probabilities. The iterative algorithm starts with randomly initialized parent perturbations. The generated offspring compete with their parent at each iteration, and the winners advance to the next step. The algorithm stops when the ground truth label probability is lower than 5%.

DeepFool [28] is a white box iterative approach in which the closest direction to the decision boundary is computed in every step. It is equivalent to finding the corresponding path to the orthogonal projection of the data point onto the affine hyperplane which separates the binary classes. The initial method for binary classifiers can be extended to a multi-class task by considering it as multiple one-versus-all binary classifications. After finding the optimal updates toward the decision boundary, the perturbation is added to the given image. The iterations continue with estimating the optimal perturbation and apply it to the perturbed image from the last step until the network decision changes.

Boundary attack is a reliable black-box attack proposed by Brendel et al. in [3]. The iterative algorithm already starts with an adversarial image and iteratively optimizes the distance between this image and the original image. It searches for an adversarial example with minimum distance from the original image.

3.2 Feature Response Estimation

The idea of visualizing CNNs through feature responses is to find out which region of the image leads to the final decision of the network. Computing feature responses enhances the interpretability of the classifier. In this paper, we use this visualization tool to track the effect of the adversarial attacks on a CNN's decision as well as to detect perturbed examples automatically.

Erhan et al. [8] used backpropagation for visualizing feature responses of CNNs. This is implemented by evaluating an arbitrary image in the forward pass, thereby retaining the values of activated neurons at the final convolutional layer, and backpropagating these activations to the original image. The feature response has higher intensities in the regions that cause larger values of activation in the network (see Fig. 2). The information of max-pooling layers in the forward pass can further improve the quality of visualizations. Zeiler et al. [45] proposed

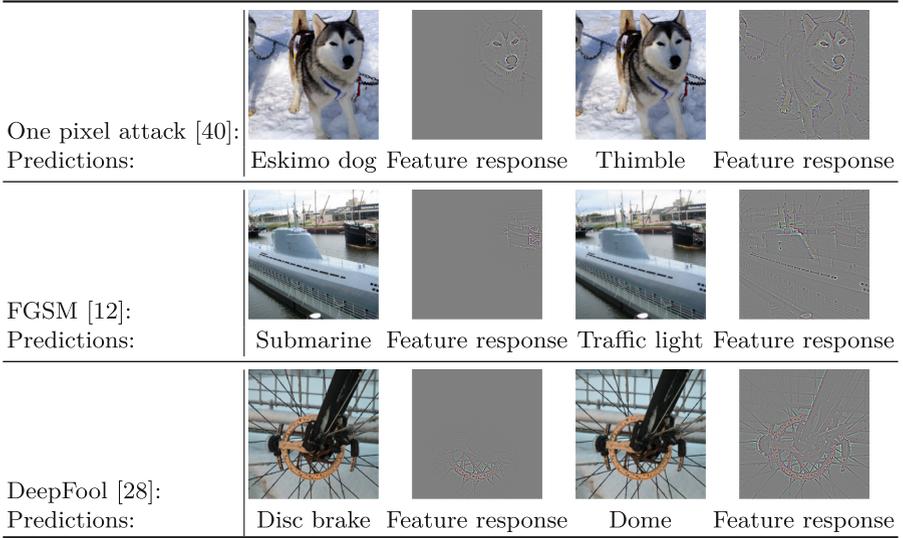


Fig. 2. Effect of adversarial attacks on feature responses: original image and feature response (left), perturbed versions (right).

to compute “switches”, the position of maxima in all pooling regions, and then construct the feature response using transposed convolutional [7] layers.

Ultimately, Springenberg et al. [37] proposed a combination of both methods called guided backpropagation. In this approach, the information of “switches” (max-pooling spatial information) is kept, and the activations are propagated backwards with the guidance of the “switch” information. This method leads to the best performance in network innards visualization, therefore we use guided backpropagation for computing feature response maps in this paper.

4 Human-Interpretable Detection of Adversarial Attacks

After reviewing the necessary background in the last Section, we will now present our work on tracing adversarial examples in feature response maps, which inspired a novel approach to automatic detection of adversarial perturbations in images. Using visual representations of the inner workings of neural network in this manner additionally provides a human expert guidance in developing deep convolutional networks with increased reliability and interpretability.

4.1 Tracing Adversarial Attacks in Feature Responses

The research question followed in this work is to obtain insight into the reasons behind misclassification of adversarial examples. Their effect in the feature response of a CNN is for example traced in Fig. 2. The general phenomenon

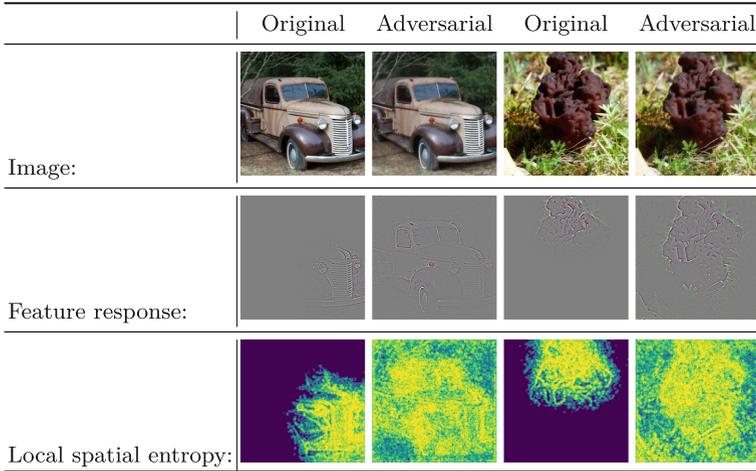


Fig. 3. Input, feature response and local spatial entropy for clean and perturbed images, respectively.

observed in all experiments is the broader feature response of adversarial examples. In contrast, Fig. 2 demonstrates that the network looks at a smaller region of the image—is more focused—in case of not manipulated samples.

The adversarial images are visually very similar to the original ones. However, they are not correctly recognizable by deep CNNs. The original idea which triggered this study is that the focus of CNNs changes during an adversarial attack and lead to the incorrect decision. Conversely, the network makes the correct decision once it focuses on the right region of the image. Visualizing the feature response provides this and other interesting information regarding the decision making in neural networks: for instance, the image of the submarine in Fig. 2 can be considered a good candidate for an adversarial attack since the CNN is making the decision based on an object in the background (see the feature response of the original submarine in Fig. 2).

4.2 Detecting Adversarial Attacks Using Spatial Entropy

Experiments for tracing the effect of adversarial attacks on feature responses thus suggested that a CNN classifier focuses on a broader region of the input if it has been maliciously perturbed. Figure 2 demonstrates this connection for decision making in case of clean inputs compared with manipulated ones. The effect of adversarial manipulation is visible in the local spatial entropy of the gray-scale feature responses as well (see Fig. 3). The feature responses are initially converted to gray scale images, and local spatial entropies are computed based on transformed feature responses as follows [4]:

$$S_k = - \sum_i \sum_j \mathbf{h}_k(i, j) \log_2(\mathbf{h}_k(i, j)) \quad (3)$$

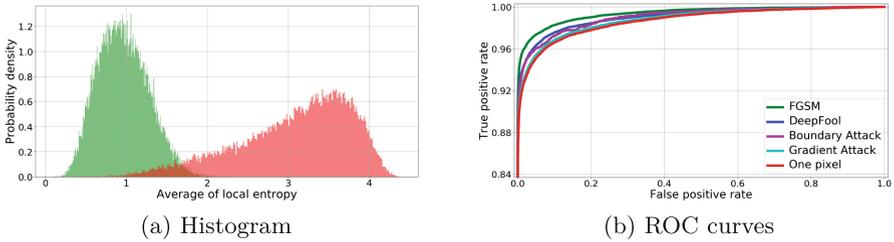


Fig. 4. (a) Distribution of average local spatial entropy in clean images (green) versus adversarial examples (red) as computed on the ImageNet validation set [34]. (b) Receiver operating characteristic (ROC) curve of the performance of our detection algorithm on different attacks. (Color figure online)

where S_k is the local spatial entropy of a small part (patch) of the input image and \mathbf{h}_k represents the normalized 2D histogram value of the k^{th} patch. The indices i and j scan through the height and width of the image patches. The patch size is 3×3 and the same as the filter size of the first layer of the used CNN (VGG19 [36]). The local spatial entropies of corresponding feature responses are presented in Fig. 3, and their difference for clean and adversarial examples suggests a likely chance to detect perturbed images based on this feature.

Accordingly, we propose to use the average local spatial entropy of an image as the final single measure to decide whether an attack has occurred or not. The average local spatial entropy \bar{S} is defined as:

$$\bar{S} = \frac{1}{K} \sum_k S_k \quad (4)$$

where K is the number of patches on the complete feature response and S_k shows the local spatial entropy as defined in Eq. 3 and depicted in the last row of Fig. 3. Our detector makes the final decision by comparing the average local spatial entropy from Eq. 4 with a selected threshold, i.e., we use this feature to measure the spatial complexity of an input image (feature response).

5 Experimental Results

To confirm the value of our final metric in Eq. 4, we first perform experiments to visually compare the approximated distribution of the averaged local spatial entropy of feature responses in clean and perturbed images. We use the validation set of ImageNet [34] with more than 50,000 images from 1,000 classes and again the VGG19 CNN [36]. Perturbations for this experiment are computed only via the Fast Gradient Sign Attack (FGSM) method for computational reasons. Figure 4(a) shows that the clean images are separable from perturbed examples although there is some overlap between the distributions.

Computing adversarial perturbations using evolutionary and iterative algorithms is demanding regarding time and computational resources. However, we

Table 1. Numerical evaluation of detection performance on the three different adversarial attacks. Column two gives the amount of tested attacks and elapsed approx. run time. Success of an adversarial attack is given if a perturbation changes the prediction. Columns four and five show average confidence values of the true (ground truth) and wrong (target) class after successful attack, respectively. The last columns show detection rates for different false positive rates.

Adversarial attack	#Images (run time [days])	Success rate	Ground truth confidence	Target class confidence	False positive rate		
					1%	5%	10%
FGSM [12]	50,014 (3)	0.925	0.022	0.588	0.954	0.974	0.983
Gradient attack [33]	50,014 (15)	0.499	0.052	0.371	0.922	0.954	0.969
One pixel attack [40]	50,014 (32)	0.620	0.037	0.463	0.917	0.951	0.966
DeepFool [28]	47,858 (42)	0.606	0.041	0.446	0.936	0.963	0.976
Boundary attack [3]	4,013 (17)	0.940	0.023	0.583	0.934	0.960	0.972

Table 2. Performance of similar adversarial attack detection methods. The Area Under Curve (AUC) is the average value of all attacks in the third and last row.

Method	Dataset	Network	Attack	Performance		
				Recall	Precision	AUC
Uncertainty density estimation [9]	SVHN [17]	LeNet [19]	FGSM	-	-	0.890
Adaptive noise reduction [22]	ImageNet (4 classes)	CaffeNet	DeepFool	0.956	0.911	-
Feature squeezing [43]	ImageNet-1000	VGG19	Several attacks	0.859	0.917	0.942
Statistical analysis [14]	MNIST	Self-designed	FGSM ($\epsilon = 0.3$)	0.999	0.940	-
Feature response (our approach)	ImageNet validation	VGG19	Several attacks	0.979	0.920	0.990

would like to apply the proposed detector to a wide range of adversarial attacks. Therefore, we have drawn a number of images from the validation set of ImageNet for each attack and present the detection performance of our method in Fig. 4. The selection of images is done sequentially by class and file name up to a total number of images per method that could be processed in a reasonable amount of time (see Table 1). We base our experiments on the FoolBox benchmarking implementation¹, running on a Pascal-based TitanX GPU.

¹ <https://github.com/bethgelab/foolbox>.



Fig. 5. Successful adversarial examples created by DeepFool [28] for binary and ternary classification tasks are only possible with notable visible perturbations.

Figure 4b presents the Receiver Operating Characteristics (ROC) of the proposed detector, and numerical evaluations are provided in Table 1. Our detection method performs better for gradient-based perturbations compared to the single pixel attack. Furthermore, Table 1 suggests that the best adversarial attack detection performance is achieved for FGSM and boundary attack perturbations, where the network confidences are changed the most. This observation suggests that the proposed detector is more sensitive to attacks which are stronger in fooling the network (i.e., change the ground truth label and target class confidence more drastically). By using feature responses, we detect more than 91% of the perturbed samples with a low false positive rate (1%).

In general, it is difficult to directly compare different studies on attack detectors since they use a vast variety of neural network models, datasets, attacks and experimental setups. We present a short overview of the performances of current detection approaches in Table 2. Our approach is most similar to the methods of Liang et al. [22] and Xu et al. [43]. The proposed detector in this paper outperforms both based on the presented results in their work; however, we cannot guarantee identical implementations and parameterizations of the used attacks (e.g., subset of used images, learning rates for optimization of perturbations). Similarly, adaptive noise reduction in the original publication [22] is applied to only four classes of the ImageNet dataset and defended a model based on CaffeNet, which differs from our experimental setup.

6 Discussion and Conclusion

The presented results demonstrate that the reality of adversarial attacks: improving the robustness of CNNs is necessary. However, we conducted further preliminary experiments on binary (cat versus dog [32]) and ternary (among three classes of cars [16]) classification tasks as proxies for the kind of few-class classifications settings frequently arising in practice. They suggest that it is more challenging to find adversarial examples in such a setting without plenty of “other classes” to pick from for misclassification. Figure 5 illustrates these results.

In this paper, we have presented an approach to detect adversarial attacks based on human-interpretable feature response maps. We traced the effect of

adversarial perturbations on the visual focus of the network in original images, which inspired a simple yet robust approach for automatic detection. This proposed method is based on thresholding the averaged local spatial entropy of the feature response maps and detects at least 91% of state-of-the-art adversarial attacks with a low false positive rate on the validation set of ImageNet. However, the results are not directly comparable with methods in the literature because of the diversity in the experimental setups and implementations of attacks.

Our results verify that feature response are informative to detect specific cases of failure in deep CNNs. The proposed detector applies to increase the interpretability of neural network decisions, which is an increasingly important topic towards robust and reliable AI. Future work, therefore, will concentrate on developing reliable and interpretable image classification methods for practical use cases based on our preliminary results for binary and ternary classification.

Acknowledgements. We are grateful for the support by Innosuisse grant 26025.1 PFES-ES “QualitAI”.

References

1. Akhtar, N., Mian, A.: Threat of adversarial attacks on deep learning in computer vision: A survey. arXiv preprint [arXiv:1801.00553](https://arxiv.org/abs/1801.00553) (2018)
2. Bojarski, M., et al.: End to end learning for self-driving cars. arXiv preprint [arXiv:1604.07316](https://arxiv.org/abs/1604.07316) (2016)
3. Brendel, W., Rauber, J., Bethge, M.: Decision-based adversarial attacks: reliable attacks against black-box machine learning models. arXiv preprint [arXiv:1712.04248](https://arxiv.org/abs/1712.04248) (2017)
4. Chanwimaluang, T., Fan, G.: An efficient blood vessel detection algorithm for retinal images using local entropy thresholding. In: International Symposium on Circuits and Systems (ISCAS), vol. 5 (2003)
5. Cireřan, D., Meier, U., Masci, J., Schmidhuber, J.: A committee of neural networks for traffic sign classification. In: IJCNN, pp. 1918–1921. IEEE (2011)
6. Cisse, M., Adi, Y., Neverova, N., Keshet, J.: Houdini: fooling deep structured prediction models. arXiv preprint [arXiv:1707.05373](https://arxiv.org/abs/1707.05373) (2017)
7. Dumoulin, V., Visin, F.: A guide to convolution arithmetic for deep learning. arXiv preprint [arXiv:1603.07285](https://arxiv.org/abs/1603.07285) (2016)
8. Erhan, D., Bengio, Y., Courville, A., Vincent, P.: Visualizing higher-layer features of a deep network. Univ. Montr. **1341**(3), 1 (2009)
9. Feinman, R., Curtin, R.R., Shintre, S., Gardner, A.B.: Detecting adversarial samples from artifacts. arXiv preprint [arXiv:1703.00410](https://arxiv.org/abs/1703.00410) (2017)
10. Fletcher, R.: Practical Methods of Optimization. Wiley, Hoboken (2013)
11. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: representing model uncertainty in deep learning. In: ICML (2016)
12. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: ICLR (2015)
13. Goodman, B., Flaxman, S.: EU regulations on algorithmic decision-making and a “right to explanation”. In: ICML Workshop on Human Interpretability in Machine Learning (WHI) (2016)

14. Grosse, K., Manoharan, P., Papernot, N., Backes, M., McDaniel, P.: On the (statistical) detection of adversarial examples. arXiv preprint [arXiv:1702.06280](https://arxiv.org/abs/1702.06280) (2017)
15. Gunning, D.: Explainable Artificial Intelligence (XAI). Defense Advanced Research Projects Agency (DARPA) (2017)
16. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3D object representations for fine-grained categorization. In: ICCV Workshops (2013)
17. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images (2009)
18. Kurakin, A., Goodfellow, I., Bengio, S.: Adversarial examples in the physical world. In: ICRL Workshop Track (2016)
19. LeCun, Y., et al.: Backpropagation applied to handwritten zip code recognition. *Neural Comput.* **1**(4), 541–551 (1989)
20. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proc. IEEE* **86**(11), 2278–2324 (1998)
21. Li, X., Li, F.: Adversarial examples detection in deep networks with convolutional filter statistics. arXiv preprint [arXiv:1612.07767](https://arxiv.org/abs/1612.07767) (2016)
22. Liang, B., Li, H., Su, M., Li, X., Shi, W., Wang, X.: Detecting adversarial examples in deep networks with adaptive noise reduction. arXiv preprint [arXiv:1705.08378](https://arxiv.org/abs/1705.08378) (2017)
23. Lu, J., Issararanon, T., Forsyth, D.: Safetynet: Detecting and rejecting adversarial examples robustly. arXiv preprint [arXiv:1704.00103](https://arxiv.org/abs/1704.00103) (2017)
24. Meng, D., Chen, H.: Magnet: a two-pronged defense against adversarial examples. In: ACM SIGSAC Conference on Computer and Communications Security (2017)
25. Metzen, J.H., Genewein, T., Fischer, V., Bischoff, B.: On detecting adversarial perturbations. In: ICLR (2017)
26. Metzen, J.H., Kumar, M.C., Brox, T., Fischer, V.: Universal adversarial perturbations against semantic image segmentation. arXiv preprint [arXiv:1704.05712](https://arxiv.org/abs/1704.05712) (2017)
27. Moosavi-Dezfooli, S.M., Fawzi, A., Fawzi, O., Frossard, P.: Universal adversarial perturbations. arXiv preprint [arXiv:1610.08401](https://arxiv.org/abs/1610.08401) (2017)
28. Moosavi-Dezfooli, S.M., Fawzi, A., Frossard, P.: Deepfool: a simple and accurate method to fool deep neural networks. In: CVPR (2016)
29. Olah, C., Mordvintsev, A., Schubert, L.: Feature visualization. *Distill* (2017). <https://doi.org/10.23915/distill.00007>
30. Olah, C., et al.: The Building Blocks of Interpretability. *Distill* (2018). <https://doi.org/10.23915/distill.00010>
31. Papernot, N., McDaniel, P., Goodfellow, I.: Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. arXiv preprint [arXiv:1605.07277](https://arxiv.org/abs/1605.07277) (2016)
32. Parkhi, O.M., Vedaldi, A., Zisserman, A., Jawahar, C.: Cats and dogs. In: CVPR (2012)
33. Rauber, J., Brendel, W., Bethge, M.: Foolbox v0.8.0: A python toolbox to benchmark the robustness of machine learning models. arXiv preprint [arXiv:1707.04131](https://arxiv.org/abs/1707.04131) (2017)
34. Russakovsky, O., et al.: Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis. (IJCV)* **115**(3), 211–252 (2015). <https://doi.org/10.1007/s11263-015-0816-y>
35. Schmidhuber, J.: Deep learning in neural networks: an overview. *Neural Netw.* **61**, 85–117 (2015)
36. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ICLR (2015)

37. Springenberg, J.T., Dosovitskiy, A., Brox, T., Riedmiller, M.: Striving for simplicity: The all convolutional net. arXiv preprint [arXiv:1412.6806](https://arxiv.org/abs/1412.6806) (2014)
38. Stadelmann, T., Tolkachev, V., Sick, B., Stampfli, J., Dürr, O.: Beyond imagenet - deep learning in industrial practice. In: Braschler, M., Stadelmann, T., Stockinger, K. (eds.) *Applied Data Science - Lessons Learned for the Data-Driven Business*. Springer (2018, to appear). <https://stdm.github.io/data-science-book/>
39. Storn, R., Price, K.: Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces. *J. Glob. Optim.* **11**(4), 341–359 (1997)
40. Su, J., Vargas, D.V., Kouichi, S.: One pixel attack for fooling deep neural networks. arXiv preprint [arXiv:1710.08864](https://arxiv.org/abs/1710.08864) (2017)
41. Szegedy, C., et al.: Intriguing properties of neural networks. In: ICLR (2014)
42. Vellido, A., Martín-Guerrero, J.D., Lisboa, P.J.: Making machine learning models interpretable. In: ESANN. vol. 12, pp. 163–172 (2012)
43. Xu, W., Evans, D., Qi, Y.: Feature squeezing: detecting adversarial examples in deep neural networks (2018)
44. Xu, X., Chen, X., Liu, C., Rohrbach, A., Darell, T., Song, D.: Can you fool AI with adversarial examples on a visual turing test? arXiv preprint [arXiv:1709.08693](https://arxiv.org/abs/1709.08693) (2017)
45. Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014*. LNCS, vol. 8689, pp. 818–833. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10590-1_53
46. Zhu, J., Liao, S., Yi, D., Lei, Z., Li, S.Z.: Multi-label CNN based pedestrian attribute learning for soft biometrics. In: *International Conference on Biometrics (ICB)*. IEEE (2015)