
EXPLAINING AWAY ATTACKS AGAINST NEURAL NETWORKS

Sean Saito¹ Jin Wang¹

ABSTRACT

We investigate the problem of identifying adversarial attacks on image-based neural networks. We present intriguing experimental results showing significant discrepancies between the explanations generated for the predictions of a model on clean and adversarial data. Utilizing this intuition, we propose a framework which can identify whether a given input is adversarial based on the explanations given by the model. Code for our experiments can be found here: <https://github.com/seansaito/Explaining-Away-Attacks-Against-Neural-Networks>.

1 INTRODUCTION

Adversarial attacks cause state-of-the-art neural networks to make misclassifications to inputs with imperceptible perturbations (Szegedy et al., 2013; Goodfellow et al., 2014). In this work, we propose a framework for detecting adversarial attacks by utilizing the explanations of the model’s predictions. Our experiments in Section 4 demonstrate that this framework can detect untargeted attacks at **99.81%** accuracy and targeted attacks at **99.87%** accuracy.

2 EXPLAINING PREDICTIONS ON ADVERSARIAL EXAMPLES

In an adversarial attack, the goal of an attacker is to construct an input \tilde{x} that is minimally perturbed from the original clean input x which causes some misclassification (an *untargeted attack*) or a desired prediction (a *targeted attack*) from a neural network. Several methods exist which achieve this, including the Basic Iterative Method (BIM) (Kurakin et al., 2016), which iteratively finds the gradient of the model’s loss function with respect to the input:

$$\begin{aligned}x_0 &= x \\ x_{i+1} &= x_i - \epsilon \cdot \text{sign}(\nabla_{x_i} J(x_i, y_t))\end{aligned}$$

where ϵ controls the magnitude of perturbations made to the image, resulting in perceptively identical images which cause large deviations in the behavior of the model. We hypothesize that the explanations of the model on each prediction would also vary significantly, which could potentially be a clue to identify whether a given prediction was caused by an adversarial attack.

¹SAP Asia, Singapore. Correspondence to: Sean Saito <sean.saito@sap.com>, Jin Wang <jin.wang02@sap.com>.

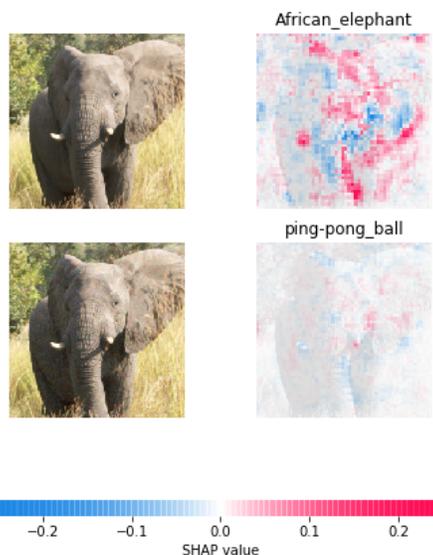


Figure 1. A comparison of SHAP Integrated Gradients computed for a pair of clean and adversarial images. The prediction for *African_elephant* had a confidence of **99.00%**, whereas the prediction for *ping-pong_ball* was **99.21%**.

To evaluate this hypothesis, we generate explanations via the SHAP Integrated Gradients framework (Lundberg & Lee, 2017), a framework which combines Integrated Gradients, a gradient-based pixel attribution method (Sundararajan et al., 2017), with SHAP, an additive feature attribution method that calculates Shapley values for each pixel:

$$IG_i(x) = \frac{x_i - b_i}{m} \sum_{h=1}^m \frac{\partial f}{\partial x_i} \left(b + \frac{h}{m} (x - b) \right)$$

In all experiments in this work, we generate explanations via the Inception-V3 model (Szegedy et al., 2016). Figure 1 displays an example of explanations generated for a pair of clean and adversarial images. Red and blue values indicate

the SHAP value of each pixel. The figure indicates qualitatively that the SHAP values for the clean image predictions are larger than those of the adversarial image.

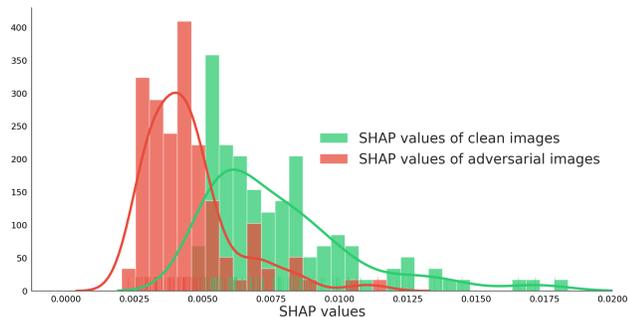


Figure 2. A comparison of SHAP Integrated Gradients distributions between clean images and adversarial images.

Figure 2 shows a distributional comparison of the absolute SHAP values across 1,000 test images from the ImageNet (Deng et al., 2009) dataset. We again observe a significant difference in the distribution between the SHAP values of clean and adversarial images. This suggests that neural networks are unable to provide sufficient “evidence” when explaining its decisions on adversarial examples. This intuition leads to our proposed framework of detecting adversarial attacks.

3 A FRAMEWORK FOR DETECTING ADVERSARIAL EXAMPLES

We propose a simple yet effective framework for detecting adversarial attacks. Based on the training images (S_{ori}), adversarial samples (S_{adv}), and neural network f , we treat the SHAP values coming from the predictions of f as feature descriptors, denoted as Z_{ori} and Z_{adv} respectively. We combine Z_{ori} and Z_{adv} to construct a binary classification dataset which trains an adversarial example detector $D(z)$. For any future sample s and prediction $f(s)$, we generate SHAP values z_s and call $D(z_s)$ to indicate whether the sample is adversarial.

4 EXPERIMENTS AND RESULTS

Our dataset contains 10,000 pairs of images from ImageNet across all classes, with 20% held out for testing. We employ three types of BIM attacks: untargeted, targeted (random class), and targeted (2^{nd} most confident class). We set ϵ as 0.1 for all attacks. Table 1 shows attack success rates.

We experiment with two off-the-shelf classifiers to assess how well they can detect adversarial examples based on the SHAP feature descriptors. Results shown in Table 2 indicate that simple classifiers can detect adversarial examples with high accuracy.

Table 1. Attack success rate of each attack

	Success Rate
Untargeted attack	99.13%
Targeted attack (random)	98.41%
Targeted attack (2^{nd} most confident)	99.87%

Table 2. Accuracy of detecting each attack

	Logistic Regression	Random Forest
Untargeted attack	99.16%	99.81%
Targeted attack (random)	98.96%	99.83%
Targeted attack (2^{nd} confident)	99.17%	99.87%

5 CONCLUSION AND FUTURE WORK

We have shown empirical results suggesting that explanations can help us detect adversarial attacks. Future work would focus on evaluating and refining this method under additional threat models.

REFERENCES

- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- Kurakin, A., Goodfellow, I., and Bengio, S. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*, 2016.
- Lundberg, S. M. and Lee, S.-I. A unified approach to interpreting model predictions. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 4765–4774. Curran Associates, Inc., 2017.
- Sundararajan, M., Taly, A., and Yan, Q. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 3319–3328. JMLR. org, 2017.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.