

ATTRIBUTION-DRIVEN CAUSAL ANALYSIS FOR DETECTION OF ADVERSARIAL EXAMPLES

Susmit Jha

Computer Science Laboratory
SRI International

Sunny Raj, Steven Lawrence Fernandes, Sumit Kumar Jha

Department of Computer Science
University of Central Florida, Orlando

Somesh Jha

Department of Computer Science
University of Wisconsin, Madison

Brian Jalaian, Gunjan Verma, Ananthram Swami

US Army Research Laboratory, Adelphi

ABSTRACT

Attribution methods have been developed to explain the decision of a machine learning model on a given input. We study the robustness of machine learning models on benign and adversarial inputs in the causal neighborhood of an input obtained by incrementally masking high attribution features. Our study indicates that benign inputs are robust to the masking of high attribution features but adversarial inputs generated by the state-of-the-art adversarial attack methods such as DeepFool, FGSM, CW and PGD, are not robust to such masking. Further, our study demonstrates that this concentration of high-attribution features responsible for the incorrect decision is even more pronounced in physically realizable adversarial examples. This difference in attribution of benign and adversarial inputs can be used to detect adversarial examples. This defense approach is independent of training data and attack method, and we demonstrate its effectiveness on digital and physically realizable perturbations.

1 INTRODUCTION

Deep learning models have been wildly successful in applications such as computer vision (Krizhevsky et al., 2017), natural language processing, speech recognition, and automatic control. These models have reached human-level performance on several benchmarks (He et al., 2015; Xiong et al., 2017). But the adoption of these models in safety-critical or high-security applications is inhibited due to two major concerns: their brittleness to adversarial attack methods that can make imperceptible modification to inputs and trigger wrong decisions (Szegedy et al., 2013; Papernot et al., 2016a), and the lack of interpretability (Gunning, 2017). Significant progress has also been made towards adversarial robustness (Papernot et al., 2016b; Madry et al., 2017; Engstrom et al., 2018) and explainability (Li & Yu, 2015; Yi et al., 2016), and a few recent theoretical studies (Kilbertus et al., 2018; Chalasani et al., 2018) indicate a strong connection between the two.

In this paper, we investigate this connection between the resilience to adversarial perturbations and the attribution-based explanation of individual decisions by a machine learning model. Our study is motivated by Kahneman’s decomposition (Kahneman & Egan, 2011) of cognition into two layers: System 1, or the intuitive system, and System 2, or the deliberate system. The original machine learning model represents the System 1 and the attribution analysis presented in this paper is the deliberative System 2 that can detect inconsistencies in System 1 by analyzing the attribution generated by the model on the input. Our central hypothesis is that adversarial inputs are imperceptibly similar to benign inputs and are still able to trigger wrong decisions because of a relatively small number of features with high attribution in the machine learning model. This concentration of causal attributions is central to their physical realizability and perceptual indistinguishability from benign inputs.

The example in Figure 1 demonstrates the shift in positive attributions and the corresponding qualitative change in saliency map due to the adversarial perturbation. We observe a similar shift in features with high magnitude attributions for the adversarial examples generated by the state-of-the-art purely digital attack methods such as DeepFool (Moosavi-Dezfooli et al., 2016), projected gradient descent (PGD) (Madry et al., 2017), fast gradient sign method (FGSM) (Goodfellow, 2014) and Carlini-Wagner (CW) (Carlini & Wagner, 2017).



Figure 1: The top row shows an original yawl image and the bottom row shows the adversarial input obtained by adding an adversarial patch (Brown et al., 2017) which causes misclassification. The leftmost image is the input, next to it is the saliency map that shows the features (pixels) with positive attributions (Szegedy et al., 2016), and the rightmost image shows the pixels with top 2% attribution masked. People in the boat are in the top 2% positive attribution features in original image but not in the adversarial image. The saliency map showing attributions are obtained using Integrated Gradient method (Sundararajan et al., 2017)

Adversarial perturbations often create a small fraction of high-attribution features responsible for the change in the output model without visible change in the input. Consequently, it is possible to identify an adversarial example by examining inputs in its causal neighborhood obtained by incrementally masking the features which have high magnitude attributions. While benign inputs are robust and the model does not change its decision on the input in this neighborhood, the decision of the model is not robust on the adversarial examples. The localized nature of existing physically realizable attacks naturally concentrates the high attribution features. The ease of detection of these adversarial examples suggests a trade off in physical realizability and indistinguishability in attribution space. We make the following novel contributions in this paper:

- We define a causal neighborhood of an input in the attribution space that can be used to measure the efficacy of an adversarial attack beyond perceptual similarity or the commonly used L_p norms. Effective attack methods should not just be able to fool the intuitive System 1 i.e. the machine learning model but also the deliberative System 2 i.e. the attribution analysis presented here.
- We use attribution methods to analyze the robustness of benign as well as adversarial inputs by incrementally masking high attribution features. Our empirical analysis includes digital attack methods such as DeepFool, FGSM, CW and PGD, and physically realizable perturbations.
- We propose a defense layer based on Kahneman’s decomposition of cognition into intuitive System 1 and deliberative System 2. This approach does not rely on analyzing training data such as manifold-based defense (Ilyas et al., 2017; Jha et al., 2018), or statistical signature of the machine learning models such as logit pairing (Engstrom et al., 2018), or methods that exploit the knowledge of specific attack for adversarial training (Tramèr et al., 2017), or robust optimization with L_p norm bounds (Madry et al., 2017; Raghunathan et al., 2018).

The rest of the paper is organized as follows. We discuss relevant background and review related work in Section 2. We present our analysis approach and derived defense against adversarial perturbation in Section 3. The experimental results from our empirical analysis on MNIST and ILSVRC datasets using different attack methods are presented in Section 4. We conclude in Section 5 by discussing the limitations of our approach, and the salient findings of our study.

2 BACKGROUND AND RELATED WORK

Szegedy et al. (2013) first introduced adversarial examples using an L-BFGS method. These generated adversarial images could be generalized to different training databases and models. L-BFGS attack uses an expensive linear search method to obtain the optimal value, but this method is time-consuming and impractical. Goodfellow (2014) proposed a fast gradient sign method (FGSM) to generate faster adversarial images as compared to the L-BFGS method; this method performed only a one-step gradient update at each pixel along the direction of the gradient’s sign. Rozsa et al. (2016) replaced the sign of the gradients with raw gradients. Dong et al. (2018) applied momentum to FGSM, and Kurakin et al. (2016) further extended FGSM to a targeted attack.

Adversarial examples (Goodfellow, 2014) which produce incorrect decision from machine learning models can be obtained using a number of techniques (Moosavi-Dezfooli et al., 2016; Madry et al., 2017; Carlini & Wagner, 2017) which are now available in tools such as Cleverhans (Papernot et al., 2016a). While these methods rely on digital transformation of the input, physically realizable attacks in form of patches or stickers that can be added to an image have also been developed. Adversarial patch (Brown et al., 2017) and localized and visible adversarial noise (LaVAN) (Karmon et al., 2018b) methods generate patches that are universal and can be used to attack any image and are targeted to force classifier to output a particular output label. Their goal is to generate universal noise “patches” that can be physically printed and put on any image, in either a black-box (when the attacked network is unknown) or white-box (when the attacked network is known) setup. A related attack is proposed in Karmon et al. (2018a) to investigate the blind-spots of state-of-the-art image classifiers, and the kinds of noise that can cause them to misclassify. More diffused universal adversarial digital perturbations (Moosavi-Dezfooli et al., 2017) have also been proposed that can be applied to any input. Further, the adversarial examples have been shown to transfer (Liu et al., 2016; Papernot et al., 2017) across models making them agnostic to availability of model parameters and effective against even ensemble approaches. In addition to these inference-time attacks, machine learning models have also been shown to be susceptible to data poisoning attacks during training. Last few years of efforts to defend against adversarial attacks have met with limited success. While empirical approaches such as logit pairing (Engstrom et al., 2018) and defensive distillation (Papernot et al., 2016b) have shown effectiveness against particular attack methods, more principled techniques such as robust optimization (Raghunathan et al., 2018; Zhang et al., 2019) and formal methods (Dutta et al., 2018) are limited to perturbations with bounded L_p norm. Schott et al. (2018) presents an insightful study of state of the art on attacks and defense, and proposes an analysis by synthesis defense. This arms race between attacks and defenses are not limited to computer vision, but extend to audio and natural language processing.

A number of attribution techniques have been recently proposed in literature that assign positive and negative importance to an input feature for a given decision of the machine learning model. Attributions can correspond to the actual decision or to counterfactuals. Many prominent attribution methods are based on the gradient of the predictor function with respect to the input (Simonyan et al., 2013; Selvaraju et al., 2017; Sundararajan et al., 2017). These approaches have been extended to counterfactual analysis (Datta et al., 2016) which has roots in cooperative game theory and revenue division. Counterfactual analysis is difficult when the number of possible output labels is large. We restrict our study to analysis of the actual attribution. Different attribution methods are compared in Adebayo et al. (2018). The sensitivity of these attributions to perturbations in the input are studied in Ghorbani et al. (2017) and adversarial attacks are presented to create perceptively indistinguishable images with the same prediction label but different attributions. Their results indicate that both System 1 and System 2 in our proposed approach can be independently attacked.

The goal of this paper is to combine both System 1 and System 2 to create a more resilient cognition model than any one of them alone. This is motivated by recent results in causal machine learning. Kilbertus et al. (2018) argues that causal mechanisms are typically continuous but most learning problems such as classification are anti-causal. Strong generalization is limited in anti-causal learning and access to the causal model can remove such limitations. Further, the lack of resilience to adversarial example is hypothesized to be the result of learning in anti-causal direction. We use the attributions and saliency map to identify the cause of the current decision of the machine learning model, and then construct a generator that can mask high-attribution features to explore the causal neighborhood of the input and observe model’s behavior in this neighborhood. While learning is still in the anti-causal direction, we add a layer of causal deliberative System 2 that reasons in forward direction to evaluate the robustness of the obtained attribution.

We study the connection between attribution and robustness. We observe that the features in adversarial inputs responsible for the wrong decision hinge on a relatively concentrated small set of features and hence, are not robust to masking of features with high attribution magnitude. This suggests two parallel and independent approaches to increase robustness of a two-level cognition system. First, the System 1 comprising of the machine learning model can be made independently robust via adversarial training or direct regularization to concentrate the attributions, and second, System 2 can rely on masking high attribute features to check the robustness of the model’s explanation for a specific instance.

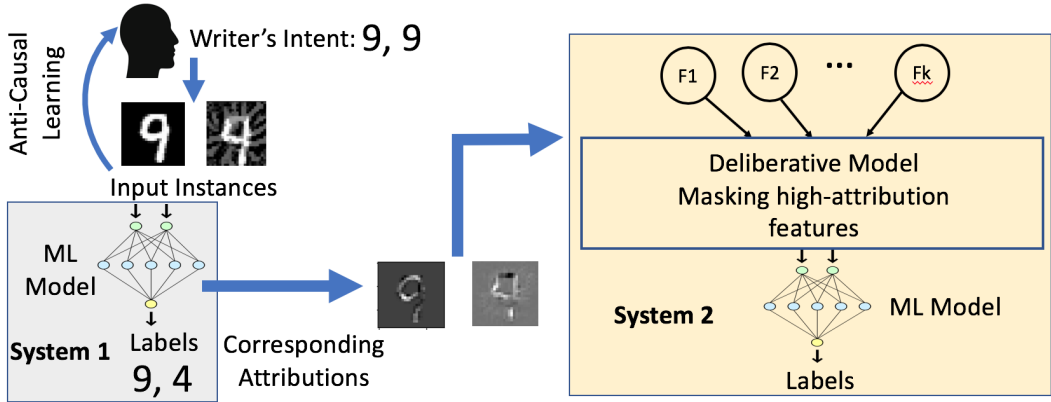


Figure 2: The overall architecture of the proposed approach motivated by the two level Kahneman’s decomposition of cognition. Typical machine learning models for classification perform anti-causal learning to determine the label from the input instance. As noted by Chalasani et al. (2018), such anti-causal reasoning lacks the natural continuity of causal mechanisms and is often not robust. But we view this model as System 1 and use attribution methods (Integrated Gradient in our experiments) to obtain features with positive and negative attributions. In this example with the MNIST dataset, we see that the adversarial perturbation that causes misclassification of 9 into 4 also significantly changes the attributions. For example, the top part of the perturbed 9 (misclassified as 4) has negative attribution. In deliberative System 2, we perform reasoning in the causal direction, and mask the high attribution features (pixels in this case) to obtain a number of input instances in the causal neighborhood of the original image. The original attributions are robust but the adversarial attributions are not robust which causes the model to assign different labels to images in the causal neighborhood of adversarial examples.

3 APPROACH

Figure 2 shows the overall architecture of the proposed approach. The anti-causal learning used in tasks such as classification are often not robust to adversarial changes (Chalasani et al., 2018). Such models can be viewed as intuitive System 1 that makes fast decisions. While these systems can definitely be made more robust and resilient to adversarial methods (Madry et al., 2017; Raghuathan et al., 2018), our focus is on analyzing the attributions generated from the machine learning models on the inputs to detect their adversarial nature. We use image benchmarks as inputs and pixels as the features. But our approach can be adapted to other application domains and to more higher-level abstract features. As discussed earlier, there are several recently proposed attribution methods that can be used to obtain the positive and negative attribution of different features for a given decision of the machine learning model. For a given neural network machine learning model M and input x , the attribution for the model output $M(x)$ is computed for the features F_1, F_2, \dots, F_k of the input.

As recognized in recent comparison of feature-attribution methods (Adebayo et al., 2018), several feature attribution methods have quirks which obscure the impact of features on the model output, and evaluation of different attribution methods is challenging. In a recent paper (Sundararajan et al., 2017), the authors identify several axioms that must be satisfied by a sound attribution method, and in particular propose a specific method that satisfies these axioms, which they call Integrated Gradients. We adopt this method in our study. The attribution in Integrated Gradient uses the notion of a baseline input. For example, the all dark image can be selected as a baseline for attribution over images. The attribution is then defined as the path integral of the gradients along the straight-line path from the baseline x' to the input x , that is, the attribution for the i -th feature is

$$IG(F_i(x)) = (x_i - x'_i) \times \int_{\alpha=0}^1 \partial_i M(x' + \alpha(x - x')) d\alpha$$

where $\partial_i M(\cdot)$ denotes the gradient of $M(\cdot)$ along the i -th feature dimension.

We adopt a simple masking model for causal generation of new inputs from the features. Given an input x with baseline x' , the masking generator G can select a subset of $n \leq k$ features $\mathcal{M} = \{m_1, m_2, \dots, m_n\}$ to be masked and generate a new input $G(x, \mathcal{M})$ such that $F_i(G(x, \mathcal{M})) = F_i(x)$ for $i \notin \mathcal{M}$ and $F_i(G(x, \mathcal{M})) = F_i(x')$ for $i \in \mathcal{M}$. For example, if the baseline is a dark image and

Algorithm 1 Evaluate robustness ρ of machine learning model M on input x with granularity ϵ

- 1: Using Integrated Gradient, compute the attributions of the features F_1, F_2, \dots, F_k of x as $IG(F_1(x)), IG(F_2(x)), \dots, IG(F_k(x))$
 - 2: Sort the features in decreasing order of attributions $F_{j_1}, F_{j_2}, \dots, F_{j_k}$ such that the attributions $IG(F_{j_1}(x)) \geq IG(F_{j_2}(x)) \geq \dots \geq IG(F_{j_k}(x))$.
 - 3: Filter the features to keep only those with positive attributions $F_{j_1}, F_{j_2}, \dots, F_{j_m}$ such that $IG(F_{j_1}(x)) \geq IG(F_{j_2}(x)) \geq \dots \geq IG(F_{j_m}(x)) > 0$.
 - 4: $\delta \leftarrow \epsilon$
 - 5: $\mathcal{M}_\delta \leftarrow \{F_{j_1}, F_{j_2}, \dots, F_{j_{\lceil \epsilon m \rceil}}\}$
 - 6: **while** M produces same outputs on x and $G(x, \mathcal{M}_\delta)$, and $\delta \leq 1$ **do**
 - 7: $\delta \leftarrow \delta + \epsilon$
 - 8: $\mathcal{M}_\delta \leftarrow \{F_{j_1}, F_{j_2}, \dots, F_{j_{\lceil \delta m \rceil}}\}$
 - 9: **end while**
 - 10: **return** $\rho(M, x) = \delta$
-

the features are pixels of the image, the masking model selects a set of pixels of an input image and makes them dark. We use the attribution over features and this masking model to define a causal neighborhood of the input x as follows:

Definition: Given an input x with features F_1, F_2, \dots, F_k and the corresponding attributions $IG(F_1(x)), IG(F_2(x)), \dots, IG(F_k(x))$, we sort the features in decreasing order of their attribution $F_{j_1}, F_{j_2}, \dots, F_{j_k}$ such that the attributions $IG(F_{j_1}(x)) \geq IG(F_{j_2}(x)) \geq \dots \geq IG(F_{j_k}(x))$. We define a δ causal neighborhood of the input x for $0 \leq \delta \leq 1$ by constructing new inputs $G(x, \mathcal{M})$ obtained via masking features $\mathcal{M} \subseteq \{F_{j_1}, F_{j_2}, \dots, F_{j_{\lceil \delta k \rceil}}\}$, that is,

$$\mathcal{N}_\delta(x) = \{G(x, \mathcal{M}) \mid \mathcal{M} \subseteq \{F_{j_1}, F_{j_2}, \dots, F_{j_{\lceil \delta k \rceil}}\}\}$$

We say that an input (benign or adversarial) is δ robust with respect to its attributions if the machine learning model produces the same output on all inputs in $\mathcal{N}_\delta(x)$. We restrict our study to positive attributions. Negative attributions can be used for counterfactual robustness testing but such a counterfactual testing will need to be done with respect to all alternate decisions. We denote this restricted neighborhood as $\mathcal{N}_\delta^+(x)$ where all features F_{j_i} in \mathcal{M} are guaranteed to have positive attributions, that is, $IG(F_{j_i}(x^{\mathcal{M}})) > 0$ for all $x^{\mathcal{M}} \in \mathcal{N}_\delta^+(x)$. The following theorem follows from the monotonicity of M with respect to the inclusion of features with positive attribution:

Theorem: If the outputs of a machine learning model M on an input x and for some input in its causal neighborhood $x^{\mathcal{M}} \in \mathcal{N}_\delta^+(x)$ are different, then the model M also produces different outputs on the input x and $G(x, \mathcal{M}_\delta)$ where $\mathcal{M}_\delta = \{F_{j_1}, F_{j_2}, \dots, F_{j_{\lceil \delta k \rceil}}\}$.

This monotonicity enables us to check the robustness of a machine learning model on an input in its causal neighborhood by just considering the furthest input. If the model produces different output, then it is not robust on this input. Algorithm 1 summarizes this robustness check procedure which forms the core of the deliberative System 2.

We can use the computation of the robustness of machine learning model M on input x to define the attribution sensitivity on a dataset X as follows:

$$\mathcal{S}(M, X, \delta) = 1 - \frac{|\{x \mid x \in X \text{ and } \rho(M, x) \leq \delta\}|}{|X|}$$

where $|\cdot|$ denotes the size of the set. The central goal of the paper is to study the connection between attributions and adversarial robustness, and to argue the need for Kahneman’s System 1 and System 2 cognition for more robust perception. While such an approach will not create a fool-proof defense against adversarial examples, it is an essential first-step towards building trusted high assurance intelligent agents. We accomplish this by computing the attribution sensitivity on original datasets and adversarially perturbed datasets. An effective adversarial attack must have smaller attribution sensitivity similar to original inputs to avoid detection.

4 ANALYSIS RESULTS

We perform the attribution sensitivity analysis on three datasets to test our hypothesis that adversarial examples are less robust (more sensitive) to masking in the attribution space when compared to the robustness of original images to such masking:

- MNIST images and adversarial datasets obtained using FGSM (with different epsilon bounds), DeepFool, PGD and CW attack methods.
- ImageNet images and adversarial datasets obtained using FGSM (with different epsilon bounds), DeepFool, PGD and CW attack methods.
- Physically realizable adversarial patch Brown et al. (2017) and LaVAN Karmon et al. (2018b) attack applied to 1000 images from ImageNet. For adversarial patch attack we used a patch size of 25% for two patch types: banana and toaster. For LaVAN we used baseball patch of size 50×50 pixels.




Figure 3: (Left) Percentage of changed labels and percentage of masked attributions for adversarial attacks FGSM, DeepFool, CW, and PGD on MNIST dataset. (Right) Percentage of changed labels and fraction of masked pixels on ImageNet database. FGSM attacks with different $\epsilon = 2, 5, 10$ are evaluated for ImageNet.

4.0.1 MNIST DATASET AND DIGITAL ATTACKS

The original MNIST images show attribution robustness where removing the top 10% attributions does not change the decision of the model on 94% of the images. Even removing 20% of the attributions does not cause change in the labels assigned to 81% of the MNIST images. In contrast, masking the top 20% attribution has a much stronger impact on the adversarial images. As shown in Figure 3, this changes the assigned labels of adversarial images generated using FGSM, DeepFool, CW and PGD by 72%, 67%, 53%, and 68% respectively.

4.0.2 IMAGENET DATASET AND DIGITAL ATTACKS

The masking of a small number of high-attribution pixels does not cause a change in the labels of the original ImageNet images as shown in Figure 3. Labels assigned to 82% of ImageNet images remain unchanged when pixels corresponding to top 0.1% of attributions are masked. The fraction of unchanged labels remains at 67% even when pixels corresponding to top 0.4% of attributions are masked. In contrast, 76%, 71% and 75% of the images changed label on masking 0.4% of pixels corresponding to top attributions for the adversarial examples generated using FGSM with different epsilon values. This percentage for PGD($\epsilon = 2$) was 71%, 97% for CW and 95% for DeepFool.

4.0.3 IMAGENET DATASET AND PHYSICAL ATTACKS

Using 1000 images from the ImageNet dataset, we created adversarial patch and LaVAN attacks for the InceptionV3 deep learning model. For adversarial patch attack we used a patch size of 25% for

two patch types: banana and toaster. A total of 347 out of 1000 images were successfully perturbed and assigned banana label on applying the banana patch. Corresponding number for the toaster patch was 376. We observe that there was a change in label for 99.71% of images having the banana patch on masking top 0.4% of high attribution pixels. Similarly, 98.14% of images having the toaster patch changed label on masking top 0.4% of high attribution pixels. For LaVAN attack we used baseball patch of size 50×50 pixels. A total of 747 out of 1000 images were successfully perturbed and assigned baseball label. We observe that 99.20% of images having baseball patch changed label when we masked top 0.4% of attributions. Detection percentages for various pixel masking percentages for images with banana, toaster and baseball patch are shown in Figure 4.

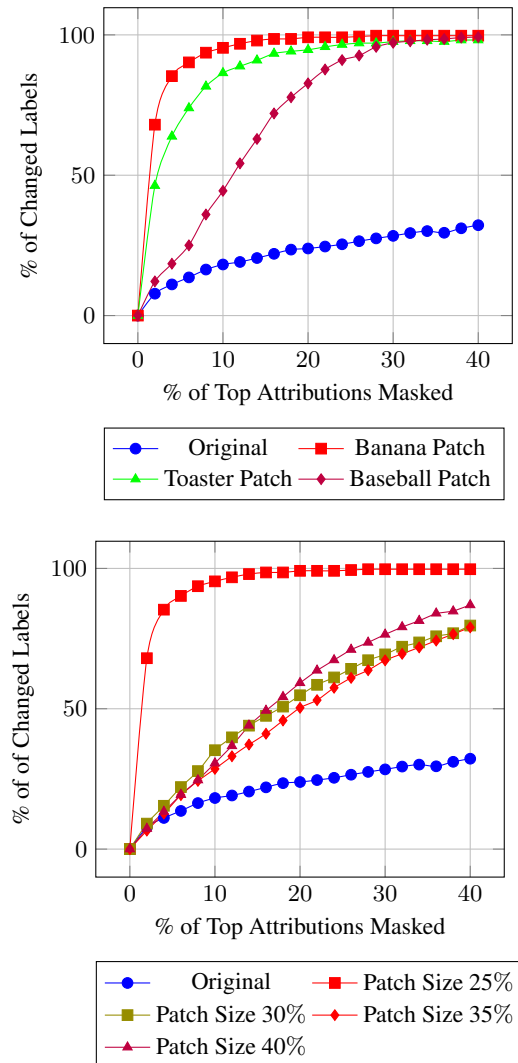


Figure 4: (Left) Percentage of correct labels and masked attributions for 1000 images from the ImageNet database is shown in blue. Banana patch and toaster patch were generated using adversarial patch method Brown et al. (2017). Baseball patch was generated using LaVAN method Karmon et al. (2018b). Dropping 0.4% of the attribution causes 99.71% of the attacks based on banana patches to be detected, 98.14% of the attacks based on toaster patch to be detected and 99.20% of the attacks based on baseball patch to be detected. (Right) Masking 0.4% of attributions causes nearly 80% of labels to change for images with adversarial patches.

We observe that most of the high attribution pixels in adversarial images is localized in the patch and masking these pixels leads to a change in the image label. An example of this behaviour can be seen for adversarial patch and LaVAN methods in Figures 5 and 6 respectively. In Figure, 5 most of the masking is localized around the banana patch; similarly, most of the masking is localized around the baseball patch in Figure 6.

A smaller size of adversarial patch is more desirable as an attacker can apply it to more surfaces. In our testing of adversarial patch method, 25% was the smallest patch size that could fool Inception classifier consistently. We applied our method on various patch sizes ranging from 25% to 40% in increments of 5% to test the robustness of our method. Results showing detection percentages for various patch sizes is presented in the right of Figure 4. Smaller patches are relatively easier to detect due to high attribute feature concentration suggesting a trade off in making attacks physically realizable and yet imperceptible, and making it robust in the attribution space.



Figure 5: (From left to right) the original image with a label of yawl; masking its top 0.2% attribution; masking its top 4% of attribution; image with a banana patch generated using adversarial patch method; masking its top 0.2% attribution; (rightmost) masking its top 4% of attribution. Most pixels are masked from the noise patch leading to a change in label. Original images are robust to small pixel masking and retain their labels.



Figure 6: (From left to right) original image; masking its top 0.2% attribution; masking its top 2% of attribution; image with a baseball patch generated using LaVAN method that causes the Inception deep neural net to label the entire image as a baseball; masking its top 0.2% attribution; (rightmost) masking top 2% of its top attributions. Most pixels are again masked from the noise patch leading to a change in label. Original images are robust to small pixel masking and retain their labels.

5 CONCLUSION

We investigated the connection between the resilience to adversarial perturbations and the attribution-based explanation of individual decisions generated by machine learning models. We observe that adversarial inputs are not robust in attribution space, that is, masking a few features with high attribution leads to change in decision of the machine learning model on the adversarial examples. In contrast, the natural inputs are robust in attribution space. This observation can be used to devise a new defense approach for machine learning models that uses attribution and incremental masking for filtering out adversarial examples. We propose such a deliberative masking based analysis approach as System 2 of Kahneman’s two layer cognition system with the actual machine learning model serving as System 1. We demonstrate the effectiveness of our approach on a set of benchmarks that include diffused digital perturbations as well as physically realizable perturbations. The physically realizable adversarial perturbations are often localized; hence, they are easier to detect using attributions. Further, attribution-based defense is agnostic to the attack method and also independent of the training data. Instead, it relies on analysis of the result and the attribution generated by a machine learning model on a specific input.

Our approach can make machine learning models more resilient to adversarial attacks because fooling this two layer cognition system requires not only attacking the original model but also ensuring that the attribution generated for the adversarial example is similar to the original examples. Both System 1 and System 2 must be simultaneously compromised for a successful adversarial attack. Our experimental results indicate that the current adversarial attack methods are not able to achieve this. The adversarial inputs generated by the state-of-the-art attack methods such as DeepFool, FGSM, CW and PGD are able to fool a machine learning model but they are not robust in the attribution space. New attack methods must be evaluated by not only measuring the observed loss of accuracy but also evaluating their robustness in attribution space.

Acknowledgement: The authors acknowledge support from the US ARL Cooperative Agreement W911NF-17-2-0196 on Internet of Battlefield Things (IoBT) and National Science Foundation(NSF) Awards #1822976, #1750009, #1740079 and #1422257. UCF authors also acknowledge an award from the Florida Cybersecurity Center, Royal Bank of Canada, and the Air Force Young Investigator Award to Sumit Jha.

REFERENCES

- Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, and Been Kim. Sanity checks for saliency maps. In *NIPS*, pp. 9525–9536, 2018.
- Tom B Brown, Dandelion Mané, Aurko Roy, Martín Abadi, and Justin Gilmer. Adversarial patch. *arXiv preprint arXiv:1712.09665*, 2017.
- Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *ISSP*, pp. 39–57. IEEE, 2017.
- Prasad Chalasanani et al. Adversarial learning and explainability in structured datasets. *arXiv preprint arXiv:1810.06583*, 2018.
- Anupam Datta, Shayak Sen, and Yair Zick. Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems. In *ISSP*, pp. 598–617. IEEE, 2016.
- Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9185–9193, 2018.
- Souradeep Dutta, Susmit Jha, Sriram Sankaranarayanan, and Ashish Tiwari. Output range analysis for deep feedforward neural networks. In *NASA Formal Methods Symposium*, pp. 121–138. Springer, 2018.
- Logan Engstrom, Andrew Ilyas, and Anish Athalye. Evaluating and understanding the robustness of adversarial logit pairing. *arXiv preprint arXiv:1807.10272*, 2018.
- Amirata Ghorbani, Abubakar Abid, and James Zou. Interpretation of neural networks is fragile. *arXiv preprint arXiv:1710.10547*, 2017.
- Ian J et al. Goodfellow. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- David Gunning. Explainable artificial intelligence (xai). *Defense Advanced Research Projects Agency (DARPA)*, *nd Web*, 2017.
- Kaiming He et al. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, pp. 1026–1034, 2015.
- Andrew Ilyas et al. The robust manifold defense: Adversarial training using generative models. *arXiv preprint arXiv:1712.09196*, 2017.
- Susmit Jha, Uyeong Jang, Somesh Jha, and Brian Jalaian. Detecting adversarial examples using data manifolds. In *MILCOM*, pp. 547–552. IEEE, 2018.
- Daniel Kahneman and Patrick Egan. *Thinking, fast and slow*, volume 1. 2011.
- Danny Karmon, Daniel Zoran, and Yoav Goldberg. Lavan: Localized and visible adversarial noise. *arXiv preprint arXiv:1801.02608*, 2018a.
- Danny Karmon, Daniel Zoran, and Yoav Goldberg. LaVAN: Localized and visible adversarial noise. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2507–2515, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018b. PMLR. URL <http://proceedings.mlr.press/v80/karmon18a.html>.

- Niki Kilbertus, Giambattista Parascandolo, and Bernhard Schölkopf. Generalization in anti-causal learning. *arXiv preprint arXiv:1812.00524*, 2018.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. *Commun. ACM*, 60(6):84–90, May 2017. ISSN 0001-0782. doi: 10.1145/3065386. URL <http://doi.acm.org/10.1145/3065386>.
- Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*, 2016.
- Guanbin Li and Yizhou Yu. Visual saliency based on multiscale deep features. In *CVPR*, pp. 5455–5463, 2015.
- Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. Delving into transferable adversarial examples and black-box attacks. *arXiv preprint arXiv:1611.02770*, 2016.
- Aleksander Madry et al. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *CVPR*, pp. 2574–2582, 2016.
- Seyed-Mohsen Moosavi-Dezfooli et al. Universal adversarial perturbations. In *CVPR*, pp. 1765–1773, 2017.
- Nicolas Papernot, I Goodfellow, et al. Cleverhans v2. 0.0: an adversarial machine learning library. *arXiv preprint arXiv:1610.00768*, 2016a.
- Nicolas Papernot, Patrick McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. Distillation as a defense to adversarial perturbations against deep neural networks. In *ISSP’16*, 2016b.
- Nicolas Papernot et al. Practical black-box attacks against machine learning. In *ACCS’17*, 2017.
- Aditi Raghunathan, Jacob Steinhardt, and Percy Liang. Certified defenses against adversarial examples. *arXiv preprint arXiv:1801.09344*, 2018.
- Andras Rozsa, Ethan M Rudd, and Terrance E Boult. Adversarial diversity and hard positive generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 25–32, 2016.
- Lukas Schott, Jonas Rauber, Matthias Bethge, and Wieland Brendel. Towards the first adversarially robust neural network model on mnist. 2018.
- Ramprasaath R Selvaraju et al. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, pp. 618–626, 2017.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *ICML*, pp. 3319–3328. JMLR. org, 2017.
- Christian Szegedy et al. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- Christian Szegedy et al. Rethinking the inception architecture for computer vision. In *CVPR*, pp. 2818–2826, 2016.
- Florian Tramèr et al. Ensemble adversarial training: Attacks and defenses. *arXiv preprint arXiv:1705.07204*, 2017.
- Wayne Xiong et al. Toward human parity in speech recognition. *TASLP*, pp. 2410–2423, 2017.
- Kwang Moo Yi et al. Lift: Learned invariant feature transform. In *ECCV*, pp. 467–483. Springer, 2016.
- Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric P Xing, Laurent El Ghaoui, and Michael I Jordan. Theoretically principled trade-off between robustness and accuracy. *arXiv:1901.08573*, 2019.