
MASSIF: Interactive Interpretation of Adversarial Attacks on Deep Learning

Nilaksh Das*

Georgia Institute of Technology
Atlanta, GA, USA
nilakshdas@gatech.edu

Haekyu Park*

Georgia Institute of Technology
Atlanta, GA, USA
haekyu@gatech.edu

Zijie J. Wang

Georgia Institute of Technology
Atlanta, GA, USA
jayw@gatech.edu

Fred Hohman

Georgia Institute of Technology
Atlanta, GA, USA
fredhohman@gatech.edu

Robert Firstman

Georgia Institute of Technology
Atlanta, GA, USA
rfirstman6@gatech.edu

Emily Rogers

Georgia Institute of Technology
Atlanta, GA, USA
Emily.Rogers@gtri.gatech.edu

Duen Horng (Polo) Chau

Georgia Institute of Technology
Atlanta, GA, USA
polo@gatech.edu

*Authors contributed equally.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CHI '20 Extended Abstracts, April 25–30, 2020, Honolulu, HI, USA.

© 2020 Copyright is held by the author/owner(s).

ACM ISBN 978-1-4503-6819-3/20/04.

<http://dx.doi.org/10.1145/3334480.3382977>

Abstract

Deep neural networks (DNNs) are increasingly powering high-stakes applications such as autonomous cars and healthcare; however, DNNs are often treated as “black boxes” in such applications. Recent research has also revealed that DNNs are highly vulnerable to adversarial attacks, raising serious concerns over deploying DNNs in the real world. To overcome these deficiencies, we are developing MASSIF, an interactive tool for deciphering adversarial attacks. MASSIF identifies and interactively visualizes neurons and their connections inside a DNN that are strongly activated or suppressed by an adversarial attack. MASSIF provides both a high-level, interpretable overview of the effect of an attack on a DNN, and a low-level, detailed description of the affected neurons. MASSIF’s tightly coupled views help people better understand which input features are most vulnerable and important for correct predictions.

Author Keywords

Deep learning interpretability, adversarial attack, visual analytics, scalable summarization, attribution graph

Introduction

Deep neural networks (DNNs) have demonstrated significant success in a wide spectrum of applications [3, 5, 6, 9]. However, they have been found to be highly vulnerable to *adversarial attacks*: typically small, human-imperceptible

perturbations on inputs that fool DNNs into making incorrect predictions [2, 4, 8, 11]. This jeopardizes many DNN-based technologies, especially in security and safety-critical applications such as autonomous driving and data-driven healthcare. To make deep learning more robust against such malicious attacks, it is essential to understand how the attacks permeate DNN models [12, 14]. Interpreting, and ultimately defending against adversarial attacks, is nontrivial due to a number of challenges:

- C1. **Scaling to large datasets.** Since DNNs contain thousands of neurons and connections, summarizing how an attack affects a model requires analyzing the model's behavior over entire datasets and developing scalable representations that characterize the attack [1, 7].
- C2. **Entangled features and connections between benign and attacked inputs.** A natural approach for understanding adversarial attacks is to compare a model's operations on benign and attacked inputs, which could help people understand *where* and *why* predictions within a model diverge. However, designing an effective comparison can be challenging because the features contributing to the differences may correlate with one or more features in both the benign and attacked classes. For example, Figure 1 shows that a feature representing “ivory face with dark eyes and nose” (center purple node) is important for both the benign *panda* class and the attacked *armadillo* class. This shared feature correlates with a feature for panda (e.g., “black & white patches”), while also correlating with a few features for armadillo (e.g., “scales”, “crossed pattern”).
- C3. **Diverse feature vulnerability.** Given an adversarial attack, some learned features may be more easily manipulated and vulnerable than others. For example,

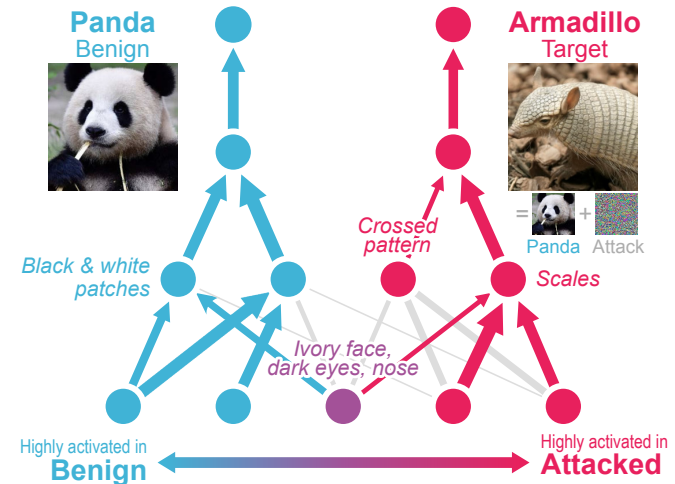


Figure 1: Adversarial attacks confuse DNNs to make incorrect predictions, e.g., attacking benign *panda* images so they are misclassified as *armadillo*. We aim to understand where such attacks occur inside the model and what features are used.

manipulating a “basketball” feature into an “orange” feature is easier (similar shape and color) than changing it into a “truck” feature. As features exhibit a spectrum of vulnerability, enabling users to visualize and understand an attack under varying levels of severity could help them design stronger countermeasures.

Contributions

To address the aforementioned challenges, we are developing **MASSIF**, an interactive visualization tool for interpreting adversarial attacks on deep learning models. Our ongoing work presents the following contributions:

1. **Novel graph-based comparison.** To discover the features and connections activated or suppressed by an attack, we adapt the recently proposed *attribution*

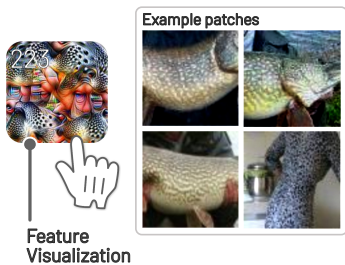


Figure 2: Each neuron in an attribution graph is represented with its feature visualization. When a user hovers over a neuron, example dataset patches are shown for context.

graph [7] in a **novel way to visualize, summarize, and compare a model’s response to benign and attacked data**. The original attribution graph aims to highlight how a model’s learned features interact to make predictions for a single class, by representing highly activated neurons as vertices and their most influential connections as edges. Our main idea for MASSIF is to generate and integrate two attribution graphs: one for the benign data and another for the attacked data, as illustrated in Figure 1. The aggregated graph helps us understand which features are shared by both benign and attacked data (e.g., purple, center feature in Figure 1), which are solely activated by the benign data (blue, far left), and which are by the attacked data (red, far right). Importantly, MASSIF also helps users more easily discover where a prediction starts to “diverge”, honing in to the **critical parts** of the model that the attack is exploiting.

2. **Fractionation of neurons based on vulnerability.** To help users prioritize their inspection of neurons, we develop a new way to sort and group them based on their vulnerability, i.e., “how easily can a neuron be activated or suppressed by an attack.” Our main idea is to vary an attack’s strength (or severity) and record all neuron activations. Neurons that are easily activated or suppressed by solely a weak attack may warrant focused inspection since they can be easily manipulated with little effort.

System Design and Implementation

This section describes MASSIF’s interface and visualization design (see Figure 3). Users can select a benign class, an attacked class, and the type and severity of an attack from the top control panel (Figure 3A). Then, MASSIF generates attribution graphs using a Python backend, and displays them to the user in a web-based interface built with HTML, CSS, JavaScript, and D3 (Figure 3B).

Nodes: Features activated (or suppressed) by attacks.

In attribution graphs, nodes represent DNN neurons which are trained to detect particular features in input data. To interpret what features a neuron detects, MASSIF represents each neuron with its feature visualization: a synthesized image that maximizes the neuron’s activation [10]. Users can hover over any neuron’s feature visualization to also display example image patches from the dataset that most activate that neuron. For example, as seen in Figure 2, the feature visualization (left) and dataset examples (right) describe a neuron that detects a dotted pattern in scales.

We divide an attribution graph’s neurons into three groups by their attack response. First, **suppressed neurons** are highly activated by benign inputs but become suppressed by adversarial inputs. These represent crucial features for the benign class, but the model fails to detect them when exposed to the attack. Second, **emphasized neurons** are not noticeably activated by benign inputs but become highly activated by adversarial inputs. These represent features that are typically not important for the benign class, but the model detects them as important features of the attacked class. Third, **shared neurons** are highly activated by and important to both benign and adversarial inputs.

We visually distinguish these three neuron groups with different colors and positions in the attribution graph view (Figure 3B). Suppressed neurons are colored **blue** and positioned on the left (Figure 3B.1). Emphasized neurons are colored **orange** and positioned on the right (Figure 3B.3). Shared neurons are colored **purple** and positioned in the middle between suppressed neurons and emphasized neurons (Figure 3B.2). The result is a visualization that disentangles and compares the DNN features and connections from the benign and attacked data.

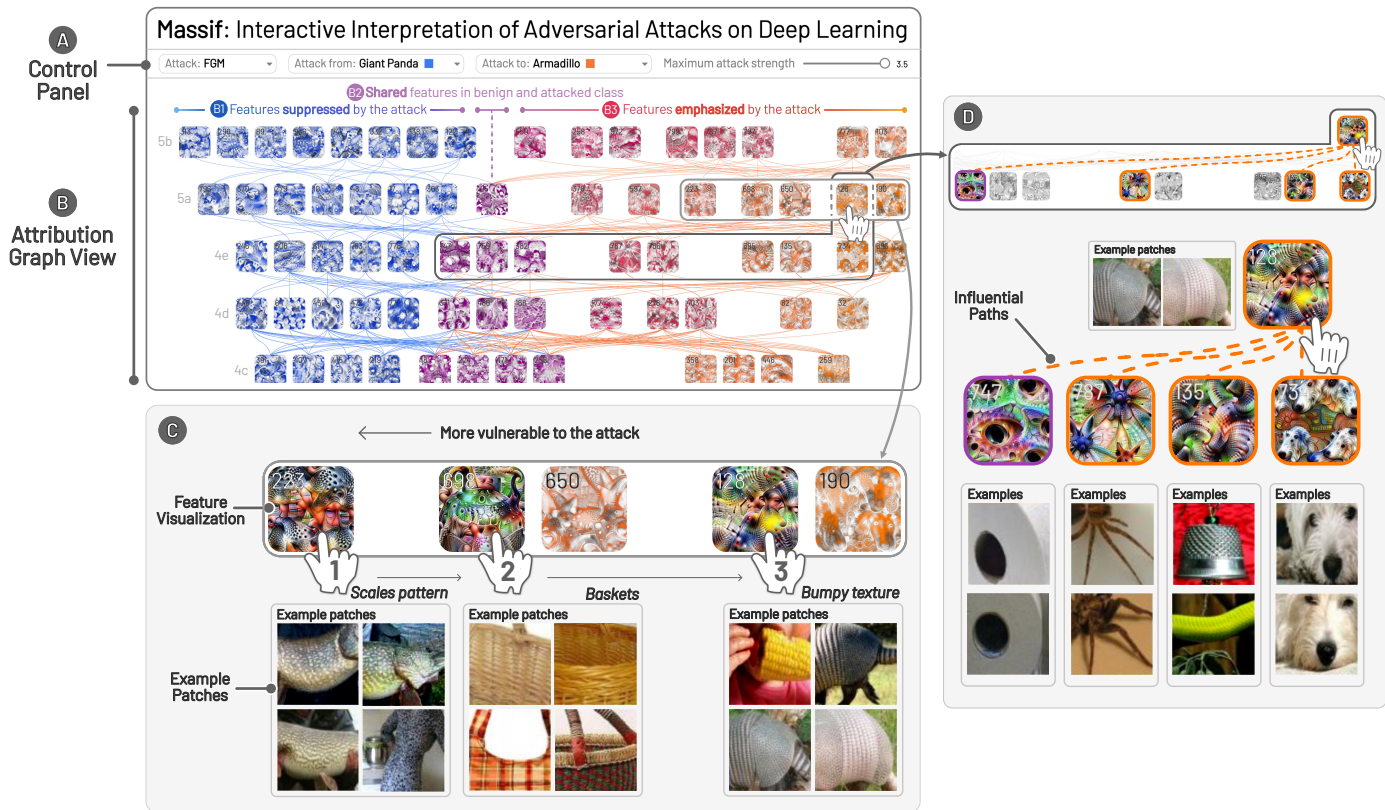


Figure 3: The MASSIF interface. A user Hailey is studying the targeted Fast Gradient Method (FGM) attack performed on the InceptionV1 model. Using the control panel (A), she selects “giant panda” as the benign class and “armadillo” as the attack target class. MASSIF generates an attribution graph (B), which shows Hailey the neurons within the network that are suppressed in the attacked images (B1, blue), shared by both benign and attacked images (B2, purple), and emphasized only in the attacked images (B3, orange). Each neuron is represented by a node and its feature visualization (C). Hovering over any neuron displays example dataset patches that maximally activate the neuron, providing stronger evidence for what a neuron has learned to detect. Hovering over a neuron also highlights its most influential connections from the previous layer (D), allowing Hailey to determine where in the network the prediction diverges from the benign class to the attacked class.

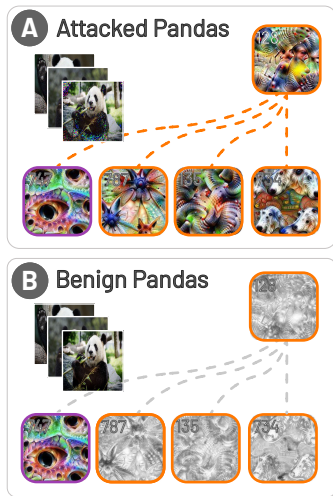


Figure 4: Edges in an attribution graph represent influential connections between neurons. (A) above shows a part of an attribution graph from Figure 3D. An emphasized neuron is connected to one shared neuron (purple) and three emphasized neurons (orange) from the previous layer. (B) shows how the same part of the attribution graph looks different for benign images. The three emphasized neurons in the previous layer are not activated by the benign inputs, which causes the emphasized neuron in the current layer to be less activated.

Fractionation: Characterizing neuron vulnerability.

Within each group of neurons, we further distinguish them based on their vulnerability. A neuron is considered more vulnerable if its activation changes greatly under weaker attacks. We encode neuron vulnerability using its position and color within its group. More vulnerable neurons are located closer to shared neurons, since they are on the border of the benign and attacked classes, and cause misclassification under weaker attacks. Suppressed neurons that are closer to the border with shared neurons are colored purple/blue, and emphasized neurons that are closer to the border with shared neurons are colored purple/orange.

Edges: Explaining why features are activated (or suppressed).

In attribution graphs, edges represent influential connections between neurons that most interact with each other to represent a particular class [7]. These connections can explain why an adversarial feature is detected and why some features are suppressed using attribution.

With MASSIF, users can interactively visualize attribution, i.e., drilling down into specific subgraphs by hovering over a neuron and highlighting its previous connections which maximally influence the neuron’s activation (Figure 3D).

When inspecting a particular emphasized neuron n , the highlighted connected neurons from the previous layer will be either shared neurons or other emphasized neurons, as seen in Figure 4A. Users can observe that the connected emphasized neurons from the previous layer contribute highly to the activation of neuron n . Similarly, to understand why a feature becomes suppressed by an attack, inspecting a particular suppressed neuron m highlights its connected neurons from the previous layer. These will be either shared neurons or suppressed neurons. The neuron m is less activated since the influential neurons from the previous layer are no longer activated, i.e., they are suppressed as well.

Preliminary Results

We present usage scenarios showing how MASSIF can help users better understand adversarial attacks on deep learning models. Our user Hailey is studying a targeted version of Fast Gradient Method [4] applied on the InceptionV1 model [13]. The model is trained on the ImageNet dataset, which contains over 1.2 million images across 1,000 classes. Using the control panel (Figure 3A), she selects “giant panda” as the benign class and “armadillo” as the target class. She sets the maximum attack strength 3.5.

Which neurons are attacked? Hailey starts by finding which specific neurons are attacked to narrow down the part of the model to investigate. In the attribution graph view (Figure 3B), she hovers over the suppressed neurons (Figure 3B.1) and the emphasized neurons (Figure 3B.3). She sees which features are emphasized and suppressed using the neuron feature visualization and dataset example patches. Exploring these features, she finds the `mixed5a` layer interesting because three emphasized neurons (223, 698, and 128) in `mixed5a` look related to armadillo skin patterns (Figure 3C). She decides to focus on these emphasized neurons in `mixed5a`.

Which neurons are easily attacked? To efficiently devise a countering defense, Hailey wants to prioritize the neurons and investigate them in order. She knows that MASSIF fractionates the neurons according to how easily they are attacked; therefore, she checks how the emphasized neurons in `mixed5a` are separated (Figure 3C). By hovering over the neurons from left to right, she observes that “scales pattern” is most vulnerable, followed by “baskets” and “bumpy texture” neurons. She decides to explore the neurons in this order, since she presumes that it is more efficient to protect more vulnerable neurons.

Why are these neurons attacked? Hailey now wants to know how to protect the attacked neurons related to armadillo skin patterns. She sequentially observes the attribution for the “scales pattern”, “baskets”, and “bumpy texture” neurons. Upon inspecting the “bumpy texture” neuron (Figure 3D), MASSIF shows that four neurons in the previous layer are highly interacting with it: a shared neuron representing “black circle,” and emphasized neurons representing “spider legs”, “granular texture”, and “a white hairy dog’s face.” As the three emphasized neurons in the previous layer can be the primary reason behind the detection of “bumpy texture”, she decides to investigate these neurons more using MASSIF.

Ongoing Work

Interactive neuron editing. MASSIF currently visualizes the neurons that are activated or suppressed by an attack under varying degrees of severity. We are working on extending MASSIF’s interactivity by allowing real-time neuron editing, e.g., deletion. This would allow a user to actively identify vulnerable neurons using our visualization and interactively remove them from the DNN to observe its effect in real-time. Neuron deletion would mask the activations of a particular neuron, potentially preventing the malicious effect of a targeted attack to propagate deeper into the network. This would enable a user to preemptively edit a DNN to enhance its robustness to adversarial attacks. For example, a user may identify and choose to delete a shared neuron that only feeds into emphasized neurons, preventing adversarially activated neurons from having any effect in the subsequent layers of the network, thus thwarting the targeted attack from succeeding.

Planned evaluation. We plan to evaluate the effectiveness of our visualization tool coupled with interactive neuron editing through in-lab user studies where participants

seek to increase the robustness of a large-scale, pretrained DNN model. We will recruit students with basic knowledge of deep learning models. All participants will be asked edit the DNN for different benign-attacked class pairs and will be evaluated on the basis of reduction in targeted attack success rate. We will also conduct pre-test and post-test surveys to evaluate whether MASSIF gave any deeper insights into the failure modes of the studied DNN and what factors the participants considered while editing the DNN to increase its robustness to adversarial attacks.

Conclusion

We present MASSIF, an interactive system we are developing that visualizes how adversarial attacks permeate DNN models and cause misclassification. MASSIF generates and visualizes multiple attribution graphs as a summary of what features are important for a particular class (e.g., benign or attacked class) and how the features are related. MASSIF enables flexible comparison between benign and attacked attribution graphs, highlighting where and why the attribution graphs start to diverge, ultimately helping people better interpret complex deep learning models, their vulnerabilities, and how to best construct defenses.

Acknowledgements

This work was supported in part by NSF grants IIS-1563816, CNS-1704701, NASA NSTRF, DARPA GARD, gifts from Intel (ISTC-ARSA), NVIDIA, Google, Symantec, Yahoo! Labs, eBay, Amazon.

REFERENCES

- [1] Shan Carter, Zan Armstrong, Ludwig Schubert, Ian Johnson, and Chris Olah. 2019. Activation atlas. *Distill* 4, 3 (2019), e15.
- [2] Shang-Tse Chen, Cory Cornelius, Jason Martin, and Duen Horng Polo Chau. 2018. Shapeshifter: Robust

- physical adversarial attack on faster r-cnn object detector. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 52–68.
- [3] Andre Esteva, Alexandre Robicquet, Bharath Ramsundar, Volodymyr Kuleshov, Mark DePristo, Katherine Chou, Claire Cui, Greg Corrado, Sebastian Thrun, and Jeff Dean. 2019. A guide to deep learning in healthcare. *Nature medicine* 25, 1 (2019), 24.
- [4] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and Harnessing Adversarial Examples. *CoRR* abs/1412.6572 (2014).
- [5] Sorin Grigorescu, Bogdan Trasnea, Tiberiu Cocias, and Gigel Macesanu. 2019. A survey of deep learning techniques for autonomous driving. *Journal of Field Robotics* (2019).
- [6] Guodong Guo and Na Zhang. 2019. A survey on deep learning based face recognition. *Computer Vision and Image Understanding* 189 (2019), 102805.
- [7] Fred Hohman, Haekyu Park, Caleb Robinson, and Duen Horng Chau. 2019. Summit: Scaling Deep Learning Interpretability by Visualizing Activation and Attribution Summarizations. *IEEE VIS* (2019).
- [8] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. 2016. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533* (2016).
- [9] Ali Bou Nassif, Ismail Shahin, Imtinan Attili, Mohammad Azzeh, and Khaled Shaalan. 2019. Speech recognition using deep neural networks: A systematic review. *IEEE Access* 7 (2019), 19143–19165.
- [10] Chris Olah, Alexander Mordvintsev, and Ludwig Schubert. 2017. Feature Visualization. *Distill* (2017). DOI:<http://dx.doi.org/10.23915/distill.00007> <https://distill.pub/2017/feature-visualization>.
- [11] Yao Qin, Nicholas Carlini, Ian Goodfellow, Garrison Cottrell, and Colin Raffel. 2019. Imperceptible, robust, and targeted adversarial examples for automatic speech recognition. *arXiv preprint arXiv:1903.10346* (2019).
- [12] Andrew Slavin Ross and Finale Doshi-Velez. 2018. Improving the adversarial robustness and interpretability of deep neural networks by regularizing their input gradients. In *Thirty-second AAAI conference on artificial intelligence*.
- [13] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1–9.
- [14] Guanhong Tao, Shiqing Ma, Yingqi Liu, and Xiangyu Zhang. 2018. Attacks meet interpretability: Attribute-steered detection of adversarial samples. In *Advances in Neural Information Processing Systems*. 7717–7728.