

# A Practical Upper Bound for the Worst-Case Attribution Deviations

Fan Wang<sup>1,2</sup>      Adams Wai-Kin Kong<sup>1</sup>

<sup>1</sup>School of Computer Science and Engineering, Nanyang Technological University

<sup>2</sup>Rapid-Rich Object Search (ROSE) Lab, IGP, Nanyang Technological University

fan005@e.ntu.edu.sg      adamskong@ntu.edu.sg

## Abstract

*Model attribution is a critical component of deep neural networks (DNNs) for its interpretability to complex models. Recent studies bring up attention to the security of attribution methods as they are vulnerable to attribution attacks that generate similar images with dramatically different attributions. Existing works have been investigating empirically improving the robustness of DNNs against those attacks; however, none of them explicitly quantifies the actual deviations of attributions. In this work, for the first time, a constrained optimization problem is formulated to derive an upper bound that measures the largest dissimilarity of attributions after the samples are perturbed by any noises within a certain region while the classification results remain the same. Based on the formulation, different practical approaches are introduced to bound the attributions above using Euclidean distance and cosine similarity under both  $\ell_2$  and  $\ell_\infty$ -norm perturbations constraints. The bounds developed by our theoretical study are validated on various datasets and two different types of attacks (PGD attack and IFIA attribution attack). Over 10 million attacks in the experiments indicate that the proposed upper bounds effectively quantify the robustness of models based on the worst-case attribution dissimilarities.*

## 1. Introduction

Attribution methods play an important role in deep learning applications as one of the subareas of explainable AI. Practitioners use attribution methods to measure the relative importance among different features and to understand the impacts of features contributing to the model outputs. They have been widely used in a number of critical real-world applications, such as risk management [2], medical imaging [24, 29] and drug discovery [13]. In particular, attributions are supposed to be secure and resistant to external manipulation such that proper explanations can be applied to safety-sensitive applications. Regulations are

also deployed in countries to enforce the interpretability of deep learning models for a ‘right to explain’ [10]. Although attribution methods have been extensively studied [18, 25, 28, 31, 36, 39], recent works reveal that they are vulnerable to visually imperceptible perturbations that drastically alter the attributions and keep the model outputs unchanged [6, 8].

Prior works [3, 4, 12, 23, 30, 32, 33] investigate the attribution robustness based on empirical and statistical estimations over entire dataset. However, current attribution robustness works are unable to evaluate how robust the model is given any arbitrary test point, perturbed or unperturbed. In this paper, we study the problem of finding the worst attribution perturbation within certain predefined regions. Specifically, given a trained model and an image sample, we propose theoretical upper bounds of the attribution deviations from the unperturbed ones. As far as we know, this is the first attempt to provide an upper bound of attribution differences.

In this paper, the general upper bound for attribution deviation is first quantified as the maximum changes of attributions after the samples are perturbed while classification results remain the same. Two cases are analyzed, including with and without label constraint, which refers to the classification labels being unchanged and changed, respectively, after the original samples are attacked. For each case, two mostly used perturbation constraints,  $\ell_2$  and  $\ell_\infty$ -norm, are considered to compute the upper bound. For  $\ell_2$ -norm constraint, our approach is based on the first-order Taylor series of model attribution, and a tight upper bound ignoring the label constraint is computed from the singular value of the attribution gradient.  $\ell_\infty$ -norm constraint is more complicated because the upper bound is a solution of a concave quadratic programming with box constraints, which is an NP-hard problem. Thus, two relaxation approaches are proposed. Moreover, a more restricted bound constrained on the unchanged label is also studied. In this study, Euclidean distance and cosine distance, which are also employed in the previous empirical studies [4, 30, 32], are used as dissimilarity functions to measure attribution difference. We

summarize the contributions of this paper as follows:

- We formally define the general upper bound for attribution deviation as the optimization problem with norm constraint and label constraint to find the maximum change of attributions after samples being attacked. According to the best knowledge of the authors, it has not been studied before.
- The tight upper bounds for  $\ell_2$ -norm constrained attacks with and without classification label constraints are proposed based on the first-order Taylor series. The proposed bound without label constraints generalizes to all gradient-based attribution methods, and the one with label constraints is applicable to all attribution methods satisfying the axiom of completeness.
- Two different approaches are provided to bound the  $\ell_\infty$ -norm constrained attacks above, which uses an  $\ell_p$ -norm relaxation and a mathematical property of the quadratic form.
- The experimental results show that the upper bounds derived in this paper can effectively bound the attribution differences between all 10 million attacked samples and their corresponding original samples from different models, datasets, attack methods and parameters choices.

The rest of this paper is organized as follows. We start with an introduction to notations and related works. The formulation of the general upper bound for attribution deviation is defined in Sec. 3. Specific methods to find the upper bounds in different scenarios are provided in Sec. 4. In Sec. 5, detailed experimental results are presented and the paper concludes in Sec. 6.

## 2. Preliminaries and related works

We consider a twice-differentiable classifier  $f$  that maps the input set  $\mathcal{D} = \{(\mathbf{x}^{(i)}, y^{(i)})\}_{i=1}^n$  to the logits,  $f: \mathbb{R}^d \rightarrow \mathbb{R}^k$ , where  $\mathbf{x}^{(i)} \in \mathbb{R}^d$  and  $y^{(i)} \in \{1, \dots, k\}$  represent the  $i$ -th sample and its ground truth label. The non-bold version  $x_k$  represents the  $k$ -th feature of  $\mathbf{x}$  and  $f_y$  is the logit at label  $y$ . The model attribution of the input sample given label  $y$  is computed by  $g^y: \mathbb{R}^d \rightarrow \mathbb{R}^d$ , and we denote the attribution of  $\mathbf{x}$  by  $g^y(\mathbf{x})$ .

### 2.1. Model attribution

The model attribution studies the importance of each input feature  $x_i$  that contributes to the final output  $f_y(\mathbf{x})$ . We can classify the most used attribution methods into two categories, perturbations-based methods [36, 39] and backpropagation-based methods [1, 25], which include gradient-based methods. In particular, in this paper, we focus on the most commonly used gradient-based attribution

methods, saliency map (SM), gradient\*input and integrated gradient (IG). Saliency map [28] is defined as the gradients of output with respect to the input. Gradient\*input [26] is computed by element-wise multiplication of input features and the gradients. Integrated gradients [31] is defined as line integral of gradients from a baseline image  $\mathbf{a}$  to the input image  $\mathbf{x}$  weighted by their difference<sup>1</sup>. It is worth noting that IG satisfies the axiom of completeness,  $\sum_i g_i^y(\mathbf{x}) = f_y(\mathbf{x})$ , which builds a direct connection between the attributions and model outputs. The mathematical expressions and examples of the attribution methods are given in Table 1.

### 2.2. Attribution robustness

It has been discovered in the literature that model attributions can be easily sabotaged by adversaries. Similar to adversarial examples [9], human-indistinguishable perturbations can be also augmented to natural images that, though classification results remain unchanged, misdirect the model attributions towards meaningless interpretations [8] or any predefined arbitrary patterns which are unrelated to the original images [6].

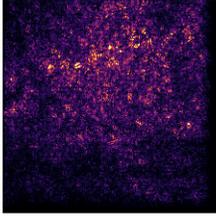
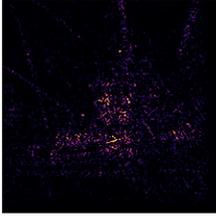
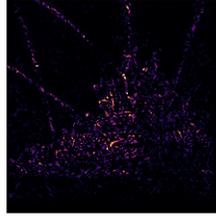
To mitigate the threat of being attacked, researchers have also worked on training attribution robust models. The most considered techniques are adapted from adversarial training [19], and they minimize the differences between original and the worst-case perturbed attributions. Boopathy *et al.* [3] and Chen *et al.* [4] consider the  $\ell_1$ -norm distance to measure the difference between attributions, and Ivankay *et al.* [12] uses Pearson correlation coefficient. Singh *et al.* [30] and Wang *et al.* [33] choose  $\ell_2$ -norm distance, where the former minimizes the spatial correlation between image and attribution using a soft-margin triplet loss, and the latter shows the smoothness of the decision surface is related to attribution robustness based on a geometric understanding. Wang & Kong [32] emphasizes the directions of attributions using the relationship between Kendall’s rank correlation and cosine similarity and protects the attribution based on the latter. However, none of the aforementioned methods defines a quantitative measurement to the attribution changes after perturbation. More clearly, the attributions are not guaranteed to be protected for all perturbations within the allowable region that do not alter the classification outputs.

## 3. Formulation of general upper bound for attribution deviations

In this section, we formally define the general upper bound for attribution deviation as a measurement of attribution robustness. Recall that adversaries incapacitate

<sup>1</sup>The baseline is chosen to be a black image ( $\mathbf{a} = \mathbf{0}$ ) in this paper if not specifically stated. Without loss of generality,  $f_y(\mathbf{a}) = 0$ .

Table 1. Mathematical expressions and visual examples of the selected attribution methods. Attributions have been taken absolute values and are presented heatmaps to reflect relative importance among pixels. The baseline  $\mathbf{a}$  of IG is chosen as a black image.  $\otimes$  denotes the element-wise multiplication.

Original image	Saliency map $\nabla f_y(\mathbf{x})$	Input*gradient $\mathbf{x} \otimes \nabla f_y(\mathbf{x})$	Integrated gradients $(\mathbf{x} - \mathbf{a}) \otimes \int_0^1 \nabla f_y(\mathbf{a} + \alpha(\mathbf{x} - \mathbf{a})) d\alpha$
			

the attributions of neural networks by adding imperceptible noises to natural images. For an attributional robust model, on the contrary, the imperceptible noises should not change the interpretability of attributions, *i.e.*, images perturbed by noises should provide *similar* attributions as the original ones. To evaluate such resistance against adversaries, it is essential to find an upper bound that represents the worst-case dissimilarity of attributions after the original images being perturbed. Thus, we define the general upper bound for attribution deviations as follows.

**Definition 1** (General upper bound for attribution deviations). Given a trained neural network  $f$ , a fixed allowable region for perturbation  $\delta$ ,  $\mathcal{B}_\varepsilon = \{\delta : \|\delta\|_p \leq \varepsilon\}$ , and an input sample  $\mathbf{x}$ , the *general upper bound for attribution deviations*  $T(\varepsilon; \mathbf{x})$  is defined that, for all perturbations  $\delta \in \mathcal{B}_\varepsilon$ , if  $\arg \max_k f_k(\mathbf{x}) = \arg \max_k f_k(\mathbf{x} + \delta)$ , then their corresponding attributions satisfy that  $D(g^y(\mathbf{x}), g^y(\mathbf{x} + \delta)) \leq T(\varepsilon; \mathbf{x})$ .

In the above definition,  $D(\cdot, \cdot)$  is a dissimilarity metric that measures the difference between two attributions, where a smaller value indicates that two attributions are more similar and represent closer meaningful interpretations.  $T$  is a function with respect to the threshold  $\varepsilon$  and  $\mathbf{x}$ . The definition formalizes the guarantee of attribution robustness, where when the model is more robust, the model attributions being attacked are less likely to be misled. More precisely, when being attacked, the change of attribution is bounded above and the smaller upper bound indicates the more attributional robust model.

Based on the above definition, the general upper bound for attribution deviations  $T(\varepsilon; \mathbf{x})$  can be found by solving the following optimization problem

$$\begin{aligned}
 & \max_{\delta} && D(g^y(\mathbf{x}), g^y(\mathbf{x} + \delta)) \\
 & \text{s.t.} && \|\delta\|_p \leq \varepsilon \\
 & && \arg \max_k f_k(\mathbf{x}) = \arg \max_k f_k(\mathbf{x} + \delta)
 \end{aligned} \tag{1}$$

We refer the first constraint  $\|\delta\|_p \leq \varepsilon$  to the norm constraint and the second one to the label constraint as it requires the unchanged label after being perturbed. In following sections, we attempt to solve the optimization problem (1) using the two mostly used norm constraints on the perturbations,  $\ell_2$  and  $\ell_\infty$ , *i.e.*,  $\|\delta\|_2 \leq \varepsilon$  and  $\|\delta\|_\infty \leq \varepsilon$ . For the dissimilarity metric, we choose from previously used attribution measurements, the Euclidean distance and the cosine distance. An alternative formulation of this problem is to find the maximum perturbation  $\varepsilon$  subject to  $D(g^y(\mathbf{x}), g^y(\mathbf{x} + \delta)) \leq \omega$  where  $\omega$  is a predefined threshold. The results obtained from problem (1) can be converted directly to the alternative formulation, and we defer the procedures to Appendix E.

## 4. A practical upper bound for attribution deviations

### 4.1. Upper bound under $\ell_2$ -norm constraint without the label constraint

To start with, the upper bound without label constraints is studied. This will provide a looser bound since, intuitively, stronger adversaries are allowed to perturb the samples that may change the classification results. While perturbations are restricted in a small region and indistinguishable to human, they could be enough to cause huge difference in attributions. The upper bound for  $\ell_2$ -norm constrained case is a straightforward derivation of the first-order Taylor series of attribution functions. The following theorem provides a tight bound for attribution robustness assuming that the attribution function is locally linear.

**Theorem 1.** Given a twice-differentiable classifier  $f : \mathbb{R}^d \rightarrow \mathbb{R}^k$ , and its attribution  $g^y$  on label  $y$ , assume that  $g^y$  is locally linear within the neighborhood of  $\mathbf{x}$ ,  $\mathcal{B}_\varepsilon(\mathbf{x}) = \{\mathbf{x} + \delta \mid \|\delta\|_2 \leq \varepsilon\}$ , then for all perturbations  $\|\delta\|_2 \leq \varepsilon$ ,

$$\|g^y(\mathbf{x} + \delta) - g^y(\mathbf{x})\|_2 \leq \xi_{max} \varepsilon,$$

where  $\xi_{max}$  is the largest singular value of  $H = \nabla g^y(\mathbf{x})$ .

*Proof.* Based on the Taylor series of  $g^y(\mathbf{x})$  and the above condition, we have

$$\|g^y(\mathbf{x} + \boldsymbol{\delta}) - g^y(\mathbf{x})\|_2^2 \leq \|\boldsymbol{\delta}^\top \nabla g^y(\mathbf{x})\|_2^2 \quad (2)$$

$$= \boldsymbol{\delta}^\top \nabla g^y(\mathbf{x}) \nabla g^y(\mathbf{x})^\top \boldsymbol{\delta} \quad (3)$$

$$= \frac{\boldsymbol{\delta}^\top}{\|\boldsymbol{\delta}\|_2} P \frac{\boldsymbol{\delta}}{\|\boldsymbol{\delta}\|_2} \cdot \|\boldsymbol{\delta}\|_2^2 \quad (4)$$

$$\leq \lambda_{max} \|\boldsymbol{\delta}\|_2^2 \leq \lambda_{max} \varepsilon^2 \quad (5)$$

where  $\lambda_{max}$  is the largest eigenvalue of  $P = HH^\top = \nabla g^y(\mathbf{x}) \nabla g^y(\mathbf{x})^\top$ , and  $\mathbf{v}_{max}$  is the corresponding eigenvector. The equality in Eq. (5) is achieved when  $\boldsymbol{\delta}$  is  $\varepsilon \mathbf{v}_{max}$  or  $-\varepsilon \mathbf{v}_{max}$ . Since the singular values of  $H$  are equal to the square root of the eigenvalues of  $P$ , then,

$$\|g^y(\mathbf{x} + \boldsymbol{\delta}) - g^y(\mathbf{x})\|_2 \leq \sqrt{\lambda_{max}} \varepsilon = \xi_{max} \varepsilon. \quad (6)$$

□

Note that the local linearity of attribution function is a weak assumption for both attribution and adversarial robust models since most of the defense methods [22, 33] attempt to smoothen the functions. In addition, when the magnitude of perturbation  $\boldsymbol{\delta}$  is constrained to small size, the magnitude of the higher-order Taylor remainders is negligible. We include the empirical results evaluating this assumption in Appendix B. Furthermore, we also provide a generalization of the theorem that bounds the attribution differences as a function of a constant  $c \geq 1$  that measures the error margin of the first-order Taylor series in Appendix B.2, which can be applied similarly to all other results in this work.

It is also noticed that the above theorem uses the gradient of attribution  $H = \nabla g^y(\mathbf{x})$ , which is also the Hessian matrix  $\nabla^2 f_y(\mathbf{x})$  when the attribution is chosen as saliency maps and can be computed easily for other gradient-based attribution methods. Moreover, the second-order derivatives can be zeros for ReLU networks [6]. In this work, the non-linearity functions are replaced by softplus function  $f(\mathbf{x}; \beta) = \frac{1}{\beta} \log(1 + e^{\beta \mathbf{x}})$  as in Dombrowski *et al.* [6]. A 2D example of the upper bound is illustrated in Fig. 1a. The optimum solution is in the same direction as the semi-major axis of the ellipse, which represents  $\boldsymbol{\delta}^\top P \boldsymbol{\delta}$ . The circle represents the 2D Euclidean ball bounded by  $T(\varepsilon; \mathbf{x})$ , which is derived from the length of the semi-major axis.

#### 4.2. Upper bound under $\ell_\infty$ -norm constraint without the label constraint

The upper bound for  $\ell_\infty$ -norm constrained case is more complicated as  $\|\boldsymbol{\delta}\|_\infty \leq \varepsilon$  defines a box constraints inequality system that  $-\varepsilon \leq \delta_i \leq \varepsilon$  for all  $i$ . If we still consider the quadratic form derived from the first-order Taylor series as in Sec. 4.1, the above optimization problem (1) turns into a concave quadratic programming with box constraints,

which is NP-hard [21]. In order to compute the upper bound efficiently, we consider a loose relaxation of  $p$ -norms.

**Corollary 1.** *Given a twice-differentiable classifier  $f : \mathbb{R}^d \rightarrow \mathbb{R}^k$ , and its attribution  $g^y$  on label  $y$ , assume that  $g^y$  is locally linear within the neighborhood of  $\mathbf{x}$ ,  $\mathcal{B}_\varepsilon(\mathbf{x}) = \{\mathbf{x} + \boldsymbol{\delta} \mid \|\boldsymbol{\delta}\|_p \leq \varepsilon\}$ , then for all perturbations  $\|\boldsymbol{\delta}\|_p \leq \varepsilon$  that  $p > 2$ ,  $\|g^y(\mathbf{x} + \boldsymbol{\delta}) - g^y(\mathbf{x})\|_2 \leq d^{\frac{1}{2} - \frac{1}{p}} \xi_{max} \varepsilon$ , where  $\xi_{max}$  is the largest singular value of  $H = \nabla g^y(\mathbf{x})$ .*

The proof of the relaxation of  $p$ -norm and Corollary 1 can be found in Appendix A.1. Note that this corollary not only avoids the NP-hard problem for  $\ell_\infty$ -norm constraint, but it is also a general upper bound for  $p$ -norm constraint on  $\boldsymbol{\delta}$  for all  $p > 2$ . However, it is also noticed that the upper bound increases with respect to the input sample dimension. The multiplication factor for  $\ell_\infty$  is  $\sqrt{d}$ . For high-dimensional input samples, the provided method would scale up to an extremely loose upper bound that can be trivial but meaningless. To better bound the attribution deviations in the  $\ell_\infty$ -norm case, we provide a tighter upper bound using the sparsity of attribution gradients.

**Theorem 2.** *Given a twice-differentiable classifier  $f$ , its attribution on label  $y$ ,  $g^y$ , and the gradient  $H = \nabla g^y$ , assume that  $g^y$  is locally linear within the neighborhood of  $\mathbf{x}$ ,  $\mathcal{B}_\varepsilon(\mathbf{x}) = \{\mathbf{x} + \boldsymbol{\delta} \mid \|\boldsymbol{\delta}\|_\infty \leq \varepsilon\}$ , then for all perturbations  $\|\boldsymbol{\delta}\|_\infty \leq \varepsilon$ ,*

$$\|g^y(\mathbf{x} + \boldsymbol{\delta}) - g^y(\mathbf{x})\|_2 \leq \varepsilon \sqrt{\sum_{i,j} |P_{ij}|}. \quad (7)$$

where  $P = HH^\top$  and the equality is taken at  $\boldsymbol{\delta} = (\pm\varepsilon, \dots, \pm\varepsilon)^\top$ .

The proof is deferred to Appendix A.2. This upper bound as the summation of absolute values of the matrix  $P = \nabla g^y(\mathbf{x}) \nabla g^y(\mathbf{x})^\top$  is shown to be tighter than that given in Corollary 1 since  $P$  is a diagonal-dominated and positive semi-definite matrix (see Fig. 1b), which implies that  $|P_{ii}| \approx \lambda_i$ .

#### 4.3. Upper bound with the label constraint

In this section, we extend our study of the upper bound to the case that labels are not changed after the samples are perturbed. Here, only attribution methods satisfying the axiom of completeness are studied as the axiom provides a direct connection between attributions and model outputs, *i.e.*,  $\sum_i g_i^y(\mathbf{x}) = f_y(\mathbf{x})$ . The following proposition gives a sufficient condition to ensure that the classification result remains unchanged after the sample is perturbed.

**Proposition 1.** *Denote the gradient-based attribution satisfying the completeness axiom of  $\mathbf{x}$  on ground truth label  $y$  by  $g^y(\mathbf{x})$ , and the attribution on a different label*

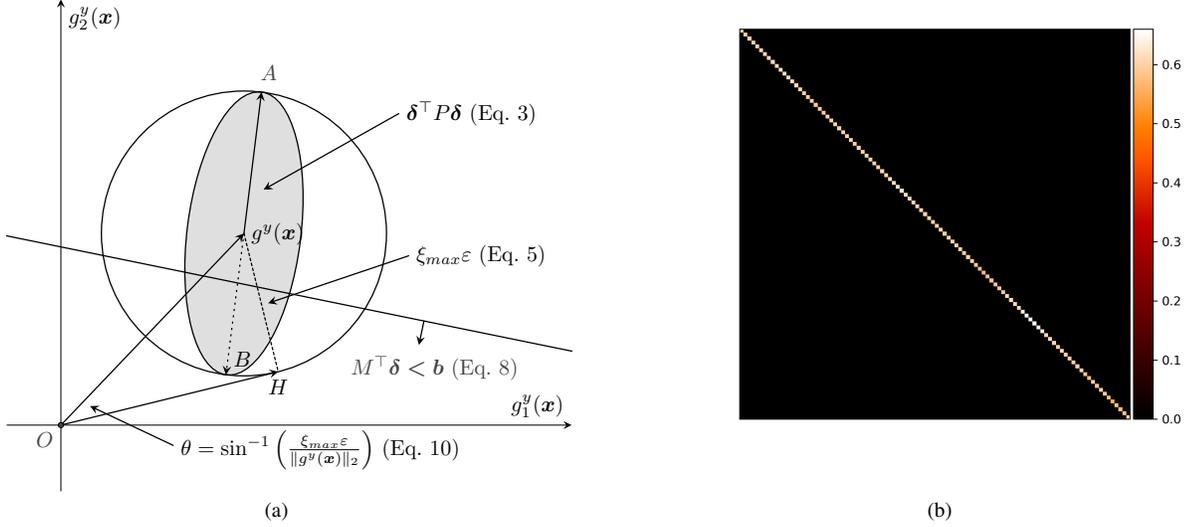


Figure 1. (a) 2D illustration of the proposed upper bound on Euclidean distance and cosine distance. (b) Visualization of the absolute values of gradient IG as a heat map. The gradient is generated using CIFAR-10 (3072  $\times$  3072), and the values are normalized to  $[0, 1]$ . Here the first 100 dimensions of each axis are plotted for better visualization. More figures and the mathematical analysis are given in Appendix C.

$y'$  by  $g^{y'}(\mathbf{x})$ . Given the perturbation  $\delta$ , assume that  $g^y$  is locally linear within the neighborhood of  $\mathbf{x}$ ,  $\mathcal{B}_\varepsilon(\mathbf{x}) = \{\mathbf{x} + \delta \mid \|\delta\|_p \leq \varepsilon\}$ , the classification result of  $\mathbf{x} + \delta$  does not change from  $y$  to  $y'$  if

$$\left( (\nabla g^{y'}(\mathbf{x}) - \nabla g^y(\mathbf{x})) \Delta \right)^\top \delta < f_y(\mathbf{x}) - f_{y'}(\mathbf{x}), \quad (8)$$

where  $\Delta$  is an all one vector,  $\Delta = (1, \dots, 1)^\top \in \mathbb{R}^d$ .

The full proof can be found in Appendix A.3. Note that the inequality is linear to  $\delta$  and we denote  $M = (\nabla g^{y'}(\mathbf{x}) - \nabla g^y(\mathbf{x}))\Delta$  and  $b = f_y(\mathbf{x}) - f_{y'}(\mathbf{x})$  for simplicity, i.e.,  $M^\top \delta < b$ . To bound the attribution differences after the sample is perturbed by noise  $\delta$  in  $\ell_2$ -norm ball, i.e.,  $\|\delta\|_2 \leq \varepsilon$ , the upper bound can be formulated by rewriting the optimization problem (1) as the optimal value of the following quadratic programming with concave objective function and a system of linear constraints for all labels different from  $y$ ,

$$\max_{\delta} \delta^\top P \delta \quad \text{s.t. } \|\delta\|_2 \leq \varepsilon \text{ and } M^\top \delta < b. \quad (9)$$

To simplify the computation, in this work, we only consider the second best label  $y'$ , i.e.,  $y' = \arg \max_{k \in \{1, \dots, c\} \setminus y} f_k(\mathbf{x})$ . In such case, the constraint  $M^\top \delta < b$  defines a half-space. Recall that Theorem 1 states that the upper bound in  $\ell_2$ -norm case without label constraint is  $\xi_{max}\varepsilon$ , and we noticed that this bound is achieved at two opposite vectors  $\delta^* = \varepsilon \mathbf{v}_{max}$  or  $\delta^* = -\varepsilon \mathbf{v}_{max}$ . Thus, at least one of these two vectors lies

in the half-space defined by the linear constraint (see point  $A$  and  $B$  in Fig. 1a). Therefore, the upper bound provided in Theorem 1 is also achieved even if the label constraint is added, i.e., the optimal value of optimization problem (9) is also  $\xi_{max}\varepsilon$ .

The bound of  $\ell_p$ -norm constrained case can be derived similarly. The less tight upper bound is still achievable at  $d^{\frac{1}{2} - \frac{1}{p}} \xi_{max}\varepsilon$  as in Corollary 1. Generalizing the  $\ell_\infty$ -norm constrained upper bound using Eq. 7 is simpler. According to Theorem 2, there are  $2^d$  different optimal solutions that achieve the optimum and they are the corners of the  $\ell_\infty$ -norm box. As long as the feasible region is non-empty, there exists at least one corner of the box lying inside the feasible region, and the optimum value is achieved. Similarly, given a  $d$ -dimensional  $\ell_\infty$ -norm box, and  $k$  label constraints that each separates the entire space into two half-spaces, if at least one corner of the box lies within the feasible region, the optimum value is attainable.

#### 4.4. Upper bound based on cosine distance

In the previous parts of this section, we discussed several practical upper bounds based on Euclidean distance. It is noticed that Wang & Kong [32] claims that cosine similarity ( $D_s$ ) is a better metric to measure the difference of attributions as it emphasizes the relative importance among different features rather than the absolute magnitude of each individual feature. Our method can be trivially extended to the scenarios using cosine distance ( $D_c = 1 - D_s(g^y(\mathbf{x} + \delta), g^y(\mathbf{x}))$ ) as the dissimilarity function  $D$

defined in the formulation (1) with simple modifications.

**Corollary 2.** *Given a twice-differentiable classifier  $f : \mathbb{R}^d \rightarrow \mathbb{R}^k$  and its attribution  $g^y$  on label  $y$ , for all perturbations  $\|\delta\|_p \leq \varepsilon$ , if the Euclidean distance of  $g^y(\mathbf{x} + \delta)$  and  $g^y(\mathbf{x})$  is upper bounded by  $T(\varepsilon; \mathbf{x})$ , and  $0 \leq T(\varepsilon; \mathbf{x}) \leq \|g^y(\mathbf{x})\|_2$ , then their cosine distance ( $D_c$ ) is upper bounded by*

$$D_c(g^y(\mathbf{x} + \delta), g^y(\mathbf{x})) \leq 1 - \sqrt{1 - \frac{T(\varepsilon; \mathbf{x})^2}{\|g^y(\mathbf{x})\|_2^2}}. \quad (10)$$

This upper bound is valid when the assumption that  $0 \leq T(\varepsilon; \mathbf{x}) \leq \|g^y(\mathbf{x})\|_2$  is satisfied, *i.e.*, the variation of attribution distance is smaller than the original attribution. As shown in Fig. 1a, the angle between original and perturbed attributions is bounded by  $\theta$  computed from the corollary.

## 5. Experimental results

In this section, we evaluate the effectiveness of proposed upper bounds by numerical experiments under both  $\ell_2$  and  $\ell_\infty$ -norms. In the following results, we compute the theoretical upper bounds for adversarial robust models and attributional robust models, including *Adversarial Training* (AT) [19], *IG-NORM* [4], *Adversarial Attributional Training* with robust training loss (AdvAAT) [12], *Attributional Robustness Training* (ART) [30], *TRADES* [37] and *Integrated Gradients Regularizer* (IGR) [32]. We follow previous attribution robustness studies to use ResNet-18 to evaluate CIFAR-10 [15], and use the neural network with four convolutional layers followed by three fully-connected layers to evaluate MNIST [17] and Fashion-MNIST [34]. The proposed upper bounds are also scalable to datasets with larger images, but due to the scalability problem of existing attribution defense methods, we provide evaluations of Flower [20] and ImageNet [5] on attributional non-robust models in Appendix D.3.

For each selected model, the theoretical upper bounds for both Euclidean distance and cosine distance are computed. We convert the cosine values to degrees for easier comparison. The theoretical bounds are compared with corresponding distance between original sample and attacked sample to verify the effectiveness of the bounds. We denote the theoretical upper bounds for Euclidean distance and cosine distance as  $T_e$  and  $T_c$ , respectively. The  $\ell_2$  PGD-20 attack [19] is implemented for  $\ell_2$ -norm bounded case. The 200-step IFIA with the top- $k$  intersection as dissimilarity function [8] is implemented for  $\ell_\infty$ -norm bounded case, where  $k$  is 100 for MNIST and Fashion-MNIST and 1000 for CIFAR-10. Each sample is attacked 20 times and the mean distance is computed. The sample mean Euclidean and cosine distances of the entire dataset under corresponding attacks are

denoted by  $\widehat{T}_e$  and  $\widehat{T}_c$ , respectively. All the experiments are implemented on NVIDIA GeForce RTX 3090.

In addition, we also provide a generalization of the proposed bounds based on the generalization of Theorem 1 (Appendix B.2) that adaptively multiply a scalar  $c$  for an given input  $\mathbf{x}$  in case that the weak assumption is violated in rare cases. Explicitly, the adaptive value of  $c$  for  $i$ -th sample is given as follows (details in Appendix B.3)

$$c^{(i)} = \max \left\{ 1, \frac{\|g^y(\mathbf{x}^{(i)} + \varepsilon \mathbf{v}_{max}^{(i)}) - g^y(\mathbf{x}^{(i)})\|_2}{\xi_{max}^{(i)} \varepsilon} \right\}. \quad (11)$$

### 5.1. Evaluation of upper bounds without the label constraint

We first evaluate the upper bounds without label constraints, which can be applied to any gradient-based attribution method. Here three methods are evaluated, saliency map, input\*gradient and integrated gradients. The upper bounds computed from TRADES+IGR on CIFAR-10 are presented in Table 2, and results from other models are deferred to Appendix D.1. We use the unlabelled upper bound introduced in Theorem 1 and 2 to compute  $T_e = \xi_{max} \varepsilon$  and extend it to  $T_c$  using Eq. 10. The perturbation size is chosen to be 0.1 for  $\ell_2$  and 0.25 for  $\ell_\infty$ . The generalized upper bound  $T'_e = c \xi_{max} \varepsilon$  (Eq. 11) is also provided for  $\ell_2$  case, and is not necessary for  $\ell_\infty$  case. As observed in the table, the proposed upper bounds are valid for different attribution methods and both Euclidean and cosine distances are well-bounded. More precisely, none of the  $\ell_\infty$  perturbed attributions is outside the theoretical bound and none of the  $\ell_2$  perturbed attributions is outside the generalized bound.

### 5.2. Evaluation of upper bounds under $\ell_2$ -norm constraint and label constraint

To evaluate the upper bound of attribution deviations after samples being attacked by  $\ell_2$ -norm constrained perturbations, we use the method provided in Sec. 4.3 and 4.4 to obtain the theoretical upper bounds for both Euclidean distance ( $T_e = \xi_{max} \varepsilon$ ) and cosine distance ( $T_c$  in Eq. 10). Integrated gradients (IG) is chosen here as the theoretical bound with the label constraint is based on the axiom of completeness. Besides, as discussed in Sec. 4.1, the percentages of attacked attribution outside  $T_e$  are provided, and the generalized bound is also calculated and denoted by  $T'_e = c \xi_{max} \varepsilon$ . Since  $T'_e$  bounds all the attacked attributions, *i.e.*, 100% for all the models, we do not report the percentages in the table. From Table 3, we observe the following results. (i) The percentages are low, which supports our assumption in Sec. 4.1 that  $g^y$  is locally linear. (ii) The computed values for both Euclidean ( $T'_e$ ) and cosine distance ( $T_c$ ) successfully bound the attribution differences above for every dataset and every model. In addition, we

Table 2. Evaluation of upper bound without the label constraint.

	SM					Input*gradient					IG				
	$\widehat{T}_e$	$T_e$	$T'_e$	$\widehat{T}_c$	$T_c$	$\widehat{T}_e$	$T_e$	$T'_e$	$\widehat{T}_c$	$T_c$	$\widehat{T}_e$	$T_e$	$T'_e$	$\widehat{T}_c$	$T_c$
$\ell_2$	0.09	0.31	0.34	6.88	7.41	0.07	0.46	0.46	0.51	2.60	0.02	0.17	0.17	1.80	3.84
$\ell_\infty$	0.41	0.85	-	21.87	27.09	0.07	0.69	-	7.03	50.59	0.25	0.52	-	23.24	35.00

Table 3. Evaluation of upper bounds under  $\ell_2$ -norm constraint and label constraint. The numbers in the brackets indicate the percentages that attacked attribution is outside the  $T_e$ .

Model	$\widehat{T}_e$	$T_e$	$T'_e$	$\widehat{T}_c(\text{deg})$	$T_c(\text{deg})$
$\varepsilon = 0.05$ MNIST					
AT	0.0685	0.1537 [2.25%]	0.1596	3.6935	7.0951
IG-NORM	0.1158	0.2888 [2.00%]	0.2967	3.3174	7.2615
ART	0.0626	0.3591 [6.00%]	0.3702	2.3923	6.8657
AdvAAT	0.0876	0.3269 [6.20%]	0.3404	1.9034	6.8992
TRADES	0.1620	0.5060 [1.68%]	0.5271	2.8374	6.9988
TRADES+IGR	0.1784	0.4964 [1.32%]	0.5145	2.9075	6.9779
$\varepsilon = 0.05$ Fashion-MNIST					
AT	0.0659	0.0700 [2.19%]	0.0869	10.1442	12.9577
IG-NORM	0.1181	0.1789 [0.00%]	0.1789	6.6002	8.8043
AdvAAT	0.1115	0.1735 [6.20%]	0.1858	5.8544	9.3692
ART	0.0940	0.1387 [0.94%]	0.1411	5.4507	9.8234
TRADES	0.0626	0.0963 [4.16%]	0.1184	8.3521	12.0991
TRADES+IGR	0.0403	0.0453 [1.91%]	0.0507	7.7302	8.8411
$\varepsilon = 0.1$ CIFAR-10					
AT	0.0392	0.2532 [0.09%]	0.2533	2.7335	4.7724
IG-NORM	0.0149	0.1582 [0.42%]	0.1621	1.6505	4.3711
AdvAAT	0.0374	0.2386 [0.06%]	0.2386	0.2847	3.8202
ART	0.0733	0.2278 [0.00%]	0.2278	0.5918	4.2123
TRADES	0.0264	0.1734 [0.16%]	0.1734	1.9084	3.8686
TRADES+IGR	0.0240	0.1692 [0.09%]	0.1692	1.8011	3.8384

also show the minimum Euclidean gaps between samples and bounds in Section 5.4 to illustrate that the tightness of the proposed bounds.

### 5.3. Evaluation of upper bounds under $\ell_\infty$ -norm constraint and label constraint

In this subsection, the upper bounds of attribution deviations under  $\ell_\infty$  attacks are presented. Instead of the much looser bound derived from  $\ell_p$ -norm relaxation (Corollary 1), the upper bound is computed from  $T_e = \varepsilon \sqrt{\sum |P_{ij}|}$  as introduced in Sec. 4.3 and 4.4. Moreover, the empirical attribution robustness is also provided using Kendall’s rank correlation [14] for comparison of theoretical and empirical protection. It should be emphasized that all the previous attribution robustness studies are based on  $\ell_\infty$ -norm constraints. Kendall’s rank correlation, which is non-differentiable, is used as one of the indexes to measure the difference between the original attributions and the attributions attacked by IFIA under  $\ell_\infty$ -norm constraints. From the results in Table 4, we see that the computed theoretical upper bounds are valid to measure the worst-case at-

Table 4. Evaluation of upper bounds  $\ell_\infty$ -norm constraint and label constraint.

Model	Kendall	$\widehat{T}_e$	$T_e$	$\widehat{T}_c(\text{deg})$	$T_c(\text{deg})$
$\varepsilon = 0.05$ MNIST					
AT	0.1846	0.3461	0.7752	15.3616	38.7024
IG-NORM	0.1562	0.6836	1.2046	16.4506	31.8791
AdvAAT	0.3791	1.6269	2.1992	11.3173	26.3000
ART	0.1439	1.4193	2.8218	12.8025	64.3115
TRADES	0.2127	1.1779	2.2216	15.9881	33.4681
TRADES+IGR	0.4537	1.2991	2.0386	17.5923	26.5748
$\varepsilon = 0.05$ Fashion-MNIST					
AT	0.1516	0.0990	0.2802	18.9720	55.1501
IG-NORM	0.3446	0.2384	0.8819	12.6023	46.4270
AdvAAT	0.5810	0.1938	0.9206	9.4499	44.9184
ART	0.2079	0.1660	0.7215	9.6582	53.7281
TRADES	0.2582	0.1042	0.4536	13.7010	51.6448
TRADES+IGR	0.6565	0.0722	0.2526	15.0947	44.4703
$\varepsilon = 0.1$ CIFAR-10					
AT	0.5578	0.4058	0.7649	26.6195	45.3223
IG-NORM	0.5811	0.1997	0.4783	21.6311	35.2981
AdvAAT	0.5484	0.2293	0.5211	28.7342	39.3981
ART	0.6875	0.3128	0.6734	31.0090	35.6422
TRADES	0.6903	0.2322	0.5001	22.9779	36.3759
TRADES+IGR	0.6940	0.2474	0.5236	23.2356	35.0009

tribution deviations. For each dataset and each model, the sample mean of attribution distances is strictly smaller than the theoretical distance. Because of the relaxation, all attacked attributions are bounded by  $T_e$ . There is no outlier, and there is no need to use the generalized bound  $T'_e$ . Moreover, the results also show that for a model with a larger Kendall’s rank correlation, the theoretical cosine distance upper bound is more likely to be smaller (see Fig. 2), which means that the model is more difficult to be attacked. This also confirms that cosine similarity is positively correlated to Kendall’s rank correlation as proposed in Wang & Kong [32].

### 5.4. Evaluation of the tightness of bounds

In addition, we further report the minimum Euclidean gaps between samples and theoretical bounds in Table 5 to measure the tightness of the provided bounds, which is defined as

$$r = \min_{0 \leq i \leq n} T_e^{(i)} - \widehat{T}_e^{(i)} \quad (12)$$

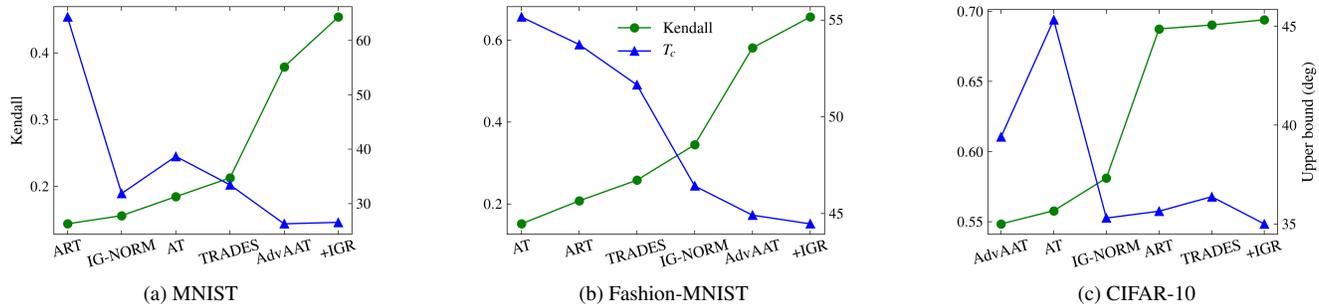


Figure 2. Comparison between Kendall’s rank correlation and the theoretical bound of cosine distance. For clear comparison, we convert the cosine distance to angles in degrees.

Table 5. Evaluation of tightness of the bounds in Euclidean distance for  $\ell_2$  and  $\ell_\infty$  cases.

$\ell_2(\varepsilon =)$	MNIST			Fashion-MNIST			CIFAR-10		
	0.05	0.1	0.2	0.05	0.1	0.2	0.1	0.2	0.3
AT	0.0260	0.0320	0.0140	0.0078	0.0130	0.1207	0.0004	0.0045	0.0134
IG-NORM	0.0391	0.0597	0.0291	0.0121	0.0171	0.0742	0.0016	0.0210	0.0448
AdvAAT	0.0290	0.0465	0.0294	0.0103	0.0167	0.0183	0.0037	0.0090	0.0181
ART	0.0178	0.0115	0.0182	0.0029	0.0239	0.0011	0.0019	0.0031	0.0141
TRADES	0.0014	0.0032	0.0104	0.0037	0.0082	0.0723	0.0028	0.0048	0.0147
TRADES+IGR	0.0010	0.0038	0.0041	0.0064	0.0134	0.0385	0.0016	0.0100	0.0126
$\ell_\infty(\varepsilon =)$	0.01	0.03	0.05	0.01	0.03	0.05	4/255	8/255	0.1
AT	0.0021	0.0035	0.0117	0.0004	0.0352	0.0708	0.0025	0.0053	0.1381
IG-NORM	0.0001	0.0062	0.0105	0.0010	0.0590	0.1069	0.0026	0.0145	0.0003
AdvAAT	0.0004	0.0223	0.0901	0.0136	0.0847	0.1665	0.0448	0.1513	0.2078
ART	0.0082	0.0112	0.0233	0.0424	0.1049	0.1467	0.0118	0.0412	0.0870
TRADES	0.0001	0.0046	0.0026	0.0014	0.0337	0.0634	0.0068	0.0043	0.0530
TRADES+IGR	0.0016	0.0143	0.1389	0.0014	0.0375	0.0682	0.0016	0.0090	0.0197

Note that the superscript  $(i)$  represents the  $i$ -th sample and  $T_e^{(i)}$  is replaced by  $T_e^{(i)}$  in  $\ell_2$ -norm cases.  $r$  is a straightforward measurement of the theoretical bound. We notice that although the mean of theoretical bounds sometimes are multiple times larger than the sample mean distance, the tightest bound can be only  $10^{-4}$  greater than the sample distance. We also observe that the values of  $r$  are all positive, which also indicates that there is no perturbed attribution that violates our theoretical bounds.

We also provide the visualizations of the distribution of the gap between theoretical bounds and attribution differences from real data in Appendix D. The values are directly computed using  $T_e^{(i)} - \hat{T}_e^{(i)}$ . As we can observe from the figures, all values are positive, which verifies the validity of our bounds, and most of the values are lying close to 0, which shows the tightness of the bounds.

## 6. Conclusion

In this paper, an upper bound that measures the worst-case attribution deviations is proposed. The bound is for-

mulated as the optimum value of a constrained optimization problem. The optimization problem is constrained on the size of perturbation and the unchanged classification label after being perturbed. For each of the two metrics of the attribution difference, Euclidean and cosine distances, the problem is solved based on the first-order Taylor series and the estimation of attribution gradients and the solution bounds the attribution deviations under both  $\ell_2$  and  $\ell_\infty$ -norm attacks. Experimental results validate the effectiveness of the proposed bounds.

## Acknowledgments and Disclosure of Funding

This work is partially supported by the Ministry of Education, Singapore through Academic Research Fund Tier 1, RG73/21.

## References

- [1] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classi-

- fier decisions by layer-wise relevance propagation. *PLoS one*, 10(7):e0130140, 2015.
- [2] Umang Bhatt, Alice Xiang, Shubham Sharma, Adrian Weller, Ankur Taly, Yunhan Jia, Joydeep Ghosh, Ruchir Puri, José MF Moura, and Peter Eckersley. Explainable machine learning in deployment. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 648–657, 2020.
  - [3] Akhilan Boopathy, Sijia Liu, Gaoyuan Zhang, Cynthia Liu, Pin-Yu Chen, Shiyu Chang, and Luca Daniel. Proper network interpretability helps adversarial robustness in classification. In *International Conference on Machine Learning*, pages 1014–1023. PMLR, 2020.
  - [4] Jiefeng Chen, Xi Wu, Vaibhav Rastogi, Yingyu Liang, and Somesh Jha. Robust attribution regularization. In *Advances in Neural Information Processing Systems*, pages 14300–14310, 2019.
  - [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
  - [6] Ann-Kathrin Dombrowski, Maximillian Alber, Christopher Anders, Marcel Ackermann, Klaus-Robert Müller, and Pan Kessel. Explanations can be manipulated and geometry is to blame. In *Advances in Neural Information Processing Systems*, pages 13589–13600, 2019.
  - [7] Chris Finlay and Adam M Oberman. Scaleable input gradient regularization for adversarial robustness. *arXiv preprint arXiv:1905.11468*, 2019.
  - [8] Amirata Ghorbani, Abubakar Abid, and James Zou. Interpretation of neural networks is fragile. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3681–3688, 2019.
  - [9] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015.
  - [10] Bryce Goodman and Seth Flaxman. European union regulations on algorithmic decision-making and a “right to explanation”. *AI magazine*, 38(3):50–57, 2017.
  - [11] Hongyu Guo, Yongyi Mao, and Richong Zhang. Mixup as locally linear out-of-manifold regularization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3714–3722, 2019.
  - [12] Adam Ivankay, Ivan Girardi, Chiara Marchiori, and Pascal Frossard. FAR: A general framework for attributional robustness. In *32nd British Machine Vision Conference, BMVC 2021, Online*, page 24. BMVA Press, 2021.
  - [13] José Jiménez-Luna, Francesca Grisoni, and Gisbert Schneider. Drug discovery with explainable artificial intelligence. *Nature Machine Intelligence*, 2(10):573–584, 2020.
  - [14] Maurice George Kendall. Rank correlation methods. 1948.
  - [15] Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009.
  - [16] Cassidy Laidlaw, Sahil Singla, and Soheil Feizi. Perceptual adversarial robustness: Defense against unseen threat models. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
  - [17] Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.
  - [18] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017.
  - [19] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.
  - [20] M-E Nilsback and Andrew Zisserman. A visual vocabulary for flower classification. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1447–1454. IEEE, 2006.
  - [21] Panos M Pardalos and Stephen A Vavasis. Quadratic programming with one negative eigenvalue is np-hard. *Journal of Global optimization*, 1(1):15–22, 1991.
  - [22] Chongli Qin, James Martens, Sven Gowal, Dilip Krishnan, Krishnamurthy Dvijotham, Alhussein Fawzi, Soham De, Robert Stanforth, and Pushmeet Kohli. Adversarial robustness through local linearization. *Advances in Neural Information Processing Systems*, 32, 2019.
  - [23] Anindya Sarkar, Anirban Sarkar, and Vineeth N Balasubramanian. Enhanced regularizers for attributional robustness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 2532–2540, 2021.
  - [24] Rory Sayres, Ankur Taly, Ehsan Rahimy, Katy Blumer, David Coz, Naama Hammel, Jonathan Krause, Arunachalam Narayanaswamy, Zahra Rastegar, Derek Wu, et al. Using a deep learning algorithm and integrated gradients explanation to assist grading for diabetic retinopathy. *Ophthalmology*, 126(4):552–564, 2019.
  - [25] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. In *International Conference on Machine Learning*, pages 3145–3153. PMLR, 2017.
  - [26] Avanti Shrikumar, Peyton Greenside, Anna Shcherbina, and Anshul Kundaje. Not just a black box: Learning important features through propagating activation differences. *arXiv preprint arXiv:1605.01713*, 2016.
  - [27] Carl-Johann Simon-Gabriel, Yann Ollivier, Leon Bottou, Bernhard Schölkopf, and David Lopez-Paz. First-order adversarial vulnerability of neural networks and input dimension. In *International Conference on Machine Learning*, pages 5809–5817. PMLR, 2019.
  - [28] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Workshop Track Proceedings*, 2014.
  - [29] Amitojdeep Singh, Sourya Sengupta, and Vasudevan Lakshminarayanan. Explainable deep learning models in medical image analysis. *Journal of Imaging*, 6(6):52, 2020.
  - [30] Mayank Singh, Nupur Kumari, Puneet Mangla, Abhishek Sinha, Vineeth N Balasubramanian, and Balaji Krishnamurthy. Attributional robustness training using input-gradient spatial alignment. In *Computer Vision—ECCV 2020*:

*16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVII 16*, pages 515–533. Springer, 2020.

- [31] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *International Conference on Machine Learning*, pages 3319–3328. PMLR, 2017.
- [32] Fan Wang and Adams Wai-Kin Kong. Exploiting the relationship between kendall’s rank correlation and cosine similarity for attribution protection. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- [33] Zifan Wang, Haofan Wang, Shakul Ramkumar, Piotr Mardziel, Matt Fredrikson, and Anupam Datta. Smoothed geometry for robust attribution. In *Advances in Neural Information Processing Systems*, volume 33, pages 13623–13634. Curran Associates, Inc., 2020.
- [34] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- [35] Yao-Yuan Yang, Cyrus Rashtchian, Hongyang Zhang, Russ R Salakhutdinov, and Kamalika Chaudhuri. A closer look at accuracy vs. robustness. *Advances in neural information processing systems*, 33:8588–8601, 2020.
- [36] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014.
- [37] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. Theoretically principled trade-off between robustness and accuracy. In *International Conference on Machine Learning*, pages 7472–7482. PMLR, 2019.
- [38] Linjun Zhang, Zhun Deng, Kenji Kawaguchi, Amirata Ghorbani, and James Zou. How does mixup help with robustness and generalization? In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [39] Luisa M. Zintgraf, Taco S. Cohen, Tameem Adel, and Max Welling. Visualizing deep neural network decisions: Prediction difference analysis. In *International Conference on Learning Representations, ICLR 2017*. OpenReview.net, 2017.