

FAR: A General Framework for Attributional Robustness

Adam Ivankay¹
aiv@zurich.ibm.com

Ivan Girardi¹
ivg@zurich.ibm.com

Chiara Marchiori¹
chi@zurich.ibm.com

Pascal Frossard²
pascal.frossard@epfl.ch

¹ IBM Research Zurich
Rüschlikon, Switzerland

² École Polytechnique Fédérale de
Lausanne (EPFL)
Lausanne, Switzerland

Abstract

Attribution maps are popular tools for explaining neural networks' predictions. By assigning an importance value to each input dimension that represents its impact towards the outcome, they give an intuitive explanation of the decision process. However, recent work has discovered vulnerability of these maps to imperceptible adversarial changes, which can prove critical in safety-relevant domains such as healthcare. Therefore, we define a novel generic *framework for attributional robustness* (FAR) as general problem formulation for training models with robust attributions. This framework consist of a generic regularization term and training objective that minimize the maximal dissimilarity of attribution maps in a local neighbourhood of the input. We show that FAR is a generalized, less constrained formulation of currently existing training methods. We then propose two new instantiations of this framework, *AAT* and *AdvAAT*, that directly optimize for both robust attributions and predictions. Experiments performed on widely used vision datasets show that our methods perform better or comparably to current ones in terms of attributional robustness while being more generally applicable. We finally show that our methods mitigate undesired dependencies between attributional robustness and some training and estimation parameters, which seem to critically affect other competitor methods.

1 Introduction

As deep neural networks (DNNs) have become larger and deeper in recent years, their complexity also increased significantly. This makes it difficult for humans to understand and reason about their decision process. Therefore, methods that provide insight and reason about the prediction outcome of DNNs are crucial for successfully deploying these complex networks in real-life scenarios. *Attribution maps* assign an importance value to each input dimension, which represents its *influence* on the outcome of the decision. While these methods are fast to compute, they do not require specific domain knowledge to give interpretable

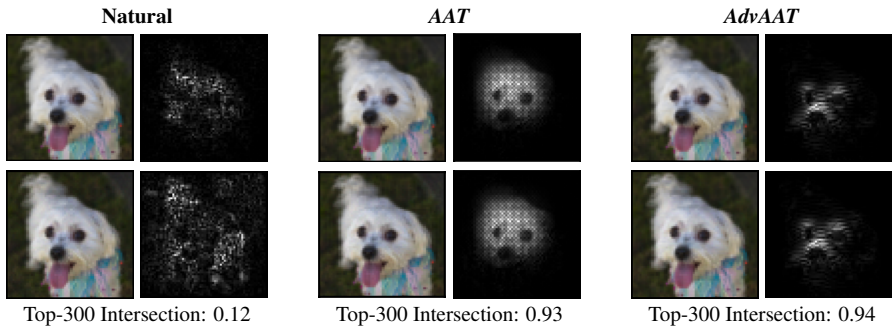


Figure 1: Original and adversarial Integrated Gradients (IG) of the natural model (left), our AAT (middle) and AdvAAT (right) method on Restricted Imagenet. For each model, the upper row contains the unperturbed image in the left column and its IG saliency map in the right column. The lower row contains the corresponding perturbed image on the left and the resulting adversarial IG saliency map on the right. Our methods yield less noisy and more robust attribution maps (measured by the *Top-300 intersection* of the highest attributed pixels of the unperturbed and perturbed image), while correctly classifying all images.

explanations - *the outcome is the given class because of the salient region in the input*, hence their popularity. Most of these methods, like Saliency [19], Integrated Gradients [23], SmoothGrad [21] or DeepLIFT [18] utilize the input gradient of the network to construct the attribution map, giving explanation on the networks behaviour in a region around the given input.

However, the authors in [5] have successfully demonstrated that these explanations can be adversarially manipulated by adding carefully crafted perturbations to the input, fundamentally changing the attributions. This is problematic in critical scenarios where the decision outcome needs to be accompanied by a sound explanation. Attackers could willingly manipulate the salient region in a digital pathology image, critically misleading the medical professional assessing the decision. In other instances, they could induce false bias in automatic credit scoring algorithms by pointing to the individuals race as reason for a low score. These are critical aspects that prevent DNNs from being adopted to solve real life problems and highlight the need for robust explanations. Figure 1 exemplifies the adversarial fragility of attribution maps.

In our work, we mitigate this fragility by making the following contributions:

- We define a *framework for attributional robustness* (FAR) as general problem formulation for training robust attributions. Key aspects of this framework are:
 - It allows for separate optimization for robust predictions and explanations,
 - It generalizes to more explanation methods and attribution distances than current methods,
 - It allows for providing ground truth explanations.
- We provide two instantiations of FAR, our AAT and AdvAAT methods that directly optimize for maximal correlation between original and adversarial feature importance within a small L_∞ neighbourhood of the input. Experiments show that our methods outperform or perform similarly to other methods on widely-used vision datasets.
- We identify undesired dependencies of gradient-based attribution maps on training and estimation parameters, which are mitigated by our methods.

2 Preliminaries

2.1 Background

Let f be a DNN classifier, $f_t(\mathbf{x})$ the output logit of f for class t , $y = \arg \max_t f_t(\mathbf{x})$ the predicted class and $g_y^{\mathbf{x}}(\mathbf{x})$ the input gradient for class y , given input \mathbf{x} . An *attribution method* S is a function that assigns a value to each element of \mathbf{x} that represents its influence towards the prediction y of DNN f .

The simplest attribution map considered in this work is *Saliency Map* (SM), defined as the element-wise absolute value of the input gradient, written as follows:

$$S(\mathbf{x}, f) = \text{SM}(\mathbf{x}, f) = |\nabla_{\mathbf{x}} f_y(\mathbf{x})| := g_{\mathbf{x}}^y(\mathbf{x})^{\text{abs}} \quad (1)$$

Integrated Gradients (IG) is an axiomatic, smoothed version of SM [23]. IG is defined as the path integral from a predefined baseline \mathbf{b} to the input \mathbf{x} written in Equation (2).

$$S(\mathbf{x}, f, \mathbf{b}) = \text{IG}(\mathbf{x}, f, \mathbf{b}) = (\mathbf{x} - \mathbf{b}) \cdot \int_{\alpha=0}^1 g_{\tilde{\mathbf{x}}}^y(\tilde{\mathbf{x}})|_{\tilde{\mathbf{x}}=\mathbf{b}+\alpha(\mathbf{x}-\mathbf{b})} d\alpha \quad (2)$$

where \mathbf{b} can be chosen as an arbitrary input signal such as $\mathbf{0}$ or noise. It is a principled approach that fulfills sensitivity, implementation invariance and completeness axioms [23], which makes it a commonly accepted and deployed explanation method.

Attributional Robustness (AR) refers to explanations' resistance towards adversarial perturbations. While there is no agreement of the definition of AR, most mathematical formulations build towards the conjecture that attribution maps should be similar for similar inputs. We define AR as written in Equation (3).

$$r(S, S^T) = 1 - \max_{\tilde{\mathbf{x}}} d_s[S(\tilde{\mathbf{x}}, f), S^T] \quad (3)$$

given the constraints in the following Equation (4),

$$\|\tilde{\mathbf{x}} - \mathbf{x}\|_p < \varepsilon \quad \text{and} \quad \arg \max_t f_t(\tilde{\mathbf{x}}) = \arg \max_t f_t(\mathbf{x}) \quad (4)$$

where $r(S, S^T)$ denotes the attributional robustness of saliency map S , S^T the target saliency map, f the classifier, $\tilde{\mathbf{x}}$ and \mathbf{x} the adversarial and original inputs, ε a small bound on the L_p -norm of the input change and d_s a (scaled) dissimilarity metric for attribution maps.

2.2 Related work

First efforts for mitigating this fragility were introduced in [4]. The authors proved theoretical connections between adversarial robustness and the *alignment* $\langle g_{\mathbf{x}}^y(\mathbf{x}), \mathbf{x} \rangle$ of prediction input gradients $g_{\mathbf{x}}^y(\mathbf{x})$ and input \mathbf{x} . Based on this consideration, it has been shown that adversarially trained networks (in their predictions) have increased attributional robustness [1, 20, 24]. On the other hand, the authors of [20] have achieved state of the art AR by maximizing alignment with a regularization term during training. However, the input-gradient alignment as an upper bound for robustness strictly holds only in a local linear neighbourhood of the input. Moreover, it depends on both the input and its gradients, as it computes the inner product between them. This is a strong assumption, as zero-input areas do not contribute to the alignment, no matter their gradient.

An axiomatic approach to achieve robust saliency maps has been introduced in [1]. The authors use Integrated Gradients (IG) [23] and its theoretical properties to attain robust attributions. In other research, the works of [2] and [24] analyzed the local geometry of the classifier, effectively reducing attribution fragility by smoothing the curvature of the classifier. This is consistent with previous observations of correlations between adversarially robust networks and robust attributions [4, 20].

While there has been significant effort in mitigating adversarial vulnerability of attribution maps, most works have crucial shortcomings. First, they jointly optimize for both adversarial and attributional robustness. This does not allow for a separate analysis of the two notions. Second, they achieve AR by regularizing curvature or gradients, therefore can not be straightforwardly defined for explanations that are not based on gradients. Third, none support defining target explanations, which is useful for scenarios described in [16] or [17]. Lastly, the method from [20] computes the inner product between input and gradients, thus coupling input data and gradient domains. This can not be straightforwardly defined for non-continuous inputs like categorical variables, text or other constraint inputs in multimodal problems.

3 Framework for Attributional Robustness

In this section, we introduce FAR, our general problem formulation of solving AR in DNNs. It consists of two generic training objectives for robust attributions and predictions. Then, we derive existing robust attribution training methods from these objectives, showcasing the general nature of our formulation.

In order to mitigate the shortcomings of current attributional robustness methods, we extend the classical notion of adversarial training [13] to attribution maps. Thus, we make the following considerations:

- *Similar attribution maps.* The main assumption of our framework is that similar inputs should give near-identical explanations, as defined in Equation (3). It is often argued that not all input features matter towards the prediction, but only a certain subset, or even only their relative ranks [5]. Therefore, we measure attribution similarity with metrics that reflect these considerations, like the *Kendall's rank order correlation* (CO) [9] or the *Top-K intersection* (IN) [5], as done in most of the related work.
- *Perceptually identical inputs and unchanged prediction outcome.* Analogously to traditional adversarial training, we utilize widely-used L_p -ball restrictions of size ϵ around the input to ensure unchanged ground truth labels of the data. Moreover, we require the same predicted class $\arg \max_t f_t(\mathbf{x}) = \arg \max_t f_t(\mathbf{x}_{\text{adv}})$ of original and perturbed inputs. The latter constraint motivates our assumption that similar inputs should have similar explanations.
- *Target attributions and identical prediction outcomes.* Generally, ground truth for attribution maps is not available, therefore, we use the attributions of the unperturbed inputs as targets. However, allowing to provide these targets could provide useful for datasets in which they are given.

3.1 Optimization Problem

Given the above points, we define our framework as a *regularization term* added to the classification loss and a *robust training loss*, which formulate generic objectives for robustifying any smooth attribution method and dissimilarity. The regularization term is used to enhance robustness of the attributions separately from the inference outcome, while the robust training loss jointly encourages adversarial and attributional robustness. Thus, the training objectives of the framework become the following *min-max* optimization problems written in Equations (5) and (6).

$$\theta^* = \arg \min_{\theta} \sum_{\mathbf{x} \in \mathcal{D}} \left\{ l(\mathbf{x}, y, f) + \lambda \cdot \max_{\|\tilde{\mathbf{x}} - \mathbf{x}\|_p < \varepsilon} d_s[S(\tilde{\mathbf{x}}, f), S^T(\mathbf{x}, f)] \right\} \quad (5)$$

or

$$\theta^* = \arg \min_{\theta} \sum_{\mathbf{x} \in \mathcal{D}} \max_{\|\tilde{\mathbf{x}} - \mathbf{x}\|_p < \varepsilon} \left\{ l(\tilde{\mathbf{x}}, y, f) + \lambda \cdot d_s[S(\tilde{\mathbf{x}}, f), S^T(\mathbf{x}, f)] \right\} \quad (6)$$

where f denotes the classifier, l the classification loss, d_s is any smooth dissimilarity metric between the saliency map S and a target saliency map S^T , p denotes an ε -bounded norm base, y the target class and θ^* the optimal parameters of the classifier f trained on dataset \mathcal{D} . λ controls the robust attribution regularization. Note that while Equation (5) allows for solely optimizing for robust attributions, the training loss described in Equation (6) encourages both robust attributions *and* predictions. The inner maximizations of Equations (5) and (6) are solved via the IFIA (Algorithm 1) and the Adversarial IFIA (Algorithm 2) algorithms respectively, written in Figure 2.

Formulating the AR problem as above has the following advantages. First, the choice of S is not fixed - the framework can be used to robustify any saliency map. Second, the domain of explanations and input data is not coupled, hence the shortcomings of current methods described in Section 2.2 do not exist for our framework. Third, the choice of S^T is not fixed, therefore target (ground truth) explanations can be provided if present, and robust explanations can be trained with respect to these. Fourth, the dissimilarity metric d_s can be chosen to any smooth d_s , depending on the use case. Lastly, by varying the regularization parameters, we can adjust the trade-off between robust attributions and predictions.

3.2 Recovering existing objectives

In this section, we show that our formulation of the robustness optimization problem is a generalization of already existing methods to train robust attribution maps.

Madry's Robust Prediction [13] can be recovered by utilizing the training objective in Equation (6) with $\lambda = 0$. It becomes as follows:

$$\theta^* = \arg \min_{\theta} \sum_{\mathbf{x} \in \mathcal{D}} \max_{\|\tilde{\mathbf{x}} - \mathbf{x}\|_p < \varepsilon} l(\tilde{\mathbf{x}}, y, f) \quad (7)$$

The **Axiomatic Attribution Regularization** terms (IG-NORM and IG-SUM-NORM) in [1] can be recovered using the regularization term in Equation (5) with the IG attribution map $S = \text{IG}(\tilde{\mathbf{x}}, \mathbf{x})$, where the baseline of IG is set to $\mathbf{b} = \mathbf{x}$ and the dissimilarity function to $d_s(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_1$. As such, Equation (5) becomes as follows.

$$\theta^* = \arg \min_{\theta} \sum_{\mathbf{x} \in \mathcal{D}} \left\{ l(\mathbf{x}, y, f) + \lambda \cdot \max_{\|\tilde{\mathbf{x}} - \mathbf{x}\|_p < \varepsilon} \|\text{IG}(\tilde{\mathbf{x}}, \mathbf{x})\|_1 \right\} \quad (8)$$

Algorithm 1 IFIA

Input: Classifier f , input \mathbf{x} , target class t , attribution map S , dissimilarity metric d_s , norm p and bound ε , step size η , iterations N , data input bounds b

Output: Adversarial example

\mathbf{x}_{adv}

```

1:  $\mathbf{x}_{\text{adv}} \leftarrow \mathbf{x}$ 
2: while  $i \leq N$  and  $\arg \max_t f_t(\mathbf{x}_{\text{adv}}) =$ 
    $\arg \max_t f_t(\mathbf{x})$  do
3:    $\mathbf{g}_t \leftarrow \nabla_{\mathbf{x}_{\text{adv}}} d_s [S(\mathbf{x}_{\text{adv}}, f), S(\mathbf{x}, f)]$ 
4:    $\mathbf{x}_{\text{adv}} \leftarrow \mathbf{x}_{\text{adv}} + \eta \cdot \text{Normalize}_p(\mathbf{g}_t)$ 
5:    $\mathbf{x}_{\text{adv}} \leftarrow \text{Project}_p(\mathbf{x}_{\text{adv}}, \mathbf{x}, \varepsilon, b)$ 
6: end while

```

Algorithm 2 Adversarial IFIA

Input: Classifier f , input \mathbf{x} , target class t , classification loss l , attribution map S , dissimilarity metric d_s , norm p and bound ε , step size η , iterations N , data input bounds b

Output: Adversarial example

\mathbf{x}_{adv}

```

1:  $\mathbf{x}_{\text{adv}} \leftarrow \mathbf{x}$ 
2: while  $i \leq N$  do
3:    $\mathbf{g}_t \leftarrow \nabla_{\mathbf{x}_{\text{adv}}} \{l(\mathbf{x}_{\text{adv}}, t, f) + \lambda \cdot$ 
      $d_s [S(\mathbf{x}_{\text{adv}}, f), S(\mathbf{x}, f)]\}$ 
4:    $\mathbf{x}_{\text{adv}} \leftarrow \mathbf{x}_{\text{adv}} + \eta \cdot \text{Normalize}_p(\mathbf{g}_t)$ 
5:    $\mathbf{x}_{\text{adv}} \leftarrow \text{Project}_p(\mathbf{x}_{\text{adv}}, \mathbf{x}, \varepsilon, b)$ 
6: end while

```

Figure 2: The IFIA (left) and Adversarial IFIA (right) attacks used to solve the inner maximizations of our framework. IFIA is also used to estimate AR during evaluation.

The IG-SUM-NORM training objective in [1] can be analogously derived from Equation (6). Note that $\text{IG}(\mathbf{x}, \mathbf{x}) = 0$ holds due to the completeness axiom of IG [1].

The input-gradient *Spatial Alignment* regularization term introduced in [20] corresponds to utilizing the sum of positive spatial alignment of true class input gradients and the negative spatial alignment of the second largest logits input gradient as attribution map S , written in the following Equation (9):

$$S(\mathbf{x}, f, y, \bar{y}) = \cos[g_{\mathbf{x}}^y(\mathbf{x}), \mathbf{x}] - \cos[g_{\mathbf{x}}^{\bar{y}}(\mathbf{x}), \mathbf{x}] \quad (9)$$

\cos denotes the *pointwise* cosine similarity between the input gradient and the image. \bar{y} is the second largest class’ logit, the rest of the notation is kept from previous sections. By using $d_s(\mathbf{x}, \mathbf{y}) = \log \{1 + \exp[-\sum_{i \in \text{dim}(\mathbf{x})} (x_i - y_i)]\}$ and omitting the use of target saliency maps, the regularization term in [20] can be recovered from Equation (5). See the technical appendix for a more detailed proof of these equations.

4 Adversarial Attributional Training

Next, we introduce *AAT* and *AdvAAT*, our instantiations of FAR that achieve robust attribution maps through optimizing for maximal correlation of explanations within a small local L_∞ neighbourhood of the input.

Using our framework, we formalize the *adversarial attributional training* objectives, consisting of a *regularization term* (*AAT*) that optimizes directly for robust attributions, and a *robust training loss* (*AdvAAT*) used to achieve both robust predictions and attributions. We choose the *Pearson correlation coefficient* [15] as attribution similarity, as it is a good proxy for optimizing for discrete rank correlations like CO and IN. These cannot be used directly due to their non-differentiable nature. We choose the aforementioned IG as attribution map, as it is a widely accepted axiomatic explanation method. Our attribution targets

are the saliency maps of the unperturbed inputs. This leads to the following optimization regularization term (10) and loss (11) respectively.

$$\theta_{\text{AAT}}^* = \arg \min_{\theta} \sum_{\mathbf{x} \in \mathcal{D}} \left\{ l(\mathbf{x}, y, f) + \lambda \cdot \max_{\|\tilde{\mathbf{x}} - \mathbf{x}\|_{\infty} < \epsilon} \text{PCL} [\text{IG}(\tilde{\mathbf{x}}, \mathbf{0}), \text{IG}(\mathbf{x}, \mathbf{0})] \right\} \quad (10)$$

$$\theta_{\text{AdvAAT}}^* = \arg \min_{\theta} \sum_{\mathbf{x} \in \mathcal{D}} \max_{\|\tilde{\mathbf{x}} - \mathbf{x}\|_{\infty} < \epsilon} \left\{ l(\tilde{\mathbf{x}}, y, f) + \lambda \cdot \text{PCL} [\text{IG}(\tilde{\mathbf{x}}, \mathbf{0}), \text{IG}(\mathbf{x}, \mathbf{0})] \right\} \quad (11)$$

with $\text{PCL} = 1 - \frac{\text{PCC} + 1}{2}$ denoting the loss derived from *Pearson correlation coefficient* PCC, the rest of the notation is kept from the previous sections.

The outer minimization is solved by standard gradient descent of the loss in the network parameter space. The inner maximizations of our AAT and AdvAAT methods during training are solved with the IFIA (Algorithm 1) and Adversarial IFIA (Algorithm 2) attacks from the previous section. During these attacks, we only approximate the second derivative of the ReLU networks with the second derivative of the Softplus activation $\nabla_{\mathbf{x}}^2 \text{ReLU}(\mathbf{x}) = \beta \cdot \text{sigmoid}(\beta \cdot \mathbf{x}) \cdot [1 - \text{sigmoid}(\beta \cdot \mathbf{x})]$, where $\text{sigmoid}(\mathbf{x}) = \frac{1}{1 + e^{-\mathbf{x}}}$ and β controls the approximation tightness of the ReLU [2]. As such, we decouple the attribution maps from the actual estimation of their robustness.

5 Experiments and Results

In this section, we report the experimental setup and evaluation of our AAT and AdvAAT methods. Utilizing the datasets MNIST [11], Fashion-MNIST [25], CIFAR-10 [10], GT-SRB [22] and Restricted Imagenet [3], we show that our methods outperform current state of the art on the former two datasets and perform comparably to state of the art on the latter three in terms of attributional robustness. Additionally, to our knowledge, we are the first to experimentally show the dependency of attribution robustness on the weight initialization of the networks and argue that training with our objectives lessens these dependencies. Moreover, we show that the tightness parameter β in the approximation of second order ReLU gradient significantly influences the robustness estimation.

Setup. We compare three state of the art attribution robustification methods (Adv, IG-SN and Align), taken from [13], [1] and [20] respectively, and a naturally trained (Nat) models' attributional robustness to networks trained with our robust training objectives from Equations (5) and (6) (AAT and AdvAAT), on the aforementioned five datasets. For MNIST and Fashion-MNIST, we train a two-layer convolutional neural network, for the other datasets we use a ResNet taken from [8]. In order to evaluate the attributional robustness of each model, the IFIA attack from [5] (Algorithm 1) is used, utilizing the proposed *Top-K intersection* attack from [5]. We use the IG attribution map and report the *Top-K intersection* (IN) of original and adversarial attribution map as well as their *Kendall rank order correlation* (CO) as robustness metrics. The natural and adversarial accuracy (NA and AA) of the models are also reported. AA is estimated with the PGD attack from the authors of [13]. A detailed description of the architectures, training and evaluation details can be found in the supplemental appendix. Table 1 contains the results of the comparison experiments. The results are run three times with different data splits and random seeds, and the average results are given.

Numerical analysis. Based on Table 1, we make the following conclusions. First, our methods outperform all other state of the art methods on MNIST and Fashion-MNIST. On the datasets CIFAR-10, GTSRB and Restricted Imagenet, our methods perform comparably to state of the art in terms of IN, while giving slightly worse results in terms of CO. Hence, we conclude that while AAT and AdvAAT do not outperform Align, they give promising results while being more general and wider applicable, as described in Section 3.1. This is backed by the phenomenon that our methods perform significantly better on MNIST and Fashion-MNIST than Align. We argue that this is due to Align being dependant on the nature of the data. A large proportion of the data are black pixels. Along these dimensions, the alignment from Equation (9) is inherently zero, independently of the gradients. Therefore, Align does not provide an optimization target along these dimensions. Moreover, white pixels are targeted to have large gradients (in alignment terms), but gradient saturation leads to small gradients for these pixels, further worsening optimization with Align on the two MNIST datasets. Our methods do not suffer from these shortcomings, as they provide optimization targets for each input dimension, independently of their values. We evaluated Align on the MNIST datasets with an input scaling between [-1, 1] as well, indicated as Align (s.) in Table 1. However, we see almost no improvement in terms of IN and CO compared to scaling between [0, 1] (Align). We believe that this is due to the arbitrary choice of input bounds. A lower bound of -1 encourages negative gradients, another arbitrary valid

Data	Model	NA (%)	AA (%)	IN	CO
MNIST	Nat	99.4	12.1	0.43	0.10
	Adv	98.9	92.7	0.52	0.19
	Align	98.7	2.6	0.52	0.40
	Align (s.)	95.2	12.3	0.58	0.43
	IG-SN	*98.3	*88.2	*0.72	*0.31
	AAT	98.4	0.0	0.76	0.72
	AdvAAT	98.7	77.1	0.77	0.73
Fashion-MNIST	Nat	91.5	11.0	0.43	0.20
	Adv	87.1	69.9	0.71	0.58
	Align	90.2	30.5	0.48	0.60
	Align (s.)	85.4	20.3	0.50	0.45
	IG-SN	*85.4	*70.3	*0.72	*0.67
	AAT	89.8	0.01	0.80	0.82
	AdvAAT	86.7	41.4	0.81	0.82
CIFAR-10	Nat	89.9	0.0	0.17	-0.02
	Adv	80.3	43.9	0.66	0.66
	Align	*89.8	*37.6	* 0.93	* 0.92
	IG-SN	-	-	-	-
	AAT	73.7	0.4	0.86	0.70
	AdvAAT	72.2	24.9	0.90	0.71
	GTSRB	Nat	98.5	14.7	0.39
Adv		94.9	66.7	0.72	0.64
Align		*98.5	*84.7	* 0.92	* 0.89
IG-SN		*95.7	*77.1	*0.74	*0.77
AAT		95.6	26.9	0.75	0.79
AdvAAT		91.7	65.9	0.84	0.80
Restr. Imagenet		Nat	89.1	0.0	0.08
	Adv	80.0	68.2	0.81	0.78
	Align	82.3	67.7	0.92	0.86
	IG-SN	-	-	-	-
	AAT	88.4	0.02	0.91	0.78
	AdvAAT	80.2	61.1	0.90	0.79

Table 1: Estimated attributional robustness (*Top-K intersection* IN and *Kendall’s rank order correlation* CO) of the models trained naturally (Nat), adversarially (Adv), *alignment*-based (Align), IG-SUM-NORM-based (IG-SN) as well as with our AAT and AdvAAT objectives. Their natural and adversarial accuracy is given in the NA and AA columns. Numbers indicated with an asterisk (*) are taken from the respective work and not reproduced by us. Align (s.) denotes the *alignment*-based method with input images scaled between -1 and 1.




Figure 3: Estimated attributional robustness (● IN and ● CO), natural (● NA) and adversarial (● AA) accuracies for our AAT (left) and AdvAAT (right) methods, evaluated on Fashion-MNIST varying the regularization parameter λ .




Figure 4: (a) Estimated attributional robustness (● IN and ● CO) of SM (Equation 1) with seven different initializations for the natural (Nat), adversarially (Adv) and attributionally (AAT) trained models on MNIST. (b) Gradient maps (SM) and their attacked maps of the natural model trained on MNIST for different weight initializations.

lower bound of 0.2 would encourage positive ones in the same dimensions. This highlights the flaws of the alignment-based method even more, namely that targets are not input shift invariant.

Our second conclusion comes from comparing our AAT method to AdvAAT. AAT achieves slightly worse attribution robustness than AdvAAT, but significantly worse adversarial accuracy for all datasets experimented on. This leads us to believe that while adversarial robustness does increase attributional robustness, the reverse is only limitedly true. We leave the theoretical analysis of this phenomenon to future work.

Dependency on the regularization parameter (λ). We examine the influence of the regularization parameter λ on the estimated robustness of attributions and predictions for our CNN trained on Fashion-MNIST. We chose this dataset because it is slightly more complex than MNIST, yet the computational burden of training is low. We train our AAT and AdvAAT models with λ -values varying from 0 to 1.5 and examine their robustness. Figure 3 contains the natural (NA) and adversarial (AA) accuracies as well as the attribution robustness metrics (IN and CO) for the AAT models to the left and AdvAAT models to the right. We observe that for both methods, higher λ values result in increased AR, with saturation occurring at values above 1. Moreover, for AdvAAT, the adversarial accuracy drops with increasing λ , while AR metrics increase, controlling the trade-off between adversarial and attribution robustness.

Dependency on network parameter initialization. Our experiments have shown that gradient-based attribution maps and their robustness estimates can depend on the initialization of the weights in the network. While resulting in nearly identical natural and adversarial accuracies, differently initialized networks yield considerably different robustness of gradient maps. We exemplify this with our natural, adversarially and AAT trained models on MNIST, by reporting their corresponding performance and attribution robustness estimates

for seven different network weight initializations. These are the default PyTorch [14] initialization (PTD), a custom initialization taken from [1] (CUST), a random uniform initialization of weights (UNI) as well as the default PyTorch He [7] and Glorot [6] uniform and normal (HU, HN, GU and GN) initializations (as listed in Figure 4(b) from left to right). Figure 4(a) reports the resulting attributional robustness estimates for the initialization methods. Both for natural and adversarially robust models, the variance of IN and CO is significant across the different initializations. We expect this behaviour, as heuristic search algorithms like SGD depends strongly on initial conditions. The gradient maps look notably different as well, as reported in Figure 4(b). This dependency is partly mitigated by our AAT method, but still present.

Dependency on the tightness parameter of the ReLU approximation (β).

Figure 5 shows the estimated *Top-K intersection* of the natural MNIST model while using different β values for the second gradient approximation. We observe that by varying this parameter, the *Top-K intersection* changes considerably. We further observe that by setting β too extreme, second gradients vanish, resulting in the IFIA attack not




Figure 5: Estimated attributional robustness (IN) of SM for the natural model trained on MNIST, varying the β parameter of the ReLU approximation.

being able to find good adversarial inputs. Previous work [2] has already shown the dependency of AR on β , however, we are the first to only use this approximation for the second order gradients. Therefore, we keep saliency maps unchanged, giving a better estimate for the true attribution robustness of ReLU networks.

6 Conclusion and Future Work

This work introduced a generalized notion of attributional robustness with FAR providing objectives for increasing the robustness of explanations in DNNs. This allows direct optimization for robust attributions, with optionally coupling it to robust predictions. We showed how current existing objectives can be instantiated from this framework. Moreover, we provided novel instantiations of FAR, AAT and AdvAAT, which directly optimize for high correlation of attributions as well as robust predictions for similar inputs. They perform comparably to or better than current state of the art methods in terms of AR, utilizing fewer assumptions and generalizing better (see Section 5). Finally, we identified parameter dependencies of robust attributions that necessitate careful assessment of methods on their dependencies on these parameters.

This work opens up many interesting directions for future research. First, we are interested in assessing the robustness of *non-differentiable maps* like Occlusion [26], utilizing gradient estimation techniques. Second, exploring the *connection between robust predictions and robust attributions* might lead to additional insights into the decision process of neural networks, enhancing interpretability. Lastly, since second order gradients in DNNs seem highly irregular, and their optimization is hard, the assessment of *additional training parameters* that influence AR would help further research in establishing fair comparison and tracking of true progress in this area.

References

- [1] 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.
- [2] 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.
- [3] Logan Engstrom, Andrew Ilyas, Hadi Salman, Shibani Santurkar, and Dimitris Tsipras. Robustness (python library), 2019. URL <https://github.com/MadryLab/robustness>.
- [4] Christian Etmann, Sebastian Lunz, Peter Maass, and Carola-Bibiane Schönlieb. On the connection between adversarial robustness and saliency map interpretability. *arXiv preprint arXiv:1905.04172*, 2019.
- [5] 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.
- [6] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feed-forward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 249–256, 2010.
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1026–1034, 2015.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [9] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.
- [10] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, Department of Computer Science, University of Toronto, 2009.
- [11] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- [12] Yann LeCun. Lenet-5, convolutional neural networks. URL: <http://yann.lecun.com/exdb/lenet>, 20(5):14, 2015.
- [13] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.

- [14] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8026–8037, 2019.
- [15] Karl Pearson. Notes on regression and inheritance in the case of two parents. *Proceedings of the Royal Society of London*, 58(347-352):240–242, 1895.
- [16] Andrew Slavin Ross, Michael C Hughes, and Finale Doshi-Velez. Right for the right reasons: Training differentiable models by constraining their explanations. *arXiv preprint arXiv:1703.03717*, 2017.
- [17] Johannes Schneider and Michalis Vlachos. Reflective-net: Learning from explanations. *arXiv preprint arXiv:2011.13986*, 2020.
- [18] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. *arXiv preprint arXiv:1704.02685*, 2017.
- [19] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- [20] Mayank Singh, Nupur Kumari, Puneet Mangla, Abhishek Sinha, Vineeth N Balasubramanian, and Balaji Krishnamurthy. On the benefits of attributional robustness. *arXiv preprint arXiv:1911.13073*, 2019.
- [21] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*, 2017.
- [22] Johannes Stalldkamp, Marc Schlipf, Jan Salmen, and Christian Igel. The german traffic sign recognition benchmark: a multi-class classification competition. In *The 2011 International Joint Conference on Neural Networks*, pages 1453–1460. IEEE, 2011.
- [23] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 3319–3328, 2017.
- [24] Zifan Wang, Haofan Wang, Shakul Ramkumar, Matt Fredrikson, Piotr Mardziel, and Anupam Datta. Smoothed geometry for robust attribution. *arXiv preprint arXiv:2006.06643*, 2020.
- [25] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- [26] Matthew Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European Conference on Computer Vision*, pages 818–833. Springer, 2014.