

Rowan University

## Rowan Digital Works

---

Henry M. Rowan College of Engineering  
Departmental Research

Henry M. Rowan College of Engineering

---

7-31-2023

# EvalAttAI: A Holistic Approach to Evaluating Attribution Maps in Robust and Non-Robust Models

Ian E. Nielsen  
*Rowan University*

Ravi Ramachandran  
*Rowan University*

Nidhal Carla Bouaynaya  
*Rowan University*

Hassan M. Fathallah-Shaykh

Ghulam Rasool  
*Rowan University*

Follow this and additional works at: [https://rdw.rowan.edu/engineering\\_facpub](https://rdw.rowan.edu/engineering_facpub)



Part of the [Biomedical Engineering and Bioengineering Commons](#), and the [Electrical and Computer Engineering Commons](#)

---

### Recommended Citation

I. E. Nielsen, R. P. Ramachandran, N. Bouaynaya, H. M. Fathallah-Shaykh and G. Rasool, "EvalAttAI: A Holistic Approach to Evaluating Attribution Maps in Robust and Non-Robust Models," in *IEEE Access*, vol. 11, pp. 82556-82569, 2023, doi: 10.1109/ACCESS.2023.3300242.

This Article is brought to you for free and open access by the Henry M. Rowan College of Engineering at Rowan Digital Works. It has been accepted for inclusion in Henry M. Rowan College of Engineering Departmental Research by an authorized administrator of Rowan Digital Works.

Received 2 June 2023, accepted 22 July 2023, date of publication 31 July 2023, date of current version 9 August 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3300242

## APPLIED RESEARCH

# EvalAttAI: A Holistic Approach to Evaluating Attribution Maps in Robust and Non-Robust Models

IAN E. NIELSEN<sup>1</sup>, RAVI P. RAMACHANDRAN<sup>1</sup>, (Senior Member, IEEE),  
NIDHAL BOUAYNAYA<sup>1</sup>, (Member, IEEE),  
HASSAN M. FATHALLAH-SHAYKH<sup>2</sup>, (Member, IEEE),  
AND GHULAM RASOOL<sup>3</sup>, (Member, IEEE)

<sup>1</sup>Department of Electrical and Computer Engineering, Rowan University, Glassboro, NJ 08028, USA

<sup>2</sup>UAB Heersink School of Medicine, University of Alabama at Birmingham, Birmingham, AL 35294, USA

<sup>3</sup>Machine Learning Department, H. Lee Moffitt Cancer Center and the Research Institute, Tampa, FL 33612, USA

Corresponding author: Ian E. Nielsen (nielsen6@students.rowan.edu)

This work was supported in part by the National Science Foundation under Award ECCS-1903466, Award OAC-2008690, and Award OAC-2234836. The work of Ian E. Nielsen was supported by the U.S. Department of Education through a Graduate Assistance in areas of National Need (GAANN) Program under Award P200A180055.

**ABSTRACT** The expansion of explainable artificial intelligence as a field of research has generated numerous methods of visualizing and understanding the black box of a machine learning model. Attribution maps are commonly used to highlight the parts of the input image that influence the model to make a specific decision. At the same time, numerous recent research papers in the fields of machine learning and explainable artificial intelligence have demonstrated the essential role of robustness to natural noise and adversarial attacks in determining the features learned by a model. This paper focuses on evaluating methods of attribution mapping to find whether robust neural networks are more explainable, particularly within the application of classification for medical imaging. However, there is no consensus on how to evaluate attribution maps. To solve this, we propose a new explainability faithfulness metric, EvalAttAI, that addresses the limitations of prior metrics. We evaluate various attribution methods on multiple datasets and find that Bayesian deep neural networks using the Variational Density Propagation technique are consistently more explainable when used with the best performing attribution method, the Vanilla Gradient. Our results suggest that robust neural networks may not always be more explainable, despite producing more visually plausible attribution maps.

**INDEX TERMS** Explainability, robustness, Bayesian neural networks, medical imaging.

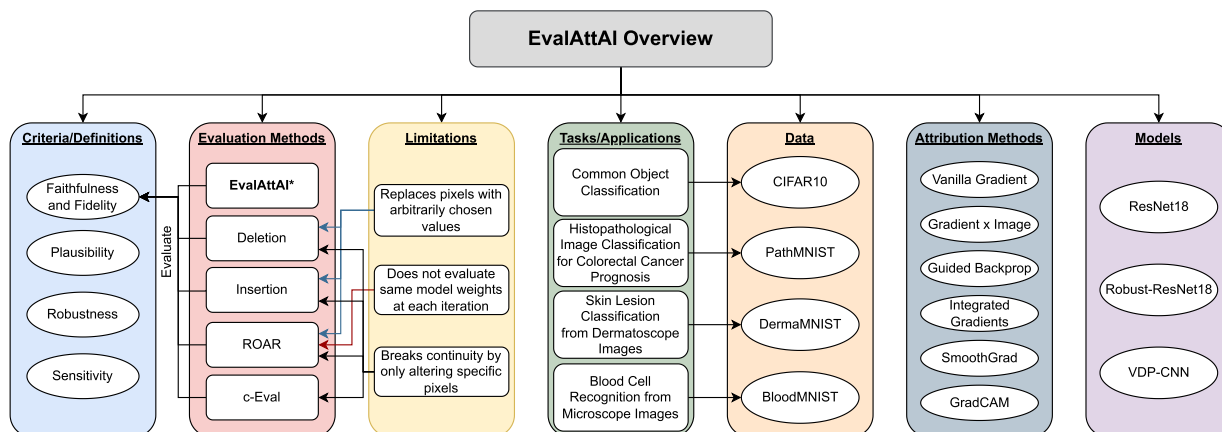
## I. INTRODUCTION

With the recent explosion of black box machine learning models over the past few years, there has been a great demand to explain how these models work so that users can understand and trust the results [1], [2], [3]. This is especially important in mission-critical and life-saving applications, such as clinical diagnosis and medical decision-making [4], [5], [6]. Explainable AI (XAI) seeks to give a user more trust

in the model explanation [7]. The problem is that we do not know if we can trust the explanations either. In order for XAI to be viable for real world applications there must be trust, thus there have been many attempts to evaluate these explanations [7], [8], [9], [10], [11], [12], [13].

These application areas also demand robust machine learning models that resist natural noise in the input data and malicious or adversarial attacks [14], [15], [16]. In real-world tasks, the data can often come with noise and artifacts that the machine learning model might not have previously seen [17]. There might even be nefarious actors trying to confuse the

The associate editor coordinating the review of this manuscript and approving it for publication was Wai-Keung Fung<sup>1</sup>.



**FIGURE 1.** Concept map illustrating the key components of the study explored in this paper, including methods, criteria, applications and models. The map provides a visual breakdown of the relationships between these components.

model with adversarial examples [14], [18]. This requires that models be built and trained to resist noise and attacks [15], [19]. There are several approaches to building robust machine learning models, including training on noisy datasets and Bayesian models [6], [14], [15], [16], [17], [18]. The first approach only focuses on model training and data processing, while the second alters the neural network architecture and introduces probability distribution functions over learnable parameters [14]. Robust models have been shown to produce more visually plausible explanations [9], [20]. However, the quantitative evaluation of the faithfulness of these explanations is a challenging task. There are many criteria for evaluating explanations. In this paper, we propose a much more direct approach (called EvalAttAI) that eliminates the errors introduced by existing approaches to evaluating the faithfulness of explanations. The contributions in this paper are as follows:

- We develop a new method of evaluating the faithfulness of an attribution map, which is much different from existing approaches and avoids the errors present in current methods. Our proposed method is denoted as Evaluating Attributions by Adding Incrementally (EvalAttAI).
- We compare our method to the state-of-the-art in the literature and show why ours is a more fair and accurate measure of faithfulness of model explanations.
- We relate the concepts of faithfulness and robustness by showing whether robust models produce more faithful explanations when evaluated using our proposed EvalAttAI method.

In Fig. 1, we provide an overview of the key components of this study. This concept map includes all methods, datasets and models used in our experiments. Most importantly, the figure highlights the fact that the evaluation methods discussed in this paper measure the faithfulness of explanations while also outlining the limitations of each approach.

The article is organized as follows. In Section II, we provide definitions of various terms that are used in the article and present various explainability faithfulness

evaluation methods, including our newly proposed, EvalAttAI. Section III provides an overview of our methods and simulation experiments. In Section IV, we present results and discuss the same in Section V before concluding in Section VI.

## II. EXPLAINABILITY METRICS AND DEFINITIONS

To evaluate the explainability method defined for a machine learning model, quantify its faithfulness, and link it to the robustness of the machine learning model, we must understand different concepts, terms, and metrics that are used in the literature. This section provides an overview of these important concepts.

### A. ROBUSTNESS

A model is said to be robust if its performance shows little decrease when the distribution of the test data is different from that used during the training. A model can be robustified (made more robust) by training it on a large variety of data that includes various types of noisy and adversarial images.

This term is also used to describe the behavior of attributions. A robust attribution method should produce a similar map for similar inputs to the model. This can be seen when generating attributions with and without adversarial attacks or imperceptible noise, since the image is similar, but a non-robust method would return a vastly different attribution map. It should also be noted that robustly trained models produce more robust (and plausible) explanations [7], [20].

### B. ATTRIBUTION MAPS

Explainability methods for image processing applications creates a pixel-wise array of attribution scores that indicate each feature’s (a pixel of the input image in most cases) importance. This is known as an attribution map or an explainability map [21].

### C. SENSITIVITY

Sensitivity of an explanation method (i.e., attribution map generation method, such as Vanilla Gradient) is defined to

be how much an explainability map changes when small perturbations are applied to the input of the model [22]. Attribution methods that produce explanations with large sensitivity scores are likely to produce dramatically different explanations when noise is introduced at the input of the model. This means that, in general, robust models tend to perform better than their non-robust counterparts when the attribution map is evaluated for sensitivity [7]. Generally, attribution methods with lower sensitivity produce more similar and consistent results in the presence of imperceptible noise than methods with higher sensitivity. More consistent explanations may be beneficial in many applications, but it does not reveal how important the pixels in the attribution map actually are to the model [7]. This motivates the examination of other metrics for evaluating explainability methods.

#### D. PLAUSIBILITY

The first instinct is to visually examine an explainability map to see whether the attribution makes sense to a human observer [7], [22]. This is called plausibility [22], which tells the users how visually convincing the explanation is to a human. For example, an attribution map that has many important pixels within the object of interest would be more plausible than an explanation that focuses on features that are irrelevant to a human. A plausibility metric has a subjective component, since humans define what are the important features according to their own visual perception. There are multiple methods which give a plausibility score corresponding to the number of important pixels that fall within a human drawn region containing the object [23], [24], [25]. It has been shown that robust and adversarially trained models tend to produce more plausible explanations [7], [9], [26], [27]. However, plausibility can be misleading [7]. The explanations that look more reasonable to a human might misrepresent the features that the trained machine learning model uses for its internal processing and decision-making. Therefore, it is important to introduce more quantitative metrics that can help us evaluate and understand the “goodness” or “faithfulness” of an explanation.

#### E. FAITHFULNESS AND FIDELITY

The faithfulness of an explanation is defined to be a measure of how accurate the explanation (an attribution map in our case) is to the model itself. An explanation that is very faithful will show the user what is *truly* most important to the model. In other words, faithfulness tells us the extent to which pixels deemed to be important in the attribution map are *actually* important to the model.

In many recent works, fidelity and faithfulness were considered synonymous [28], [29], [30]. We argue that fidelity metrics are used to measure the faithfulness of the attribution map (or the explanation). Many different techniques have been proposed in the literature to measure the faithfulness or fidelity [9], [11], [12], [13], [28], [29], [30], [31], [32], [33],

[34]. However, there is no current consensus on the best way to evaluate faithfulness of attribution maps.

The *Fidelity* metric measures the correlation between the pixel attribution scores and the drop in prediction or accuracy scores when the pixel is altered or removed. This shows how well the attribution map ranks the pixels by their importance. The most common implementations of fidelity metrics involve replacing important pixels with another pixel value (most often black or the image mean) [9], [11], [12], [13], [33], [34]. This introduces error since the model has not been trained to understand arbitrary pixels being introduced in the input image. There have been attempts to rectify this by retraining the model for possible changes in the pixels [8]. However, methods based on retraining introduce another type of error, the explanation (attribution map) is no longer being evaluated using the same model parameters (that is, the model has been modified due to retraining). Later in this work, we show that our proposed EvalAttAI eliminates the errors present in these existing methods. The reason is linked to the fact that EvalAttAI does not remove pixels, but rather perturbs pixels to a small, almost imperceptible degree, maintaining continuity.

#### F. MODEL ROBUSTNESS AND ATTRIBUTION FAITHFULNESS, FIDELITY AND PLAUSIBILITY

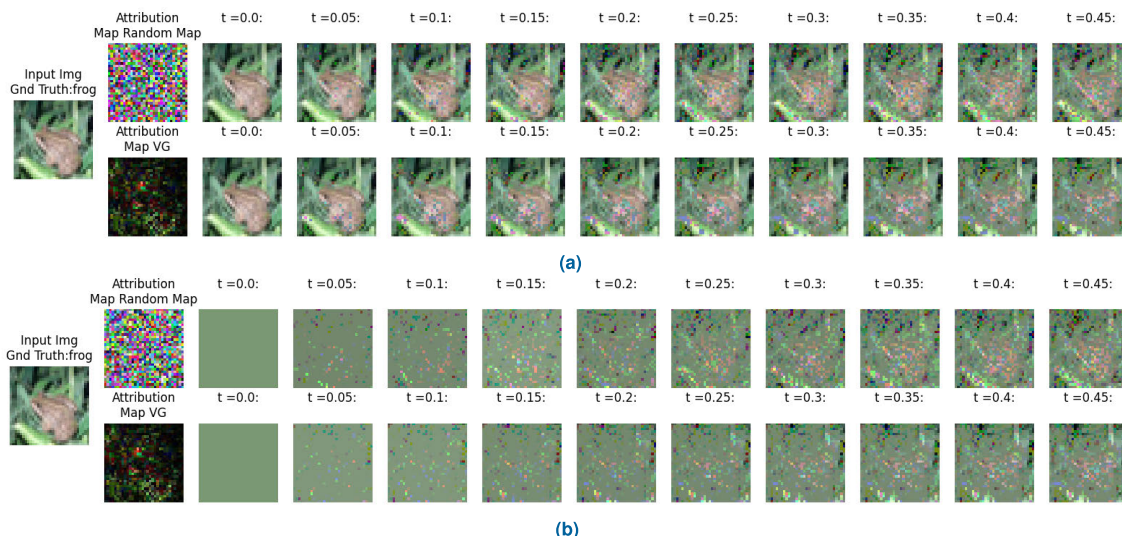
Using faithfulness, fidelity and plausibility, we will seek to shed more light on whether robust models are more explainable. As mentioned earlier, robust models have been shown to be more plausible. However, there is no consensus on whether the attribution maps of robust models are more faithful or whether there is a standard approach to evaluating faithfulness.

#### G. ATTRIBUTION EVALUATION METHODS

The most prominent faithfulness evaluation metrics currently include Insertion [34] and Deletion [13], [33], [34]. These methods are based on replacing pixels, and recently it was reported that their results appear contradictory [9]. That is, when Deletion performed well, Insertion did not and vice versa.

##### 1) DELETION

Deletion is the most widely used method of evaluating faithfulness of attribution maps and has many variations [13], [33], [34]. Deletion involves (1) incrementally removing pixels, (2) replacing removed pixel spaces with an arbitrary value, and (3) evaluating the resulting change in model predictions. The order in which the pixels are chosen is based on attribution scores of these pixels. The faithfulness of the attribution map is evaluated based on the change in the model prediction. However, when replacing pixels with an arbitrary value, an error is introduced. The replaced pixels create a sharp contrast in the image, which the model was never trained to process. There have been attempts to fix this by retraining the model for each increment [8]. However, retraining introduces new errors since the attributions are being evaluated



**FIGURE 2.** Two different metrics (Insertion and Deletion) for assessing various attribution maps are presented. The variable  $t$  indicates the percentage of pixels deleted/inserted, where  $t = 0$  refers to zero pixels and  $t = 1$  refers to all pixels. Removed pixels are replaced by the per channel mean. (a) Deletion and (b) Insertion examples using Vanilla Gradient (VG) and random Gaussian noise as the attribution maps.

on a newly trained model. EvalAttAI avoids this issue by perturbing the pixels rather than removing them and replacing with arbitrary values.

### 2) INSERTION

Insertion is another method of evaluating the faithfulness of attribution maps. This method uses the same process as Deletion, but starts with all pixels removed from the input image and incrementally inserts the most important ones back. All the issues that pertain to Deletion also apply to Insertion. Even worse, these two appear to produce contradictory results showing that these methods may not be evaluating the same thing, that is, the faithfulness of the attribution methods [9].

### 3) c-EVAL

The c-Eval approach [35] evaluates attribution maps by perturbing the most important pixels using an adversarial attack to determine how this perturbation affects model predictions. However, not all pixels in an adversarial attack have the same impact on the output. Therefore, selecting only the most important pixels to perturb is prone to error. Moreover, the adversarial attacks are designed to be applied to all pixels in the image, so only applying these to some pixels breaks the continuity. EvalAttAI ameliorates this by applying the perturbation to all pixels. EvalAttAI also takes a much more direct approach than c-Eval by perturbing the pixels using attribution scores, rather than adversarial attacks that are unrelated to the attribution maps and explanations being evaluated.

## H. EVALUATING ATTRIBUTIONS by ADDING INCREMENTALLY (EvalAttAI)

Our novel metric for evaluating faithfulness is based on the concept of unification between the explanation and input image. By observing the decrease in accuracy resulting from

the explanation’s influence, we can assess its faithfulness. All other faithfulness metrics discussed in this section segment the attributions, thereby breaking the continuity of the explanations in the process. Instead of segmenting the explanation, our approach unites a scaled attribution map to the input image thereby maintaining continuity at every iteration. In doing so, EvalAttAI emphasizes important pixels and de-emphasizes unimportant pixels based on the importance scores that the model directly attributes to each feature. Using our novel evaluation, a more accurate analysis of how robustness relates to faithfulness can be achieved.

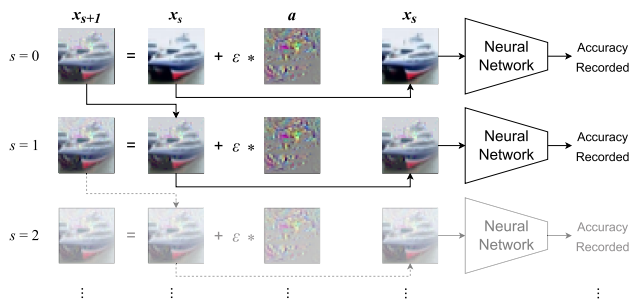
In this work, we focus on gradient-based explainability methods to show the applicability of EvalAttAI. However, EvalAttAI can be used for all other attribution generation methods without any changes. In general, for gradient-based methods, the attribution scores are calculated by taking the gradient of the activation score (logit or soft-max class probability values) with respect to the input features via backpropagation [21]. Since we will be perturbing pixels to decrease activation scores, it makes sense to look at gradient-based methods. Moreover, gradient-based attribution methods are derived directly from the loss function, which is used to train the model. Thus, these methods can be considered appropriate and effective for our proposed evaluation metric. We consider that the Vanilla Gradient method of explainability [36] will be the most faithful, since it directly leverages the gradient of the trained model without any further modification or changes [7].

In our experiments, we show that EvalAttAI can serve as a reliable new metric to significantly improve user trust in machine learning models. This is especially critical in medical and clinical applications.

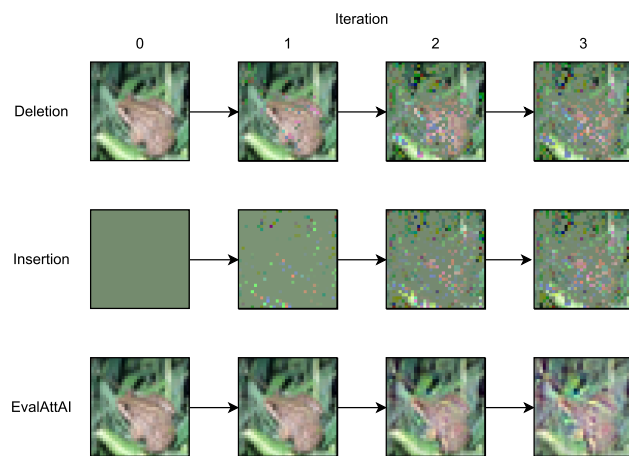
The novelites of EvalAttAI are as follows:

- **Uniqueness:** EvalAttAI defines and tests faithfulness in a way that has not been done before.





**FIGURE 3.** A schematic layout of the process of testing an attribution map using the proposed Evaluating Attributions by Adding Incrementally (EvalAttAI) method. The column labeled  $x_{s+1}$  depicts the modified input image, and  $x_s$  ( $s = 0$ ) shows the clean image in the first row and the modified image from the previous iteration on all subsequent rows. The attribution map is indicated by  $a$ .



**FIGURE 4.** A visualization of each evaluation method is shown. Each image is formatted as depicted before being input into the neural network. Deletion incrementally removes and replaces pixels from the original image, and Insertion adds the original pixels back from a pixel removed image. EvalAttAI takes an entirely different approach by shifting pixel values according to the importance scores on the attribution map.

- **Continuity:** EvalAttAI integrates all pixels in the explanation into the input image, maintaining continuity at every iteration. In contrast, other methods explored in this paper break continuity by altering some pixels, but not all, at each iteration.
- **No arbitrary parameters:** EvalAttAI eliminates errors present in Deletion and Insertion by not removing and replacing pixels with arbitrary values.
- **Consistent evaluation:** EvalAttAI evaluates all iterations on the same model weights, unlike ROAR which retrains the model each time, making each iteration an evaluation of an entirely different trained model.
- **Intelligent targeting:** EvalAttAI emphasizes and de-emphasizes important and unimportant pixels, respectively. It utilizes the model explanations themselves to intelligently weigh important (high scoring) pixels to evaluate the change in accuracy.

### III. METHODS AND EXPERIMENTS

This section describes the models, datasets and experimental protocol used in our study.

#### A. MODELS AND DATASETS

There are three models considered in our study. The first two are ResNet18 [37] with different training scenarios. (1) The standardly trained ResNet18 uses a dataset with no modifications to the training data. (2) For the robustly trained ResNet18, random Gaussian noise is added to each input with a signal-to-noise ratio (SNR) of 5 dB. (3) The third model is a standardly trained Bayesian deep neural network, based on Variational Density Propagation (VDP-CNN) [14], [16], [17]. This model follows the same architecture as ResNet18, but propagates both the mean and the covariance of the probability distribution function defined over the model parameters through each layer of the model.

All three models were trained on the CIFAR10 dataset [38]. The ResNet18 standard model was trained until it achieved an 84% validation accuracy. The robustly trained ResNet18 model achieved a validation accuracy of 70%. The VDP-CNN was trained on CIFAR10 to achieve a validation accuracy of 68%.

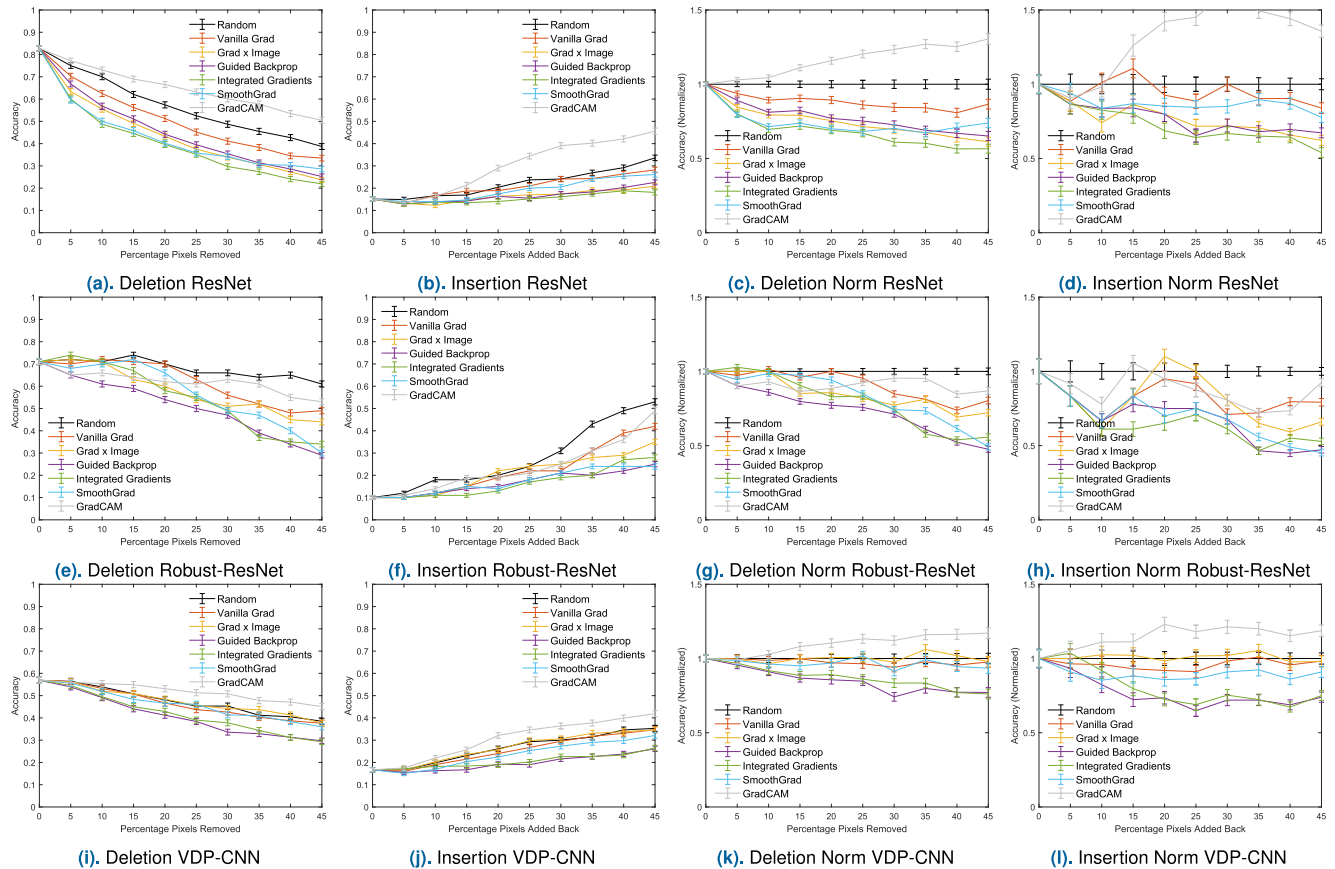
The EvalAttAI method was also evaluated on three medical imaging datasets, which are part of MedMNIST. The dataset included PathMNIST, DermaMNIST, and BloodMNIST [39], [40]. All three are multi-class prediction datasets consisting of 9, 7, and 8 labels for PathMNIST, DermaMNIST, and BloodMNIST, respectively. The datasets consist of images that are classified to aid in the diagnosis of colon, skin and blood cancers. The images from the datasets used all 3 color channels. This was chosen because the ResNet architecture is designed to classify colored image data.

The standard trained ResNet18 model was trained on PathMNIST to 85%, DermaMNIST to 73% and BloodMNIST to 92% validation accuracy. The robustly trained ResNet18 was trained on PathMNIST to 86%, DermaMNIST to 73% and BloodMNIST to 93% validation accuracy. Lastly, the VDP model was trained on PathMNIST, DermaMNIST and BloodMNIST to get 76%, 72%, and 84% validation accuracies, respectively.

We evaluated six attribution methods on each model including (1) Vanilla Gradient [36], (2) Grad x Image [21], (3) Guided Backprop [41], (4) Integrated Gradients [42], (5) SmoothGrad [43] and (6) GradCAM [44]. For the GradCAM method, we selected the first convolution layer. This means that the gradient was backpropagated to the first layer before performing global average pooling, linear combination and ReLU, resulting in the GradCAM attribution map. Attribution maps were generated using the Captum library [45], [46], along with built-in PyTorch functions.

#### B. DELETION AND INSERTION METHODS

The Deletion metric evaluates how much the accuracy changes when important pixels are removed from the input image [13], [33], [34]. Pixels can be replaced with various values, but are often replaced with the image mean. The pixel importance is determined by the magnitude of the attribution scores on the map being evaluated. One starts with a clean



**FIGURE 5.** Evaluating faithfulness of multiple different explainability attribution methods using Deletion and Insertion metrics. Each attribution method was evaluated using 5000 CIFAR10 images with three trained models. The top, middle, and bottom rows depict the standard trained ResNet-18 model, robustly-trained ResNet-18 (Robust-ResNet), and a Bayesian CNN trained using the VDP method respectively. The black line in all sub-figures show the behavior of the model when pixel are randomly deleted/inserted. The 95% confidence interval is shown in each sub-figure using error bars. (Left) Deletion. (Center left) Insertion. (Center right) Deletion after normalization (abbreviated as Norm) with random baseline. (Right) Insertion after normalization with random baseline. A method is considered to be performing well on the Deletion method if the accuracy drops significantly more than random. Inversely, a method that is performing well on the Insertion metric will show the accuracy increase more than random.

image and incrementally removes pixels starting at the highest attribution score. The attribution maps are considered to be performing well if the accuracy drops significantly faster than the random baseline, which implies a smaller area under the curve (AUC). In our opinion, the AUC satisfactorily captures how well each method is performing overall by accounting for the performance over all increments using a single score. The random baseline consists of random Gaussian noise that is tested in place of an attribution map. The process of Deletion can be seen visually in Figs. 2a and 4.

The Deletion experiment begins with a clean image. Then, starting from the most important according to attribution score, the pixels are removed in 5% increments. The first increment 0% evaluates the accuracy of the image with no pixels removed (i.e., the clean image). The next increment finds the accuracy with 5% of the pixels in the image removed. This is continued until 45% (almost half) of the pixels in the image are removed. The pixels are replaced with the mean of each color channel, respectively.

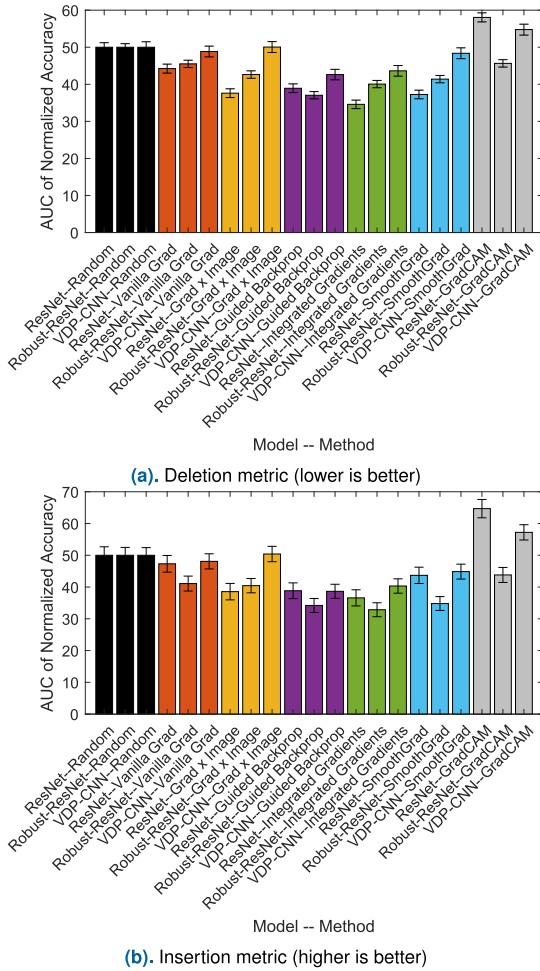
Insertion [34] uses the same premise as Deletion. The difference is that one starts with a blank image consisting

of all removed pixels, and incrementally introduces the most important pixels. The more significantly the accuracy increases compared to the random baseline, the better the attribution map is performing with Insertion. The Deletion and Insertion metrics may introduce error, since the models are not trained to be able to interpret missing and replaced pixels. In fact, the values that are chosen to replace the pixels can drastically alter the model accuracy in unintended ways. Insertion is depicted visually in Figs. 2b and 4.

For Insertion, we begin at 0% with an image where all pixels are replaced by the per channel mean of the dataset. We then introduce pixels in 5% increments until 45% of the image is restored. The entire image can be restored, but the most meaningful change in accuracy occurs at the beginning of the replacement with the most important pixels.

### C. EvalAttAI

Our metric overcomes The limitations of the other faithfulness metrics by (1) being continuous over all iterations and (2) intelligently targeting pixels based on the explanations themselves. Our method of evaluating the faithfulness of



**FIGURE 6.** Area under the curve (AUC) calculated using normalized accuracy curves (from Figs. 5c, 5d, 5g, 5h, 5k and 5l) for various models and attribution methods are presented. The models include a non-robust model (ResNet), a robustly-trained model (Robust-ResNet), and a Bayesian robust model (VDP-CNN). For the Deletion metric, lower values means the evaluation method and model are more explainable. For the Insertion metric, higher scores are better. The 95% confidence interval is shown in each sub-figure.

attributions takes a much different approach than Deletion and Insertion by not removing pixels at all and maintaining continuity of the explanation. Our approach is to perturb pixels proportionally to the pixel scores of the attribution map. This is done by joining a scaled down attribution map together with the original input image and then passing the new image through the machine learning model. The faithfulness of each attribution method is assessed by comparing the accuracy of the trained model before and after adding the perturbations generated using attribution scores.

We can see how EvalAttAI looks visually in comparison to Deletion and Insertion in Fig. 4. We expect that the more faithful the attribution method is, the more the accuracy will decrease after the perturbation. This implies that the maps which cause the steepest drop in accuracy are considered the most faithful. This is because highly faithful attribution maps will identify and perturb the most important pixels

first, which is verified by the significant drop in the model accuracy. However, if unimportant pixels are perturbed, then the accuracy will not drop as much, perhaps not at all. This will indicate that the attribution map produced by the explainability method is less faithful. The experiments that we performed also evaluate a baseline map, which consists of random Gaussian noise with a standard deviation of 0.25 and a mean of 0.

The EvalAttAI method is described in Eq. 1. In our experiments, the attribution map ( $\mathbf{a}$ ) is multiplied by a scaling variable epsilon ( $\epsilon$ ), which we set to 0.1, resulting in a scaled down attribution map. The variable  $\epsilon$  functions like a learning rate, allowing for control over the step size of each iteration. Although  $\epsilon$  can be assigned any positive value, we discovered that values close to zero work best. Starting with an image ( $\mathbf{x}_s$ ), we add the scaled attribution, resulting in the modified image ( $\mathbf{x}_{s+1}$ ). This is done in an iterative fashion until the desired number of steps ( $s$ ) are completed. The process always begins at  $s = 0$  with the clean image ( $\mathbf{x}_0$ ). Figure 3 and Eq. 1 describe the process, which is repeated until the desired number of steps ( $s$ ) are completed.

$$\mathbf{x}_{s+1} = \mathbf{x}_s + \epsilon * \mathbf{a} \quad (1)$$

EvalAttAI maintains continuity of the explanations over all iterations and intelligently targets pixels based on the importance scores attributed to them. The most important pixels according to the attribution map will be perturbed the most. At the same time, the least important pixels, which have attribution scores closest to 0, will not significantly alter the corresponding pixels on the input image. This method avoids the error introduced by Deletion and Insertion, which breaks continuity due to pixels being entirely removed and replaced. Our method can also be tuned using  $\epsilon$  so that pixels are altered in even smaller increments. The EvalAttAI method can be thought of as a more continuous approach to evaluating faithfulness.

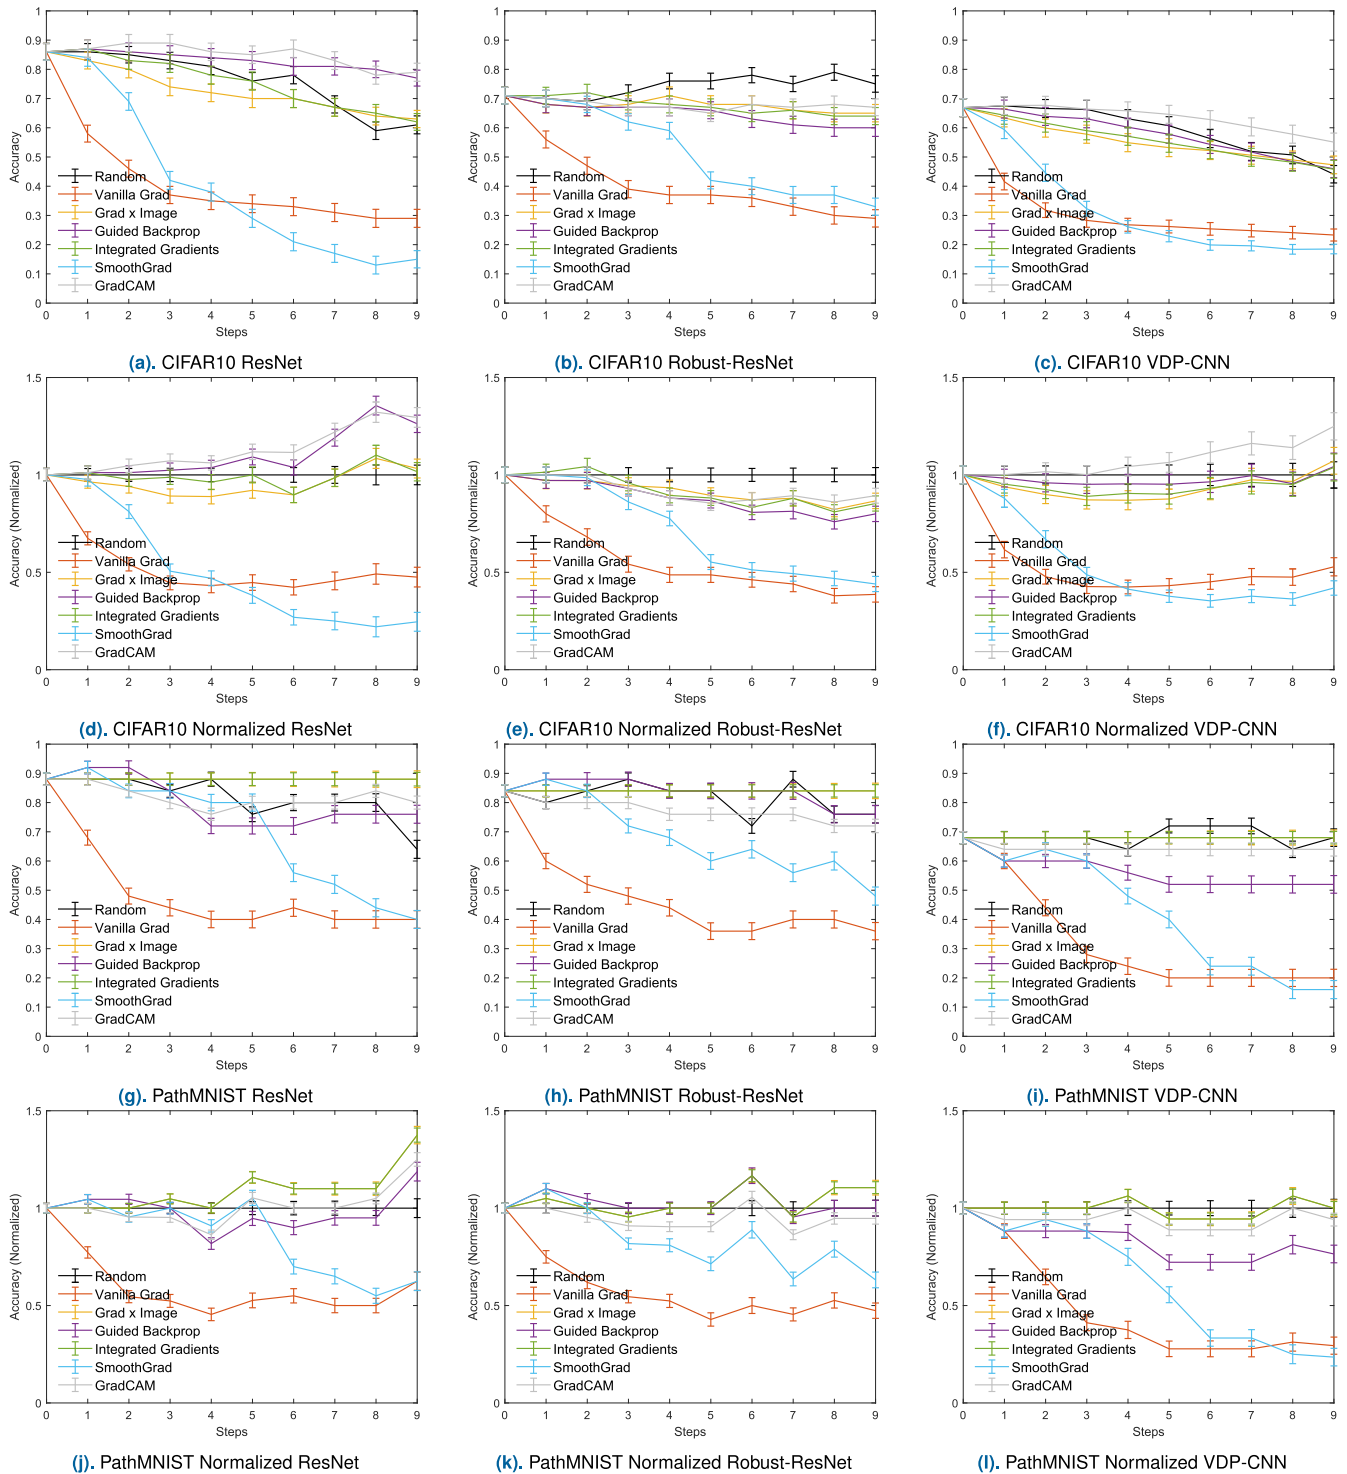
#### IV. RESULTS

In this section, we present results to show how well various attribution evaluation methods including EvalAttAI perform. In order to ensure a fair comparison, we normalize the recorded output accuracy such that the random line (baseline method) is always equal to one. Any time that normalization is discussed in this paper, it is being used to explain how the results were formatted after they were collected. For instance, in Figs. 5, 6, 7 and 8, we show the results before and after normalization of the data such that we can use the random line as a baseline across all models.

##### A. DELETION AND INSERTION

The goal of this first experiment is to quantify the behavior of Deletion and Insertion metrics for evaluating various types of attribution methods using three different types of trained models (ResNet, Robust-ResNet, and VDP-CNN). These models represent a non-robust model (ResNet), a robustly-trained model (Robust-ResNet), and a Bayesian

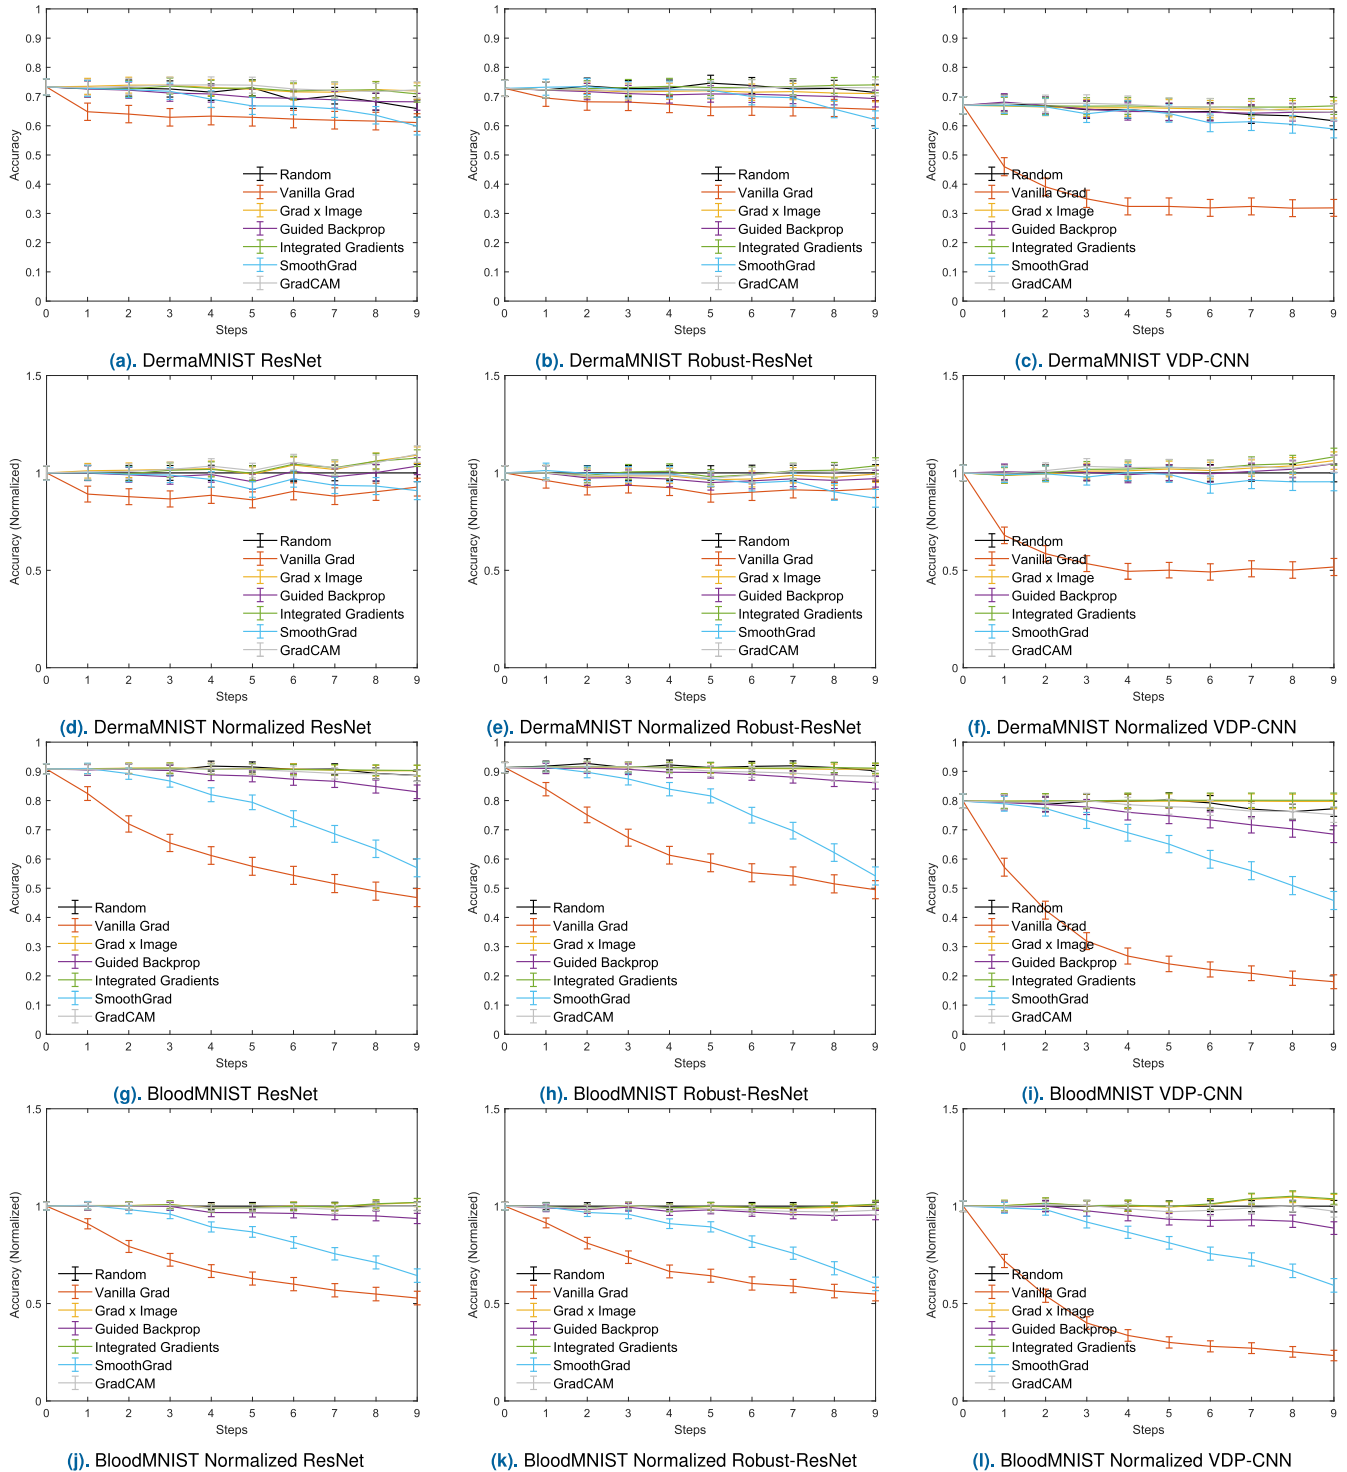




**FIGURE 7.** Results of the EvalAttAI metric tested using CIFAR10 and PathMNIST on three models, ResNet18, Robust-ResNet18 and VDP-CNN. The y-axis of all sub-figures shows average accuracy as attributions are added to the input image, which is captured by the x-axis. Lower accuracy values represent that the attribution method is faithful, and the metric is meant to capture this aspect. The 95% confidence interval is shown in each plot using error bars. (a)-(c) and (g)-(i) Non-normalized accuracy. (d)-(f) and (j)-(l) Normalized accuracy using random as the baseline.

model trained using the VDP technique (VDP-CNN). The test accuracy and normalized test accuracy are presented in Fig. 5. The AUC of normalized test accuracy are presented in Fig. 6. The error bars represent the 95% confidence

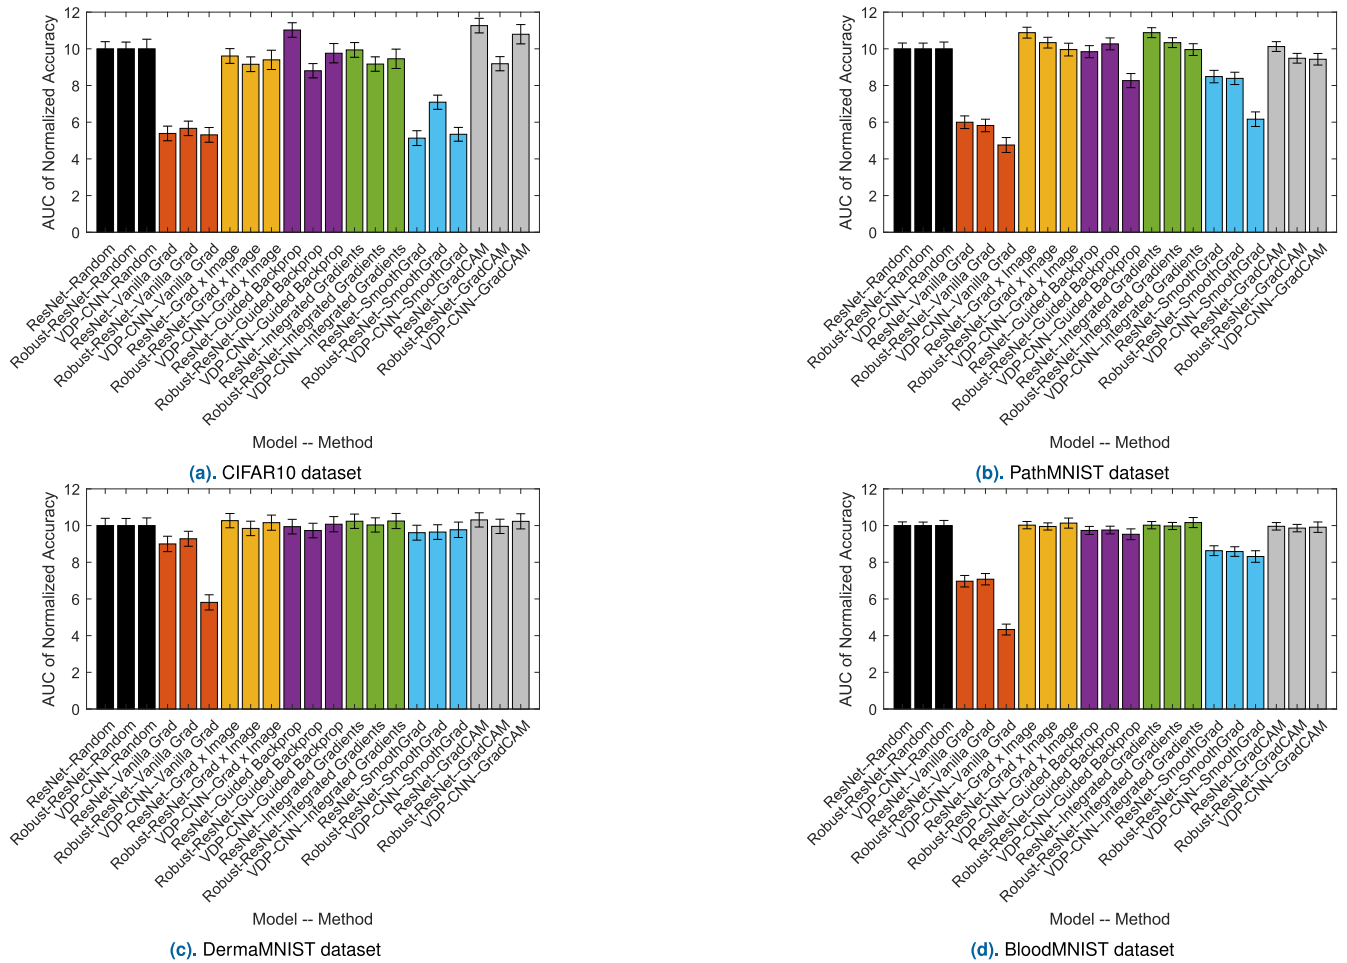
interval in both figures. We also performed p-value calculations for these results and determined that all had a p-value of less than 0.0001, indicating that the results are statistically significant.



**FIGURE 8.** Results of the EvalAttAI metric tested using DermaMNIST and BloodMNIST on three models, ResNet18, Robust-ResNet18 and VDP-CNN. The y-axis of all sub-figures shows average accuracy as attributions are added to the input image, which is captured by the x-axis. Lower accuracy values correspond to the attribution method being more faithful. The 95% confidence interval is shown in each plot using error bars. (a)-(c) and (g)-(i) Non-normalized accuracy. (d)-(f) and (j)-(l) Normalized accuracy using random as the baseline.

The results for the Deletion metric as presented in Fig. 5 ((a) and (c)) show each of the methods performing better than randomly removing pixels. We see this most clearly in Fig. 6a where AUC is presented. The one exception

to that is GradCAM, which performs worse (with statistical significance) than random. However, when we look at the Insertion results in Fig. 5 (b) and (d) and Fig. 6b, we see that GradCAM is the *only* method which performs better than



**FIGURE 9. Faithfulness results for EvalAttAI evaluation. Bar plots of the area under the curve (AUC) for the normalized CIFAR10, PathMNIST, DermaMNIST and BloodMNIST data (from MedMNIST dataset [39], [40]) results shown in Figs. 7 and 8. Both subfigures include the 95% confidence interval. The lower the bar, the more faithful the explanations.**

random (with statistical significance). In fact, when ranking the performance of the methods, we find that both metrics show contradictory behavior. Various attribution methods are ranked according to their performance and presented in Table 1. This table clearly shows that the methods that are ranked best by the Deletion metric may be the worst when evaluated using the Insertion metric and vice versa.

When further examining the results of Table 1, it should be noted that the peculiar behavior of GradCAM is likely because the method generates an attribution map of the same resolution as the convolutional layer that it is visualizing, which is then scaled up to the size of the input image. We utilized the unmodified GradCAM maps, although it is a widespread practice in research to apply a blur effect to the image, resulting in a more refined map. To keep the analysis fair and since we did not apply blur to any other attributions, we did not apply blur to GradCAM. Another likely reason for this method’s behavior is that unlike all other methods tested in this paper, GradCAM does not backpropagate all the way to the input when generating its attributions, resulting in a significant loss of information from the early layers of the

model. However, all other attribution methods tested in this paper do backpropagate all the way back to the input.

### B. EvalAttAI

The goal of this second experiment is two-fold. The first is to test the consistency and validity of the proposed EvalAttAI method. The second is to test whether the proposed methods capture the relationship between the faithfulness of the attribution method and the robustness of the model. All results with the EvalAttAI approach are presented in Figs. 7, 8 and 9.

In Figs. 7 and 8, we note the slope of the accuracy drop for each attribution method. The attribution methods that cause the steepest drop, and therefore lowest AUC, are deemed to be the most faithful. The best performing methods can be clearly identified when looking at Fig. 9. We can see that the Vanilla Gradient and SmoothGrad are the only two that consistently perform better than random baseline, with statistical significance across all models. Other attribution methods, including Grad x Image, Guided Backprop, Integrated Gradients, and GradCAM do not perform well as compared to the random

**TABLE 1.** Rankings of various attribution methods for faithfulness evaluation based on AUC using CIFAR10 dataset and three different models.

	ResNet	Robust-ResNet	VDP-CNN
<b>Deletion</b>			
1	IG	GBP	IG
2	SG	IG	VG
3	IxG	SG	GBP
4	GBP	IxG	SG
5	VG	GC	IxG
6	GC	VG	GC
<b>Insertion</b>			
1	GC	GC	GC
2	VG	VG	IxG
3	SG	IxG	VG
4	GBP	SG	SG
5	IxG	GBP	IG
6	IG	IG	GBP
<b>EvalAttAI</b>			
1	SG	VG	VG
2	VG	SG	SG
3	IxG	GBP	GBP
4	IG	IxG	IxG
5	GBP	IG	IG
6	GC	GC	GC

Method names are represented as follows: IG = Integrated Gradient, SG = SmoothGrad, IxG = Input x Grad, GBP = Guided Backprop, VG = Vanilla Gradient, GC = GradCam.

baseline. We also observe that EvalAttAI produces consistent results for each attribution method across all models in most of the cases. We also calculated the p-values for all EvalAttAI results and determined that  $p < 0.0001$  for all the results.

## V. DISCUSSION

We aimed to answer whether robust neural networks were more explainable, considering the potential usefulness of robust models and their visual explanations in medical imaging [7]. However, the question is difficult to answer due to three interrelated phenomena that need consideration. The first one is the robustness of the machine learning models, which is difficult to define and quantify [14], [16], [17], [18]. The second deals with various explainability methods and their internal complexities [7]. The third is related to metrics or measures employed to evaluate the plausibility and faithfulness of these explainability methods that try to explain the behavior of neural networks [7]. These three aspects of explainability research in deep neural networks are intimately intertwined. It may be challenging to disentangle these three (i.e., the robustness of models, internal dynamics of explainability methods, and metrics to evaluate this relationship) to understand the underlying relationship between robustness and explainability.

We started by limiting ourselves to two types of robust neural networks, (1) ResNet models trained using noisy datasets and (2) Bayesian deep neural networks trained using VDP technique. We acknowledge that the robustness of deep neural networks can be defined in many different ways, and consequently, many types of robust models can be built. However, we argue that these two methods represent a large class of robust models [6], [14], [15], [16], [17], [47]. On the other hand, for the explainability methods, we restricted ourselves to well-known gradient-based methods that are routinely used in image applications [7]. There are many different approaches to building explanations for elucidating the behavior of deep neural networks on test datasets. However, gradient-based methods represent a significantly large class of explainability methods. Finally, we wanted to use Insertion and Deletion metrics to evaluate the “goodness” or faithfulness of various attribution methods and quantify the explainability of robust neural networks. However, given the contradictory results produced by these two metrics due to their inherent nature, we proposed a new metric, EvalAttAI, to evaluate attribution methods.

Recently, Nourelahi et al. analyzed the faithfulness of various attribution methods for robust and non-robust CNNs [9]. The authors used Deletion and Insertion metrics for the evaluation of faithfulness. The faithfulness results appear to show methods and models which perform best on Deletion, perform the worst on Insertion, and vice versa [9]. We get similar results. Based on the working principles of these metrics (refer to Fig. 2), the logical conclusion would be that both methods might quantify different things in an attribution method. Perhaps, Insertion and Deletion may not be capturing any helpful information, as evident in Figs. 5 and 6. The contradictory findings for Deletion and Insertion are likely the result of an error introduced by removing and replacing the pixels, since the machine learning model is not trained on image data containing modified features (considering each pixel as a feature). We argue that given such a discrepancy, Deletion and Insertion may not be reliable metrics for evaluating faithfulness. The proposed EvalAttAI introduces perturbations smoothly and continuously, thus, avoiding abrupt changes in feature (pixel) values. These continuous and smooth changes are controlled using the  $\epsilon$  parameter as defined in Eq. 1.

A summary of our results is presented in Fig. 9 which compares three models (ResNet, Robust-ResNet, and Bayesian VDP-CNN) for four datasets (three of medical images and one of natural images), six attribution methods and a random baseline. In all sub-figures, the y-axes present AUC numbers calculated using test accuracy values for the proposed EvalAttAI metric and error bars represent 95% confidence interval. We do not observe any significant trend showing that any model (among ResNet, Robust-ResNet, or VDP-CNN) is more explainable than others across all tested attribution methods and datasets. For some attribution methods, robust models are more explainable, but not for others. On the other hand, we note that the Vanilla Gradient (orange color



bars) and SmoothGrad (blue color bars) consistently perform better than all other attribution methods. Amongst these two attribution methods, we note that Bayesian neural networks (VDP-CNN) are more explainable. Based on these results, we can conclude that Bayesian CNNs (trained using the VDP technique [14]) are more explainable than the standard and robustly trained neural networks when attributions are generated using Vanilla Gradient.

We also observed that the Vanilla Gradient consistently performed better than all other methods on our evaluation metric. We argue that this is expected since this method works directly on the neural network without any alterations or modifications.

These findings expand upon the best practices presented in our previous work [7], where we discussed considerations for the researchers when choosing an attribution method. We also discussed how robustness could play a significant role in generating plausible-looking explanations that may not be faithful. Based on this work, we suggest that the Vanilla Gradient should be used as the primary method of generating attributions. Instead of developing new ways to create visually appealing explanations using various operations, the community should focus on improving the robustness of the machine learning models using Bayesian approaches or other training methods.

## VI. CONCLUSION

In this work, we introduce a new metric for testing the faithfulness of attribution methods while showing the inconsistency and unreliability of the current state-of-the-art approaches. Our experiments are performed on both natural and medical image datasets. Our proposed faithfulness evaluation metric, EvalAttAI, shows consistent results. Our evaluation found that the Vanilla Gradient and SmoothGrad performed consistently better than all other attribution methods. Explanations from more complex attribution methods might appear convincing, but they should be approached with caution. In many instances, these methods are not faithful and can create a false sense of confidence in the model's predictions. Considering the limitations of current attribution methods, a meaningful future direction for the machine learning community to prioritize would be the development of models that are not only inherently explainable but also accurate, robust, efficient, scalable and versatile to improve interpretability, rather than creating more post-hoc attribution methods.

We could not find compelling evidence that all robust models are more explainable across the board. However, our experiments consistently show that Bayesian CNNs (trained using the VDP framework) were more explainable than all other models when used with the best performing attribution method (Vanilla Gradient). For this reason, we recommend that future research be focused on creating models with more robust architecture and training methods, rather than developing more complex post-hoc methods of generating attributions.

## REFERENCES

- [1] R. Confalonieri, L. Coba, B. Wagner, and T. R. Besold, "A historical perspective of explainable artificial intelligence," *WIREs Data Mining Knowl. Discovery*, vol. 11, no. 1, p. e1391, 2021.
- [2] D. Minh, H. X. Wang, Y. F. Li, and T. N. Nguyen, "Explainable artificial intelligence: A comprehensive review," *Artif. Intell. Rev.*, vol. 55, no. 5, pp. 3503–3568, Jun. 2022.
- [3] J. R. Epifano, R. P. Ramachandran, S. Patel, and G. Rasool, "Towards an explainable mortality prediction model," in *Proc. IEEE 30th Int. Workshop Mach. Learn. Signal Process. (MLSP)*, Sep. 2020, pp. 1–6.
- [4] Y. Zhang, Y. Weng, and J. Lund, "Applications of explainable artificial intelligence in diagnosis and surgery," *Diagnostics*, vol. 12, no. 2, p. 237, Jan. 2022.
- [5] B. H. M. van der Velden, H. J. Kuijf, K. G. A. Gilhuijs, and M. A. Viergever, "Explainable artificial intelligence (XAI) in deep learning-based medical image analysis," *Med. Image Anal.*, vol. 79, Jul. 2022, Art. no. 102470.
- [6] S. Ahmed, D. Dera, S. U. Hassan, N. Bouaynaya, and G. Rasool, "Failure detection in deep neural networks for medical imaging," *Frontiers Med. Technol.*, vol. 4, Jul. 2022, Art. no. 919046.
- [7] I. E. Nielsen, D. Dera, G. Rasool, R. P. Ramachandran, and N. C. Bouaynaya, "Robust explainability: A tutorial on gradient-based attribution methods for deep neural networks," *IEEE Signal Process. Mag.*, vol. 39, no. 4, pp. 73–84, Jul. 2022.
- [8] S. Hooker, D. Erhan, P.-J. Kindermans, and B. Kim, "A benchmark for interpretability methods in deep neural networks," in *Proc. 33rd Int. Conf. Neural Inf. Process. Syst. (NIPS)*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds. Red Hook, NY, USA: Curran Associates, 2019, pp. 9737–9748. [Online]. Available: <https://proceedings.neurips.cc/paper/2019/file/fe4b8556000d0f0cae99daa5c5c5a410-Paper.pdf>
- [9] M. Nourelahi, L. Kotthoff, P. Chen, and A. Nguyen, "How explainable are adversarially-robust CNNs?" 2022, *arXiv:2205.13042*.
- [10] A. Mamalakis, E. A. Barnes, and I. Ebert-Uphoff, "Investigating the fidelity of explainable artificial intelligence methods for applications of convolutional neural networks in geoscience," *Artif. Intell. Earth Syst.*, vol. 1, no. 4, 2022, Art. no. e220012. [Online]. Available: <https://journals.ametsoc.org/view/journals/aies/1/4/AIES-D-22-0012.1.xml>
- [11] N. Hama, M. Mase, and A. B. Owen, "Deletion and insertion tests in regression models," 2022, *arXiv:2205.12423*.
- [12] J. Zhou, A. H. Gandomi, F. Chen, and A. Holzinger, "Evaluating the quality of machine learning explanations: A survey on methods and metrics," *Electronics*, vol. 10, no. 5, p. 593, Mar. 2021.
- [13] H. Phan and A. Nguyen, "DeepFace-EMD: Re-ranking using patch-wise earth mover's distance improves out-of-distribution face identification," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 20227–20237.
- [14] D. Dera, N. C. Bouaynaya, G. Rasool, R. Shterenberg, and H. M. Fathallah-Shaykh, "PremiUm-CNN: Propagating uncertainty towards robust convolutional neural networks," *IEEE Trans. Signal Process.*, vol. 69, pp. 4669–4684, 2021.
- [15] G. Carannante, D. Dera, N. C. Bouaynaya, H. M. Fathallah-Shaykh, and G. Rasool, "SUPER-Net: Trustworthy medical image segmentation with uncertainty propagation in encoder-decoder networks," 2021, *arXiv:2111.05978*.
- [16] D. Dera, G. Rasool, and N. Bouaynaya, "Extended variational inference for propagating uncertainty in convolutional neural networks," in *Proc. IEEE 29th Int. Workshop Mach. Learn. Signal Process. (MLSP)*, Oct. 2019, pp. 1–6.
- [17] D. Dera, G. Rasool, N. C. Bouaynaya, A. Eichen, S. Shanko, J. Cammerata, and S. Arnold, "Bayes-SAR net: Robust SAR image classification with uncertainty estimation using Bayesian convolutional neural network," in *Proc. IEEE Int. Radar Conf. (RADAR)*, Apr. 2020, pp. 362–367.
- [18] A. Waqas, H. Farooq, N. C. Bouaynaya, and G. Rasool, "Exploring robust architectures for deep artificial neural networks," *Commun. Eng.*, vol. 1, no. 1, p. 46, Dec. 2022.
- [19] A. Waqas, D. Dera, G. Rasool, N. C. Bouaynaya, and H. M. Fathallah-Shaykh, "Brain tumor segmentation and surveillance with deep artificial neural networks," in *Deep Learning for Biomedical Data Analysis: Techniques, Approaches, and Applications*. Springer, 2021, pp. 311–350.
- [20] D. Tsipras, S. Santurkar, L. Engstrom, A. Turner, and A. Madry, "Robustness may be at odds with accuracy," in *Proc. Int. Conf. Learn. Represent.*, 2019, pp. 6–9. [Online]. Available: <https://openreview.net/forum?id=SyxAb30cY7>

- [21] M. Ancona, E. Ceolini, C. Öztireli, and M. Gross, "Gradient-based attribution methods," in *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, W. Samek, G. Montavon, A. Vedaldi, L. K. Hansen, and K.-R. Müller, Eds. Cham, Switzerland: Springer, 2019, pp. 169–191, doi: 10.1007/978-3-030-28954-6\_9.
- [22] M. Nauta, J. Trienes, S. Pathak, E. Nguyen, M. Peters, Y. Schmitt, J. Schlötterer, M. van Keulen, and C. Seifert, "From anecdotal evidence to quantitative evaluation methods: A systematic review on evaluating explainable AI," 2022, *arXiv:2201.08164*.
- [23] J. Zhang, S. A. Bargal, Z. Lin, J. Brandt, X. Shen, and S. Sclaroff, "Top-down neural attention by excitation backprop," *Int. J. Comput. Vis.*, vol. 126, no. 10, pp. 1084–1102, Oct. 2018.
- [24] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2921–2929.
- [25] L. Raatikainen and E. Rahtu, "The weighting game: Evaluating quality of explainability methods," 2022, *arXiv:2208.06175*.
- [26] T. Zhang and Z. Zhu, "Interpreting adversarially trained convolutional neural networks," in *Proc. 36th Int. Conf. Mach. Learn.*, in Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97, Jun. 2019, pp. 7502–7511. <https://proceedings.mlr.press/v97/zhang19s.html>
- [27] P. Chalasani, J. Chen, A. R. Chowdhury, X. Wu, and S. Jha, "Concise explanations of neural networks using adversarial training," in *Proc. 37th Int. Conf. Mach. Learn.*, in Proceedings of Machine Learning Research, vol. 119, H. Daumé III and A. Singh, Eds., Jul. 2020, pp. 1383–1391. [Online]. Available: <https://proceedings.mlr.press/v119/chalasani20a.html>
- [28] A. F. Markus, J. A. Kors, and P. R. Rijnbeek, "The role of explainability in creating trustworthy artificial intelligence for health care: A comprehensive survey of the terminology, design choices, and evaluation strategies," *J. Biomed. Informat.*, vol. 113, Jan. 2021, Art. no. 103655.
- [29] A.-P. Nguyen and M. R. Martínez, "On quantitative aspects of model interpretability," 2020, *arXiv:2007.07584*.
- [30] M. Velmurugan, C. Ouyang, C. Moreira, and R. Sindhgatta, "Evaluating fidelity of explainable methods for predictive process analytics," in *Intelligent Information Systems*, S. Nurcan and A. Korthaus, Eds. Cham, Switzerland: Springer, 2021, pp. 64–72.
- [31] C.-K. Yeh, C.-Y. Hsieh, A. Suggala, D. I. Inouye, and P. K. Ravikumar, "On the (in)fidelity and sensitivity of explanations," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds. Red Hook, NY, USA: Curran Associates, 2019. [Online]. Available: <https://proceedings.neurips.cc/paper/2019/file/a7471fd77b3435276507cc8f2dc2569-Paper.pdf>
- [32] Y. Ge, S. Liu, Z. Li, S. Xu, S. Geng, Y. Li, J. Tan, F. Sun, and Y. Zhang, "Counterfactual evaluation for explainable AI," 2021, *arXiv:2109.01962*.
- [33] V. Petsiuk, A. Das, and K. Saenko, "RISE: Randomized input sampling for explanation of black-box models," 2018, *arXiv:1806.07421*.
- [34] W. Samek, A. Binder, G. Montavon, S. Lapuschkin, and K.-R. Müller, "Evaluating the visualization of what a deep neural network has learned," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 28, no. 11, pp. 2660–2673, Nov. 2017.
- [35] M. N. Vu, T. D. Nguyen, N. Phan, R. Gera, and M. T. Thai, "C-eval: A unified metric to evaluate feature-based explanations via perturbation," in *Proc. IEEE Int. Conf. Big Data (Big Data)*, Dec. 2021, pp. 927–937.
- [36] K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2014, pp. 1–8.
- [37] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [38] A. Krizhevsky, "Learning multiple layers of features from tiny images," M.S. thesis, Dept. Comput. Sci., Univ. Toronto, Toronto, ON, Canada, 2009.
- [39] J. Yang, R. Shi, and B. Ni, "MedMNIST classification decathlon: A lightweight AutoML benchmark for medical image analysis," in *Proc. IEEE 18th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2021, pp. 191–195.
- [40] J. Yang, R. Shi, D. Wei, Z. Liu, L. Zhao, B. Ke, H. Pfister, and B. Ni, "MedMNIST v2—A large-scale lightweight benchmark for 2D and 3D biomedical image classification," *Sci. Data*, vol. 10, no. 1, p. 41, Jan. 2023.
- [41] J. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller, "Striving for simplicity: The all convolutional net," in *Proc. ICLR*, 2015, pp. 7–14. [Online]. Available: <http://lmb.informatik.uni-freiburg.de/Publications/2015/DB15a>
- [42] M. Sundararajan, A. Taly, and Q. Yan, "Axiomatic attribution for deep networks," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 3319–3328.
- [43] D. Smilkov, N. Thorat, B. Kim, F. Viégas, and M. Wattenberg, "SmoothGrad: Removing noise by adding noise," 2017, *arXiv:1706.03825*.
- [44] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 618–626.
- [45] N. Kokhlikyan, V. Miglani, M. Martin, E. Wang, B. Alsallakh, J. Reynolds, A. Melnikov, N. Kliushkina, C. Araya, S. Yan, and O. Reblitz-Richardson, "Captum: A unified and generic model interpretability library for PyTorch," 2020, *arXiv:2009.07896*.
- [46] N. Kokhlikyan, V. Miglani, M. Martin, E. Wang, J. Reynolds, A. Melnikov, N. Lunova, and O. Reblitz-Richardson, "Pytorch Captum," GitHub Repository, Captum Developed Meta Open Source, Menlo Park, CA, USA, Tech. Rep., Captum v0.6.0 Release, 2019. [Online]. Available: <https://github.com/pytorch/captum> and <https://arxiv.org/pdf/2009.07896.pdf>
- [47] G. Carannante, D. Dera, G. Rasool, and N. C. Bouaynaya, "Self-compression in Bayesian neural networks," in *Proc. IEEE 30th Int. Workshop Mach. Learn. Signal Process. (MLSP)*, Sep. 2020, pp. 1–6.



**IAN E. NIELSEN** received the B.S. degree (cum laude) in electrical and computer engineering from Rowan University, in 2020, where he is currently pursuing the Ph.D. degree in electrical and computer engineering. He is a Research Assistant and a Teaching Assistant. His current research was focused on robust machine learning and how it relates to explainable artificial intelligence. His work mainly focuses on computer vision and cancer diagnosis tasks. He conducts research as a part of the Rowan's Artificial Intelligence Laboratory (RAIL). He currently coordinates a group of 19 undergraduate student researchers alongside. He through the Rowan University Engineering Clinic Program. Since the start of his graduate education, he has taught machine learning using PyTorch and Python to dozens of students through this program. His tutorial on robust explainability was recently published in *IEEE Signal Processing Magazine*. He was a recipient of the GAANN Teaching Fellowship under the U.S. Department of Education, as of January 2021.



**RAVI P. RAMACHANDRAN** (Senior Member, IEEE) received the B.Eng. degree (Hons.) from Concordia University, in 1984, and the M.Eng. and Ph.D. degrees from McGill University, in 1986 and 1990, respectively. From October 1990 to December 1992, he was with the Speech Research Department, AT&T Bell Laboratories. From January 1993 to August 1997, he was a Research Assistant Professor with Rutgers University. He was also a Senior Speech Scientist with T-Netix, from July 1996 to August 1997. Since September 1997, he has been with the Department of Electrical and Computer Engineering, Rowan University, where he has been a Professor, since September 2006. He was a Consultant with T-Netix, Avenir Inc., Motorola, and FocalCool. His research interests are in digital signal processing, speech processing, biometrics, pattern recognition, machine learning, and filter design. Since May 2002, he has been on the Digital Signal Processing Technical Committee for the IEEE Circuits and Systems Society. Since May 2012, he has been on the Education and Outreach Technical Committee for the IEEE Circuits and Systems Society. From September 2002 to September 2005, he was an Associate Editor for the IEEE Transactions on Speech and Audio Processing and was on the Speech Technical Committee for the IEEE Signal Processing Society. From September 2000 to December 2015, he was on the Editorial Board of the *IEEE Circuits and Systems Magazine*. He is currently an Associate Editor of *Circuits, Systems and Signal Processing* journal.



**NIDHAL BOUAYNAYA** (Member, IEEE) received the Ph.D. degree in electrical and computer engineering (ECE) and the M.S. degree in pure mathematics from the University of Illinois at Chicago.

She is a Professor of ECE and the Director of the Rowan's Artificial Intelligence Laboratory (RAIL). She is currently the Associate Dean of Research and Graduate Studies with the Henry M. Rowan College of Engineering. Previously, she was a Faculty Member with the University of Arkansas at Little Rock. She has coauthored more than 100 refereed journal articles, book chapters, and conference proceedings, such as IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, IEEE SIGNAL PROCESSING LETTERS, *IEEE Signal Processing Magazine*, and *PLOS Medicine*. Her research interests are in big data analytics, machine learning, artificial intelligence, and mathematical optimization.

Dr. Bouaynaya won numerous best paper awards, the most recent was at the 2019 IEEE International Workshop on Machine Learning for Signal Processing. She is also the winner of the top algorithm at the 2016 Multinomial Brain Tumor Segmentation Challenge (BRATS). She has been honored with numerous research and teaching awards, including the Rowan Research Achievement Award, in 2017, and the University of Arkansas at Little Rock Faculty Excellence Award in research. Her research was primarily funded by the National Science Foundation (NSF CCF, NSF ACI, NSF DUE, NSF I-Corps, NSF ECCS, NSF OAC, and NSF HRD), the National Institutes of Health (NIH), the U.S. Department of Education (USED), the New Jersey Department of Transportation (NJ DoT), the U.S. Department of Agriculture (USDA), the Federal Aviation Administration (FAA), Lockheed Martin Inc., and other industry. She is also interested in entrepreneurial endeavors. In 2017, she co-founded and is the Chief Executive Officer (CEO) of MRIMATH, LLC, a start-up company that uses artificial intelligence to improve patient oncology outcomes and treatment response. MRIMath is funded by the NIH SBIR Program.

**HASSAN M. FATHALLAH-SHAYKH**, photograph and biography not available at the time of publication.



**GHULAM RASOOL** (Member, IEEE) received the B.S. degree in mechanical engineering from the National University of Sciences and Technology (NUST), Pakistan, in 2000, the M.S. degree in computer engineering from the Center for Advanced Studies in Engineering (CASE), Pakistan, in 2010, and the Ph.D. degree in systems engineering from the University of Arkansas at Little Rock, in 2014. He was a Postdoctoral Fellow with the Rehabilitation Institute of Chicago and

Northwestern University, from 2014 to 2016. Before joining Moffitt, he was an Assistant Professor with the Department of Electrical and Computer Engineering, Rowan University. He is an Assistant Member with the Department of Machine Learning, H. Lee Moffitt Cancer Center and the Research Institute, Tampa, FL, USA. His current research focuses on building trustworthy multimodal machine learning and artificial intelligence model for cancer diagnosis, treatment planning, and risk assessment. His research efforts are currently funded by two National Science Foundation (NSF) awards. His research was supported by the National Institute of Health (NIH), the U.S. Department of Education, NSF, the New Jersey Health Foundation (NJHF), Google, NVIDIA, and Lockheed Martin, Inc. His work on Bayesian machine learning won the Best Student Award at the 2019 IEEE Machine Learning for Signal Processing Workshop.

• • •