

Relative Attributing Propagation: Interpreting the Comparative Contributions of Individual Units in Deep Neural Networks

Woo-Jeoung Nam,¹ Shir Gur,³ Jaesik Choi,⁵ Lior Wolf,^{3,4} Seong-Whan Lee^{1,2*}

¹Department of Computer and Radio Communications Engineering, Korea University

²Department of Artificial Intelligence, Korea University

³The School of Computer Science, Tel Aviv University

⁴Facebook AI Research

⁵Graduate School of Artificial Intelligence, KAIST

Abstract

As Deep Neural Networks (DNNs) have demonstrated superhuman performance in a variety of fields, there is an increasing interest in understanding the complex internal mechanisms of DNNs. In this paper, we propose Relative Attributing Propagation (RAP), which decomposes the output predictions of DNNs with a new perspective of separating the relevant (positive) and irrelevant (negative) attributions according to the relative influence between the layers. The relevance of each neuron is identified with respect to its degree of contribution, separated into positive and negative, while preserving the conservation rule. Considering the relevance assigned to neurons in terms of relative priority, RAP allows each neuron to be assigned with a bi-polar importance score concerning the output: from highly relevant to highly irrelevant. Therefore, our method makes it possible to interpret DNNs with much clearer and attentive visualizations of the separated attributions than the conventional explaining methods. To verify that the attributions propagated by RAP correctly account for each meaning, we utilize the evaluation metrics: (i) Outside-inside relevance ratio, (ii) Segmentation mIOU and (iii) Region perturbation. In all experiments and metrics, we present a sizable gap in comparison to the existing literature.

Introduction

Despite the impressive performance, the adoption of Deep Neural Networks (DNNs) is sometimes hindered by a transparency issue that arises from the complex internal structure of DNNs. Many studies have recently attempted to resolve the lack of transparency in DNNs. The attributing methods (Bach et al. 2015; Montavon et al. 2017; Kindermans et al. 2017; Shrikumar, Greenside, and Kundaje 2017; Montavon, Samek, and Müller 2018; Lapuschkin et al. 2019; Lundberg and Lee 2017; Sundararajan, Taly, and Yan 2017) reveal the significant factors of the input in making decisions by assigning a relevance score to the input layer.

To consider the positive and negative contributions of each image location to the output of a DNN, (Bach et al. 2015) introduced the layer-wise relevance propagation rule,




Figure 1: Comparison of the conventional explaining methods and RAP applied to VGG-16. In the previous methods, the attributions are similarly distributed across the entire image. Our RAP clearly distinguishes relevant (red) and irrelevant pixels (blue), placing the relevant attributions on the object, and the irrelevant ones on the background.

which propagates the relevance from the prediction. However, propagating the positive and negative relevance without considering the amount and direction of the contribution, may lead to defective interpretation. It is required to clarify the actual influences of individual units to the output, since the components of the complex inner structure shift and switch the conveyance of value. Furthermore, the relevance of each neuron is highly dependent on the absolute amount of contribution, resulting in both positive and negative relevance types to be correlated.

In this paper, we propose a new perspective for interpreting the relevance of each neuron, accounting for each neuron’s influence among connected neighbors and allocating it with the relative importance. The main idea of this paper is changing the perspective of the relevance, from the sign of contribution to the influence among the neurons. Our method redistributes the relevance by changing the priority and rearranging it into positive and negative while preserving the conservation. This way, the relevance is assigned to each neuron directionally in line with the degree of importance to the output.

Fig. 1 illustrates the comparison between RAP and conventional methods. While previous work considers directionality based on the sign of the neuron’s contribution, which leads to a similar distribution of the positive and negative attributions, our method assigns the relevance according to the importance of neuron, which is highly focused on the object. The main contributions of this work are as follows:

*Corresponding author: Seong-Whan Lee
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

- We propose relative attributing propagation (RAP), a method for attributing the positive and negative relevance to each neuron, according to its relative influence among the neurons. We address the phenomenon of the relevance dependency, which is highly dependant on the amount of neuron contribution and present the necessity of the new perspective to approach the relevance from the priority. We also prevent the risk of degeneracy during propagation by setting the criterion of separation according to the actual contribution between the intermediate layers.
- We apply the Intersection of Union, Outside-Inside relevance ratio (Lapuschkin et al. 2016) and region perturbation (Samek et al. 2017) to assess whether the propagated attributions are meaningful. The evaluation shows that attributions from RAP provide a high objectness score with a clear separation of irrelevant regions, compared to the other explaining methods as well as to the literature of objectness methods.

Related Work

There has been many recent studies on understanding of what a DNN model has learned. From the standpoint of interpreting a DNN model, the manner in which a DNN works can be visualized by maximizing the activation of hidden layers (Erhan et al. 2009) or generating salient feature maps (Dabkowski and Gal 2017; Simonyan, Vedaldi, and Zisserman 2013; Mahendran and Vedaldi 2016; Zhou et al. 2016; Zeiler and Fergus 2014; Zhou et al. 2018). (Foerster et al. 2017) introduced the input switched affine network, which can decompose the contributions of previous characters to the current prediction, and (Koh and Liang 2017) proposed the influence function to understand model behavior, debug models, detect dataset errors, and even create visually indistinguishable training-set attacks. (Ribeiro, Singh, and Guestrin 2016) proposed LIME, an algorithm that explains the predictions of the classifier, by learning an interpretable model locally around the prediction.

From the standpoint of explaining the decision of a DNN, the contributions of the input are propagated backward, resulting in a redistribution of relevance in the pixel space. Sensitivity analysis visualizes the sensitivities of input images classified by a DNN while explaining the factors that reduce/increase the evidence for the predicted results (Baehrens et al. 2010). (Zeiler and Fergus 2014) proposed a deconvolution method to identify the patterns of a predicted input image from a DNN. Layer-wise relevance propagation (LRP) (Bach et al. 2015) was introduced to backpropagate the relevance, also called as attribution, by making the network output become fully redistributed throughout the layers of a DNN. (Samek et al. 2017) showed that the LRP algorithm qualitatively and quantitatively provides a better explanation than do either the sensitivity-based approach or the deconvolution method.

Guided BackProp (Springenberg et al. 2014) and Integrated Gradients (Sundararajan, Taly, and Yan 2017) each compute the single and average partial derivatives of the output to attribute the prediction of a DNN. Deep Taylor Decomposition (Montavon et al. 2017) is an extension of

LRP for interpreting the decision of a DNN, by decomposing the activation of a neuron in terms of the contributions from its inputs. DeepLIFT (Shrikumar, Greenside, and Kundaje 2017) decomposes the output prediction by assigning the differences of contribution scores between the activation of each neuron to its reference activation. (Ancona et al. 2018) approached the problem of the attribution value from a theoretical perspective and formally proved the conditions of equivalence and approximation between four attribution methods: Guided Backprop (Springenberg et al. 2014), Integrated Gradients (Sundararajan, Taly, and Yan 2017), LRP and DeepLIFT. (Lundberg and Lee 2017) proposed SHAP explaining methods with unifying the conventional explaining methods and approximate the shapley value.

However, there are no studies which analyze the problem of ambiguous visualization in dealing with negative relevance. We bring out the fundamental causes of this problem and address the solution to handle the priority of neurons, resulting in the clear separation of the (ir)relevant objects.

Background

Notations Throughout this paper, the letter $f(x)$ is used to denote the value of the network output before passing through the softmax layer for input x . R represents the value of $f(x)$ corresponding to the prediction class, which constitutes the input relevance for the attributing procedure. A neuron j in the layer $l + 1$ receives the value $z_{ij}^{(l+1)}$ from a neuron i in the layer l , which is obtained by multiplying the activation of the neuron in layer l , denoted $m_i^{(l)}$, and the weight $w_{ij}^{(l,l+1)}$. These contributions are summed over the relevant neurons to obtain $z_j^{(l+1)}$, which becomes $m_j^{(l+1)}$ after adding the bias $b_j^{(l,l+1)}$ and applying the activation function $a(\cdot)$.

$$z_{ij}^{(l+1)} = m_i^{(l)} w_{ij}^{(l,l+1)}, z_j^{(l+1)} = \sum_i z_{ij}^{(l+1)},$$

$$m_j^{(l+1)} = a(z_j^{(l+1)} + b_j^{(l,l+1)})$$
(1)

We consider the positive and negative parts of the contributions: $z_{ij} = z_{ij}^+ + z_{ij}^-$, where $z_{ij}^+ = \max(z_{ij}, 0)$ and $z_{ij}^- = \min(z_{ij}, 0)$.

Layerwise Relevance Propagation (LRP) LRP finds the parts with high relevance in the input, by propagating the relevance R from the output to the input. The algorithm is based on the conservation principle, which maintains the relevance in each layer. Let $R_i^{(l)}$ denote the relevance of a neuron i in a layer l and $R_j^{(l+1)}$ is associated with a neuron j of the layer $l + 1$, this conservation takes the form:

$$\sum_i R_i^{(l)} = \sum_j R_j^{(l+1)}$$
(2)

(Bach et al. 2015) introduced two relevance propagation rules that satisfy Eq. 2. The first rule, called LRP- ϵ , is defined as

$$R_i^{(l)} = \sum_j \frac{z_{ij}}{\sum_i z_{ij} + \epsilon} R_j^{(l+1)}$$
(3)

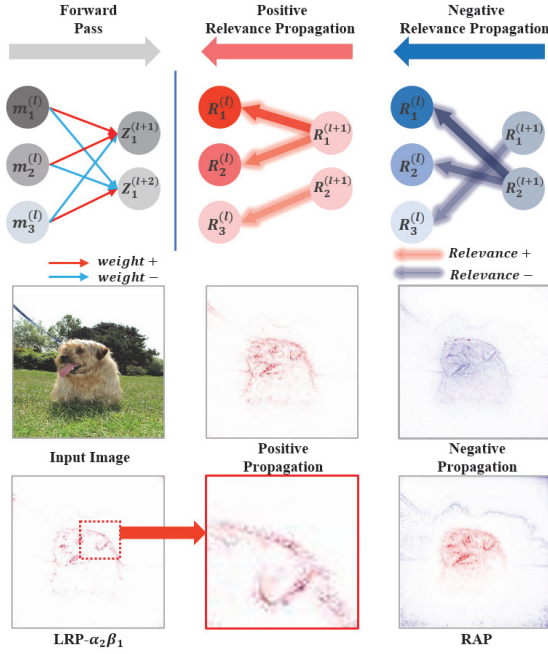


Figure 2: An illustration of the way positive and negative contributions are handled in LRP. See text for details.

In this rule, a neuron i in the layer l receives the relevance, according to their contribution to the activation of the neurons in the layer $l + 1$. The constant ϵ prevents the numerical instability for the case in which the denominator becomes zero. The second rule LRP- $\alpha\beta$ enforces the conservation principle, while separating between the positive and negative activations in the relevance propagation process.

$$R_i^{(l)} = \sum_j \left(\alpha \cdot \frac{z_{ij}^+}{\sum_i z_{ij}^+} - \beta \cdot \frac{z_{ij}^-}{\sum_i z_{ij}^-} \right) R_j^{(l+1)} \quad (4)$$

Recall that $z_{ij}^+ + z_{ij}^- = z_{ij}$. To maintain the total relevance, the parameters are chosen such that $\alpha - \beta = 1$. We refer to the part of the relevance that is multiplied by α (β) and which is related to the positive (negative) activations as the positive (negative) relevance.

The Shortcoming of Current Relevance Propagation Methods

To motivate our method, we employ a toy sample to understand how relevance is propagated in the current methods. While it is presented in the context of LRP- $\alpha\beta$, similar pathologies can be found in other methods, such as LRP- ϵ , integrated gradients, and pattern attribution. Fig. 2 presents an example of forward pass and backward relevance propagation between the two intermediate layers. For illustration purposes, the forward process does not include a bias term nor batch normalization, and all neurons in layer l are non-negative. For simplicity, the absolute values of all weights are identical. The darker color means the higher value of the neuron. The positive part of $R_1^{(l+1)}$ is propagated back to

the neurons $m_1^{(l)}$ and $m_2^{(l)}$ at the ratio of $\{m_1^{(l)}:m_2^{(l)}\}$. Similarly, the negative relevance of $m_1^{(l)}$ receives in the propagation the lion's share of the relevance of $R_2^{(l+1)}$, resulting in very low relevance value for $R_1^{(l)}$. When summing the high positive and the negative contributions of neuron $i = 1$ in Eq. 4, these cancel out. However, in terms of the amount of the contribution, it is clear that the relevance of this neuron should be high, since it plays a major role in the activations of layer $l + 1$.

Another related pathology is illustrated in the dog sample in Fig. 2, which further illustrates this phenomenon. The relevance is propagated recursively using Eq. 4 from the output layer to the input, obtaining the positive propagation image ($\alpha = 1, \beta = 0$). When doing the same process, but only for the negative value (similarly to $\alpha = 0, \beta = 1$), we obtain the negative propagation image. It seems that the same locations of the object receive both high positive relevance and high negative relevance. When these are combined, using Eq. 4, they tend to cancel each other. The combination of the positive and negative contributions is illustrated in the third image row, which demonstrated the results of the LRP- $\alpha_2\beta_1$ method (similar to previous work this notation refers to the LRP method with $\alpha = 2$ and $\beta = 1$). Many of the positive relevance values are canceled out by equally large negative values, except for specific locations in which one contribution dominates. However, these locations can be either positive or negative and appear in close proximity to each other, as the zoomed-in subfigure demonstrates.

In RAP, we consider the absolute contribution of each neuron and propagating that relevance according to a novel method we introduce in the next section. The result of our method is shown on the bottom right of Fig 2. As shown in the result, the positive and negative relevances appear in different parts of the image, corresponding to regions of high and low importance.

Relative Attributing Propagation

Motivated by the above issues, our goal is to separate the relatively (un)important neurons according to their influence across the layers. The method has three main steps: (i) absolute influence normalization, (ii) propagation of the relevance, and (iii) uniform shifting. Fig. 3 illustrates the three steps.

Absolute Influence Normalization We first propagate the relevance value to neuron i of the penultimate layer p , according to its actual contribution value $m_i w_{ij}$ from the prediction node j in the final layer q . Because the bias $b_j^{(p,q)}$ is a single value, it is possible to consider the relevance of the bias $b_j^{(p,q)}$ to the previous layer p , by increasing the contribution of each neuron.

$$R_i^{(p)} = \left(\sum_i z_{ij}^+ + \sum_i z_{ij}^- \right) * \frac{R_j^{(q)} + b_j^{(p,q)}}{R_j^{(q)}} \quad (5)$$

After applying the following, relevance values in the penultimate layer p are composed of both positives and negatives.

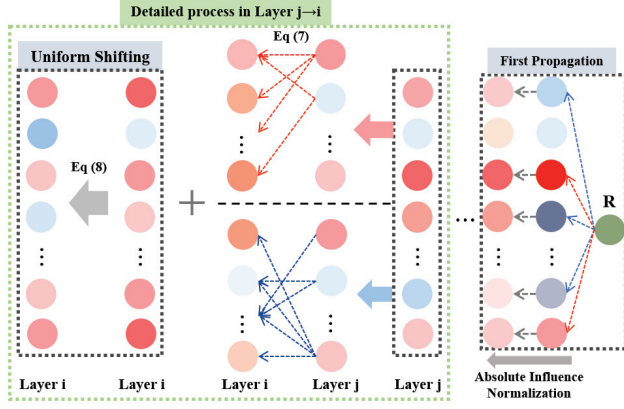


Figure 3: Overall structure of RAP algorithm.

Next, we normalize the entire value $R_i^{(p)}$ by the ratio of the absolute positive and negative values $|R_i^{(p)+}| : |R_i^{(p)-}|$.

$$R_i^{(p)} = |R_i^{(p)}| * \frac{\sum_i R_i^{(p)}}{\sum_i |R_i^{(p)}|} \quad (6)$$

$R_i^{(p)}$ is the new input relevance to the next propagation and all relevance values in layer p are distributed by their relative importance to the output layer, from most influenced to rarely influenced. The greater the neuron’s influence on the contribution, the more positive relevance it is assigned. Eq. 6 is only applied in the first relevance propagation process.

Criterion of Relevant Neuron After changing the relevance to the amount of absolute contribution each neuron has, i.e., when all relevance scores are positive, it is possible to propagate the relevance while maintaining a degree of relative influence. We then apply a uniform shifting to all activated neurons, which causes low influential neurons to have negative relevance.

Next, the relevance propagation redistributes the relevance to the next layer through the positive contributions $\mathcal{P} = \{i, j | z_{ij}^+ > 0\}$ and $\mathcal{N} = \{i, j | z_{ij}^- < 0\}$. It is possible to propagate the relevance of layer l through the formal case \mathcal{P} , which makes each relevance to be redistributed according to the degree of the positive influence. For the negative case, we apply the same procedure, considering the influence among the negatively contributed neuron. Here, the propagated relevance originally means the ratio between the whole negative contribution in the forward pass. We utilize this amount of the relevance to uniformly shift the whole activated neurons, which makes the neurons to be converted as negative, in the order in which they are close to zero. For each relevance value in the layer l , it is possible to compute the ratio of the positive and negative contribution with $\{\sum_j z_{ij}^+ : \sum_j z_{ij}^-\}$. However, when we compute the ratio with the original sign, the degeneracy problem has occurred because the absolute value of the numerator $\sum_j z_{ij}^+, \sum_j z_{ij}^-$ has a much larger value than the denominator $\sum_j z_{ij}^+ + \sum_j z_{ij}^-$. Therefore, the contribution ratio is computed after normalizing with the ab-

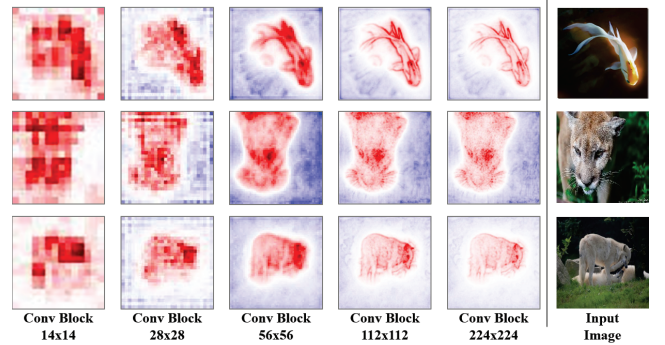


Figure 4: The visualization of the relevance map of the intermediate layers of the VGG-16 network.

solute value of each contribution.

$$\nu_j^{(l+1)} = R_j^{(l+1)} * \frac{\sum_i |z_{ij}^-|}{\sum_i (|z_{ij}^+| + |z_{ij}^-|)} \quad (7)$$

$$\bar{R}_{i \in \mathcal{P}, \mathcal{N}}^{(l)} = \sum_j \left(\frac{z_{ij}^+}{\sum_i (z_{ij}^+)} R_j^{(l+1)} + \frac{z_{ij}^-}{\sum_i (z_{ij}^-)} \nu_j^{(l+1)} \right)$$

After propagating the relevance to the negatively contributed neurons, all neurons in layer R_j receive the relevance value, according to the inner ratio of each contribution. However, relevance is not conserved, and there is an over-allocation in layer l in comparison to layer $l + 1$. This over-allocated relevance, which originally means the negative contribution, is utilized for separating the important (unimportant) neurons by uniformly shifting all activated neurons. Let $\Gamma = \{m_i^{(l)} \neq 0\}$ be the number of the activated neurons in layer i . We evenly subtract the mean over allocation over these neurons. This shifts some of the relevancy scores to the negative region. Specifically, the shift is given by

$$\Psi_i^l = \begin{cases} \sum_i \left(\bar{R}_{i \in \mathcal{N}}^{(l)} \right) * \frac{1}{\Gamma} & , m_i^{(l)} \text{ is activated} \\ 0 & , \text{otherwise} \end{cases} \quad (8)$$

In this case, the relevance of the neurons in both groups $i \in \mathcal{P}$ and $i \in \mathcal{N}$ becomes

$$R_i^{(l)} = \bar{R}_{i \in \mathcal{P} \cup \mathcal{N}}^{(l)} - \Psi_i^l \quad (9)$$

From Eq. 8,9, it is easy to verify that the relevance value is preserved as in Eq. 2. We emphasize that the criterion for separating the important and unimportant neurons is the amount of each contribution, not the direction. Since the activated parts are different between the feature maps of intermediate layers, we evenly subtract the same value to all activated neurons for not losing the important parts of each feature map. The negative input relevance for the next propagation indicates a relatively lower priority for the prediction result. Therefore, Eq. 7 propagate the negative relevance to the connected neurons, which contributed to the relatively unimportant neurons in positive and negative directions.

For the final relevance propagation to the input image layer, we utilized the Z^β rule (Bach et al. 2015) which is

commonly used for propagating to the input layer in methods derived from LRP.

$$R_i^{(0)} = \sum_j \left(\frac{x_i w_{ij} - l_i w_{ij}^+ - h_i w_{ij}^-}{\sum_i (x_i w_{ij} - l_i w_{ij}^+ - h_i w_{ij}^-)} R_j^{(1)} \right) \quad (10)$$

where x is the input image and $\{l, h\}$ denotes the minimum and maximum values of x .

We investigate the variation of the relevance maps during the propagating process. Fig. 4 shows the relevance map of activated neurons in the intermediate layers of VGG-16 net, where each pixel represents the sum of the relevance scores along the channel-wise axis. Here, we note that the intensity of the represented becomes lighter, since the size of the feature map is increased when passing the max-pooling layers, resulting in scattered to more neurons. As expected, the positive/negative maps change gradually from the classification layer (left) toward input ones (right).

Experimental Evaluations

We extensively verified our method on large scale CNNs, including the popular VGG-16 and ResNet-50 architectures. We utilize the Large Scale Visual Recognition Challenge 2012 (ILSVRC 2012) dataset (Russakovsky et al. 2015) and Pascal VOC 2012 dataset (Everingham et al. 2015), which are widely employed and easily accessible. We also use the Imagenet segmentation mask provided by (Guillaumin, Küttel, and Ferrari 2014a). We implement RAP with both Pytorch and Keras and visualize the explanation as a heatmap. For the evaluation, we utilized the Keras version to fairly compare with other explaining methods. We utilize the implementation introduced in (Alber et al. 2019) for the other explaining methods. As is customary in the field (but maybe unintuitive), the visualized heatmap is represented by seismic colors, where red and blue colors denote positive and negative values, respectively.

The results of our method are compared with those of existing attribution methods, including integrated gradients, gradient* gradient, Guided BackProp, pattern attribution and LRP- $\alpha\beta$ with $\{\alpha_1\beta_0, \alpha_2\beta_1\}$. Since (Ancona et al. 2018) proved that LRP- ϵ and DeepLift (Rescale version) are equivalent to the input* gradient when the model utilized the ReLU activation function, the result of input* gradient is the same as both methods. Since Pattern Attribution is developed for the VGG network, we do not report the result of it for experiments with Resnet or on the Pascal VOC dataset.

Qualitative Evaluation of Heatmap

For qualitatively evaluating the positive attributions generated by RAP, we compare the results by examining how the areas in which positive attributions converge are similar to those of the other methods. As the existing methods propagate the positive relevance well, we can utilize them to assess whether our method is consistent in attributing positive relevance. Fig. 5 presents the heatmaps generated from the various methods for the predicted images by the VGG-16 network. Fig. 6 illustrates the comparison between LRP- $\alpha_1\beta_0$ and RAP in ResNet-50. More qualitative comparisons are illustrated in the supplementary material.

To qualitatively evaluate the negative attributions, we regard the attributions allocated in the parts that are not related to the prediction as to the negative relevance. While our results clearly distinguish between the object and the irrelevant parts as positive and negative attributions, the attributions from other methods overlap each other and appear as purple, as shown in Fig. 5. The results shown are typical: we qualitatively assessed all images in the validation set of ILSVRC 2012 and Pascal VOC 2012 dataset, and most of them appear to show similarly satisfactory results.

Quantitative Assessment of Attributions

It is not trivial to assess the quantitative performance of the methods designed for explaining DNN models, since each method poses a different assumption and is designed for slightly different objectives. In this work, we utilized three methods that are commonly used to evaluate objectness and relevance: (i) Outside-Inside ratio, (ii) Pixel accuracy and Intersection of Union, and (iii) Region perturbation.

(Lapuschkin et al. 2016) introduce a method for assessing how the attributions are focused on the object, by computing the relevance inside and outside of the bounding box. We extend this method to consider the effect of correctly and wrongly distributed negative relevance. However, since the bounding box is not perfect for the object corresponding to the prediction, we additionally utilize the segmentation masks and metrics to assess how the positive attributions are correctly distributed in the target object. (Samek et al. 2017) introduce the method for quantitatively assessing the explanations methods, which utilizes the region perturbation process that progressively distorts the pixels from the heatmap and formalized this method as Area over the perturbation curve (AOPC). To evaluate how the negative attributions are distributed to the irrelevant regions of the prediction, we perturbate *least relevant first* (LeRF) and investigate the degradation of the accuracy.

In our experiment, we extract 10,000 correctly classified images from the validation set of imagenet and employ the specified bounding boxes (for the perturbation test). For the objectness scores, we utilized the 4,276 images from the imagenet dataset with segmentation masks and 1,449 images from the Pascal VOC validation set.

Objectness of Positive Attributions To verify how the attributions are distributed on the prediction object, we assess the outside-inside relevance ratio of attributions (Lapuschkin et al. 2016), utilizing the bounding box. We extend the original metric to evaluate the positive and negative relevance simultaneously.

$$\mu = \frac{\frac{1}{|P_{out}|} \sum_{q \in P_{out}} R_q^{(0)+} + \frac{1}{|P_{in}|} \sum_{p \in P_{in}} R_p^{(0)-}}{\frac{1}{|P_{in}|} \sum_{p \in P_{in}} R_p^{(0)+} + \frac{1}{|P_{out}|} \sum_{q \in P_{out}} R_q^{(0)-}} \quad (11)$$

Here, $|\cdot|$ is the cardinality operator and $P_{in,out}$ denotes the set of pixels inside and outside of the bounding box, respectively. When the positive (negative) relevance is attributed out of the bounding box, the value of μ is increased (decreased). By contrast, if the positive (negative) relevance is

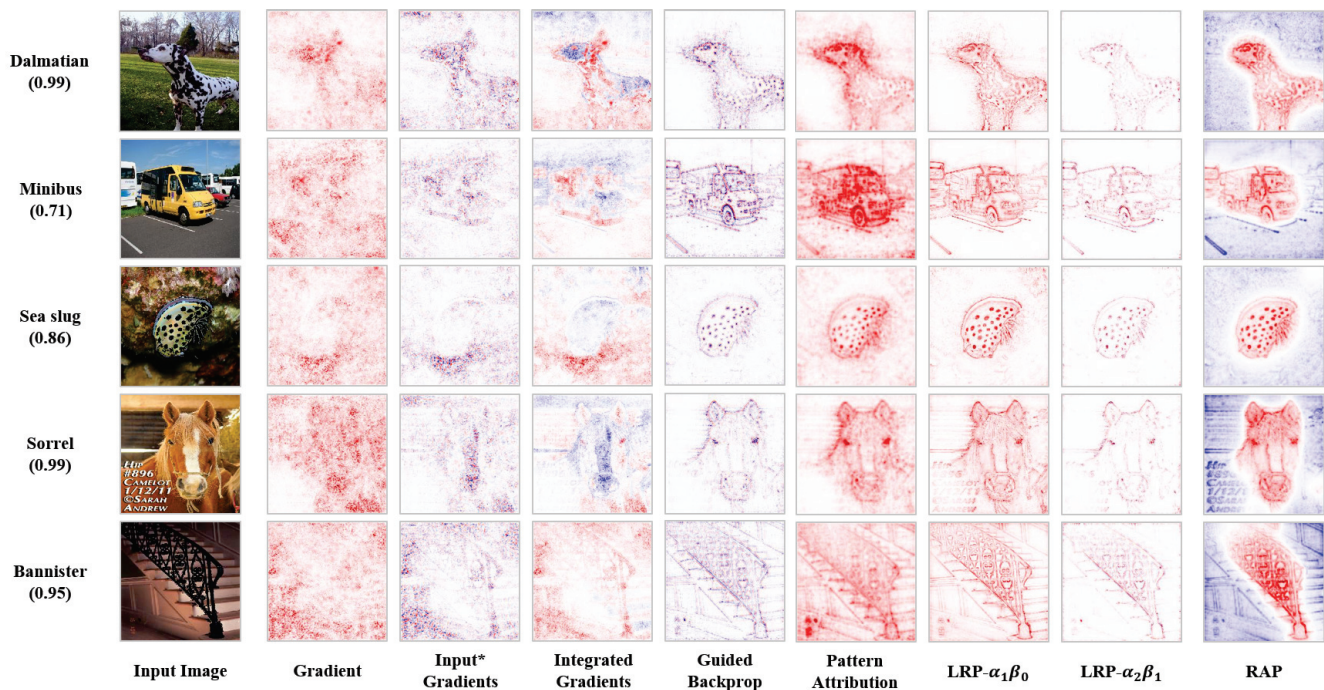


Figure 5: Comparison of the results of conventional methods and our RAP in VGG-16 network.

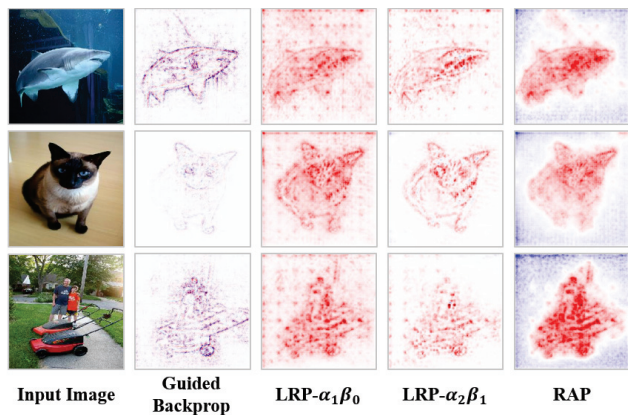


Figure 6: Comparison of visualization results applied on ResNet-50.

distributed in the bounding box, the ratio is decreased (increased). For the conventional methods which only consider the positive relevance, this metric becomes identical to the original metric of (Lapuschkin et al. 2016).

We present the results with(out) considering the negative relevance. In Tab. 1, this is denoted as ALL when considering both and POS when discarding the negative relevance. For both cases in Tab. 1, RAP provides the best scores with indicating that the relevance attribution is better distributed inside/outside the bounding box than the existing methods.

Furthermore, we utilize the segmentation masks provided for the ImageNet and Pascal VOC 2012 validation sets

to evaluate our explainability performance in terms of objectness. The evaluation is done by first producing an explainability map, followed by thresholding, considering two cases: (i) only positive relevance, in which case the mean value is taken as a threshold, and (ii) positive and negative relevance, where the threshold is set to zero. By thresholding, we produce a segmentation mask, considering values above the threshold as 1, and 0 otherwise. Tab. 1 shows that our method greatly outperforms previous work in both Pixel Accuracy and mean-IOU. These results are stable for other choices of the threshold as well, such as various percentiles and using the median. As can be seen, the output relevance is highly correlated with the input image objectness.

Interestingly, as shown in Tab. 2, our method is extremely competitive with respect to the objectness literature methods, which are trained with additional supervision, such as bounding boxes. This is demonstrated using the acceptable metric of mIOU on imagenet, which is commonly used for benchmarking the generic object segmentation (Cho, Yuille, and Lee 2017).

Interfering Negative Attributions When a DNN makes a correct prediction, removing the contribution of irrelevant pixels from the prediction should not change significantly the prediction accuracy and relevance values. A small decrement of accuracy is inevitable because the distortion of color and shape could affect the classification, resulting in unpredictable noise and attack. We carefully discuss these issues in more details in supplementary materials. Also, it is important to note that removing pixels corresponding to the negative attributions does not always bring an increment of the accuracy, because the negative relevance of incorrect predic-

Outside-Inside Ratio		RAP	$LRP_{\alpha_1\beta_0}$	$LRP_{\alpha_2\beta_1}$	Gradient	Input* Gradient	Integrated Gradients	Pattern Attribution	Guided Backprop
VGG-16	ALL	0.252	-	0.616	-	0.989	1.230	-	1.069
	POS	0.341	0.474	0.524	0.619	0.691	0.827	0.415	0.427
Res-50	ALL	0.164	-	0.302	-	0.996	1.195	-	1.035
	POS	0.166	0.429	0.299	0.597	0.689	0.698	-	0.296

Segmentation Mask		RAP	$LRP_{\alpha_1\beta_0}$	$LRP_{\alpha_2\beta_1}$	Gradient	Input* Gradient	Integrated Gradients	Pattern Attribution	Guided Backprop
Imagenet	PIX ACC	79.23	75.40	72.95	70.01	66.38	66.52	76.84	71.98
	mIOU	62.23	55.78	50.86	49.30	44.01	45.90	58.05	49.87
Pascal VOC	PIX ACC	73.91	70.86	69.43	68.14	50.01	52.38	-	66.92
	mIOU	55.60	49.82	46.85	46.07	31.69	34.39	-	43.63

Table 1: In Outside-inside relevance ratio result, the first (second) row is the ratio when considering all (only positive) relevance in Imagenet dataset. Segmentation mask result shows the pixel accuracy and mIOU of relevance heatmap in VGG-16 network.

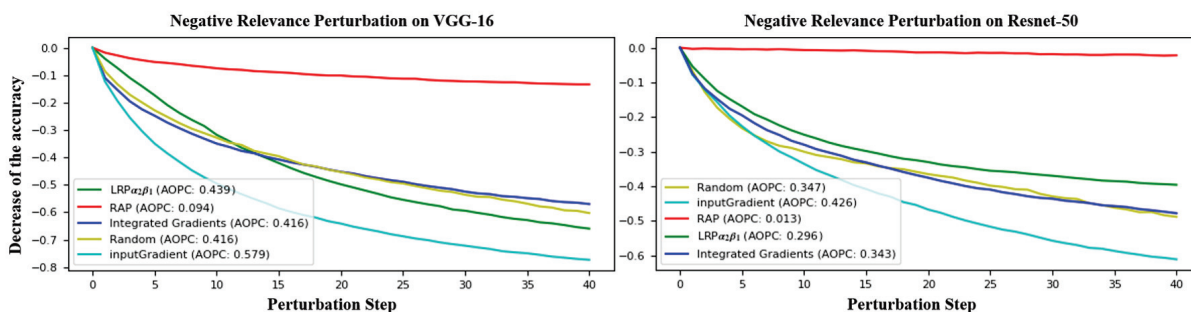


Figure 7: This graph illustrates the results of the negative perturbation on VGG-16 and Resnet-50. For each step, 100 pixels corresponding to the LeRF is perturbed as zero. RAP shows the unique characteristics of the robustness to the perturbation.

Method	mIOU
Guillaumin et al. (Guillaumin, Küttel, and Ferrari 2014b)	57.30
DeepMask (Pinheiro, Collobert, and Dollár 2015)	58.69
DeepSaliency (Li et al. 2016)	62.12
Xiong et al. (Xiong, Jain, and Grauman 2018)	64.22
Ours	62.23

Table 2: Quantitative mIOU results on the ImageNet Segmentation task. Our method is highly comparable to state of the art, despite not using the additional supervision.

tion does not denote the positive relevance of true label.

Fig. 7 shows the results when applying the LeRF perturbation on the VGG-16 and Resnet-50. For each step, we perturbate 100 pixels corresponding to the negative relevance, a total of 4,000 pixels distortion. As shown in the results, while other methods that consider the negative relevance show the rapid decrements of the accuracy, RAP rarely affects the prediction result during the negative attributions removal process. Thus, RAP distinguishes between the relevant and unimportant parts of the input image without overlapping each other.

Conclusion

In this paper, we propose RAP, a new method for interpreting the neurons in terms of importance to the predictions of DNNs, by assigning the relevance score according to the influence of each neuron. By approaching the relevance in terms of influence among the neurons, it is possible to separate relevant and irrelevant regions to the prediction. We evaluate our methods in quantitative and qualitative ways to verify that the attributions correctly account for the meaning. For the quantitative evaluation, we utilize the metrics: Outside-Inside ratio, mIOU and region perturbation to confirm how the attributions focused on the (ir)relevant object according to their assigned relevance scores. Overall, the experiments show that the RAP method leads to the desired characteristics: (i) a clear distinction of positive (relevant) and negative (irrelevant) attributions and (ii) it is indicative of objectness and can separate the main image object from the other irrelevant regions.

Acknowledgments

This work was supported by Institute for Information & communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No.2017-0-01779, A machine learning and statistical inference frame-

work for explainable artificial intelligence & No.2019-0-01371, Development of brain-inspired AI with human-like intelligence) and the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant ERC CoG 725974).

References

- Alber, M.; Lapuschkin, S.; Seegerer, P.; Hägele, M.; Schütt, K. T.; Montavon, G.; Samek, W.; Müller, K.-R.; Dähne, S.; and Kindermans, P.-J. 2019. Investigate neural networks! *Journal of Machine Learning Research* 20(93):1–8.
- Ancona, M.; Ceolini, E.; Oztireli, C.; and Gross, M. 2018. Towards better understanding of gradient-based attribution methods for deep neural networks. In *Proceedings of the International Conference on Learning Representations*.
- Bach, S.; Binder, A.; Montavon, G.; Klauschen, F.; Müller, K.-R.; and Samek, W. 2015. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS one* 10(7):e0130140.
- Baehrens, D.; Schroeter, T.; Harmeling, S.; Kawanabe, M.; Hansen, K.; and Müller, K.-R. 2010. How to explain individual classification decisions. *Journal of Machine Learning Research* 11(Jun):1803–1831.
- Cho, N.-G.; Yuille, A.; and Lee, S.-W. 2017. A novel linelet-based representation for line segment detection. *IEEE transactions on pattern analysis and machine intelligence* 40(5):1195–1208.
- Dabkowski, P., and Gal, Y. 2017. Real time image saliency for black box classifiers. In *Advances in Neural Information Processing Systems*, 6970–6979.
- Erhan, D.; Bengio, Y.; Courville, A.; and Vincent, P. 2009. Visualizing higher-layer features of a deep network. *University of Montreal* 1341(3):1.
- Everingham, M.; Eslami, S. M. A.; Van Gool, L.; Williams, C. K. I.; Winn, J.; and Zisserman, A. 2015. The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision* 111(1):98–136.
- Foerster, J. N.; Gilmer, J.; Sohl-Dickstein, J.; Chorowski, J.; and Sussillo, D. 2017. Input switched affine networks: An rnn architecture designed for interpretability. In *Proceedings of the International Conference on Machine Learning*, 1136–1145.
- Guillaumin, M.; Küttel, D.; and Ferrari, V. 2014a. Imagenet auto-annotation with segmentation propagation. *International Journal of Computer Vision* 110:328–348.
- Guillaumin, M.; Küttel, D.; and Ferrari, V. 2014b. Imagenet auto-annotation with segmentation propagation. *International Journal of Computer Vision* 110(3):328–348.
- Kindermans, P.-J.; Schütt, K. T.; Alber, M.; Müller, K.-R.; Erhan, D.; Kim, B.; and Dähne, S. 2017. Learning how to explain neural networks: Patternnet and patternattribution. *arXiv preprint arXiv:1705.05598*.
- Koh, P. W., and Liang, P. 2017. Understanding black-box predictions via influence functions. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 1885–1894. JMLR. org.
- Lapuschkin, S.; Binder, A.; Montavon, G.; Müller, K.-R.; and Samek, W. 2016. Analyzing classifiers: Fisher vectors and deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2912–2920.
- Lapuschkin, S.; Wäldchen, S.; Binder, A.; Montavon, G.; Samek, W.; and Müller, K.-R. 2019. Unmasking clever hans predictors and assessing what machines really learn. *Nature Communications* 10:1096.
- Li, X.; Zhao, L.; Wei, L.; Yang, M.-H.; Wu, F.; Zhuang, Y.; Ling, H.; and Wang, J. 2016. Deepsaliency: Multi-task deep neural network model for salient object detection. *IEEE Transactions on Image Processing* 25(8):3919–3930.
- Lundberg, S. M., and Lee, S.-I. 2017. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, 4765–4774.
- Mahendran, A., and Vedaldi, A. 2016. Visualizing deep convolutional neural networks using natural pre-images. *International Journal of Computer Vision* 120(3):233–255.
- Montavon, G.; Lapuschkin, S.; Binder, A.; Samek, W.; and Müller, K.-R. 2017. Explaining nonlinear classification decisions with deep Taylor decomposition. *Pattern Recognition* 65:211–222.
- Montavon, G.; Samek, W.; and Müller, K.-R. 2018. Methods for interpreting and understanding deep neural networks. *Digital Signal Processing* 73:1–15.
- Pinheiro, P. O.; Collobert, R.; and Dollár, P. 2015. Learning to segment object candidates. In *Advances in Neural Information Processing Systems*, 1990–1998.
- Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. Why should I trust you?: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1135–1144. ACM.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision* 115(3):211–252.
- Samek, W.; Binder, A.; Montavon, G.; Lapuschkin, S.; and Müller, K.-R. 2017. Evaluating the visualization of what a deep neural network has learned. *IEEE transactions on neural networks and learning systems* 28(11):2660–2673.
- Shrikumar, A.; Greenside, P.; and Kundaje, A. 2017. Learning important features through propagating activation differences. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 3145–3153. JMLR. org.
- Simonyan, K.; Vedaldi, A.; and Zisserman, A. 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*.
- Springenberg, J. T.; Dosovitskiy, A.; Brox, T.; and Riedmiller, M. 2014. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*.
- Sundararajan, M.; Taly, A.; and Yan, Q. 2017. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 3319–3328. JMLR. org.
- Xiong, B.; Jain, S. D.; and Grauman, K. 2018. Pixel objectness: learning to segment generic objects automatically in images and videos. *arXiv preprint arXiv:1808.04702*.
- Zeiler, M. D., and Fergus, R. 2014. Visualizing and understanding convolutional networks. In *European conference on computer vision*, 818–833. Springer.
- Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2921–2929.
- Zhou, B.; Bau, D.; Oliva, A.; and Torralba, A. 2018. Interpreting deep visual representations via network dissection. *IEEE transactions on pattern analysis and machine intelligence*.