

Interpreting Deep Neural Networks with Relative Sectional Propagation by Analyzing Comparative Gradients and Hostile Activations

Woo-Jeoung Nam,¹ Jaesik Choi,³ Seong-Whan Lee^{1,2*}

¹Department of Computer and Radio Communications Engineering, Korea University, Seoul, Republic of Korea

²Department of Artificial Intelligence, Korea University, Seoul, Republic of Korea

³Graduate School of Artificial Intelligence, KAIST, Daejeon, Republic of Korea
nwj0612@korea.ac.kr, jaesik.choi@kaist.ac.kr, sw.lee@korea.ac.kr

Abstract

The clear transparency of Deep Neural Networks (DNNs) is hampered by complex internal structures and nonlinear transformations along deep hierarchies. In this paper, we propose a new attribution method, Relative Sectional Propagation (RSP), for fully decomposing the output predictions with the characteristics of class-discriminative attributions and clear objectness. We carefully revisit some shortcomings of backpropagation-based attribution methods, which are trade-off relations in decomposing DNNs. We define hostile factor as an element that interferes with finding the attributions of the target and propagate it in a distinguishable way to overcome the non-suppressed nature of activated neurons. As a result, it is possible to assign the bi-polar relevance scores of the target (positive) and hostile (negative) attributions while maintaining each attribution aligned with the importance. We also present the purging techniques to prevent the decrement of the gap between the relevance scores of the target and hostile attributions during backward propagation by eliminating the conflicting units to channel attribution map. Therefore, our method makes it possible to decompose the predictions of DNNs with clearer class-discriminateness and detailed elucidations of activation neurons compared to the conventional attribution methods. In a verified experimental environment, we report the results of the assessments: (i) Pointing Game, (ii) mIoU, and (iii) Model Sensitivity with PASCAL VOC 2007, MS COCO 2014, and ImageNet datasets. The results demonstrate that our method outperforms existing backward decomposition methods, including distinctive and intuitive visualizations.

Introduction

As Deep Neural Networks (DNNs) have shown a remarkable performance in various fields, many studies have attempted to resolve the basis of network predictions. However, it still lacks clear transparency about the myriad of components and the complex inner structure of DNNs. The problem of attribution, also called relevance, seeks the most relevant factors with respect to the predictions of DNNs and characterizes them as a supporting basis for the decision. Grad-CAM (Selvaraju et al. 2017) is the most popular and widely used method in the field of weakly supervised segmentation and detection due to the easily applicable property



Figure 1: Relative Sectional Propagation (RSP) aims to fully decompose the network predictions with taking advantages of i) strong objectness, ii) class-discriminateness, and iii) detailed descriptions of neuron activations.

and its high performance for localizing the primary objects. Despite these advantages, there is a limitation that the feature extraction stage of DNNs cannot be decomposed (Rebuffi et al. 2020). Also, detailed information about the neuron activation is lost due to the interpolation of the class activation map. To fully interpret the network, many studies based on the modified backpropagation algorithm (Bach et al. 2015; Kindermans et al. 2017; Montavon et al. 2017; Zhang et al. 2018; Nam et al. 2019) attempted to identify the significant parts in the input image with each of their own perspectives by decomposing the output predictions in a backward manner. By visualizing attributions as saliency maps, the important objects are highlighted as the basis for the predictions. Despite many studies on such methods, it is challenging to address the class-discriminateness, incorrect distribution of attributions that are derived in unrelated objects, and model sensitivity. Furthermore, the different mechanisms among various types of DNNs and the variations of feature attention according to the layer depth make the interpretation of the network more difficult.

The class discriminative issue was addressed in recent papers, and the contrastive perspective was presented as a countermeasure (Zhang et al. 2018; Gu, Yang, and Tresp 2018). The idea of contrasting is erasing the duplicated attributions among classes by backpropagating twice from the target and all other classes. By efficiently removing the duplicate relevance of the activated neurons, it is possible to obtain the attributions of the target class within the object area. However, unreasonable positive or negative attributions in irrelevant regions, such as background and watermark, are easily found in the results of this concept.

In this paper, we propose an attribution method, Rela-

*Corresponding author: Seong-Whan Lee
Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

tive Sectional Propagation (RSP), by analyzing the relative gradient activation maps between the target and hostile classes and propagating corresponding relevance according to the sectional influence of individual neurons. We carefully investigate the reasons for the non-suppressed characteristics of neuron activations according to the different classes and address these issues by assigning the bi-polar relevance scores: from highly related with target to highly relevant to hostile classes. Inspired by the traits of **winner always wins** (Zhang et al. 2018) among activated neurons, the main idea is to separately compute the relative gradient activation maps, which contain the neuron importance, and purge the conflicting units to the channel attributions along the channel-axis, thereby preventing the decrement of the gap between bi-polar relevance scores. Our method preserves the conservation rule (Bach et al. 2015) to prevent degeneracy problem, and allocates the relevance scores aligned with contributions.

Fig. 1 illustrates the samples that summarize the advantages of RSP. The attributions are fully decomposed from the output to input with the characteristics of i) the detailed visualizations of neuron activations, ii) strong objectness to output predictions, and iii) discriminativeness among classes with bi-polar relevance scores (positive and negative). As these characteristics are the trade-off relations in previous attribution literature, we mainly focus on overcoming the limitations during decomposition. The main contributions of this work are as follows:

- We propose a new method for decomposing the output predictions with the relative gradient activation maps and backward sectional propagation according to the individual influences of neurons. By hostilely changing the priority of attributions corresponding to the non-target, it is possible to properly distribute the bi-polar relevance scores between the predicted classes while maintaining the irrelevant attributions as negative.
- We carefully address the phenomenon of non-suppressed characteristics of activated neurons and the contrary influences of the conflicting units to the channel attribution map, both of which prevent the attributions from being distinctive. We present a purging process to account for these neurons and to sustain the gap between positive and negative attributions.
- For evaluation, we apply Pointing Game (Lapuschkin et al. 2016), sanity check with Model Sensitivity (Adebayo et al. 2018), and mIOU to assess the quality of attributions. We report the performance in two cases of model decision (either only correct or all labels) to confirm the efficacy of interpreting the models. The evaluation demonstrates that our method outperforms other backpropagation-based attribution methods in circumstances of complete decomposition, including the advantages of strong objectness, class-discriminativeness, and detailed descriptions of activations.

Related Work

As DNNs are applied to a variety of traditional computer vision methodologies (Roh et al. 2007; Roh, Shin, and Lee

2010; Yang and Lee 2007; Bulthoff et al. 2003), there are many attempts to improve the transparency issues of DNNs. As the manner of interpreting a DNN model itself, intermediate features are visualized by maximizing the activated neurons in intermediate layers (Erhan et al. 2009) or generating saliency maps (Simonyan, Vedaldi, and Zisserman 2013; Zeiler and Fergus 2014; Mahendran and Vedaldi 2016; Zhou et al. 2016; Dabkowski and Gal 2017; Zhou et al. 2018). (Ribeiro, Singh, and Guestrin 2016) proposed LIME, which explains the black-box models, by locally approximating them as simpler interpretable models.

A perturbation-based approach directly analyzes the variations of decision when distorting the input of the network. (Zeiler and Fergus 2014; Petsiuk, Das, and Saenko 2018) investigate the variations of the output as applying occlusions to images with specified patterns. (Fong, Patrick, and Vedaldi 2019) introduced the concept of extremal perturbation to understand network behavior with theoretically based masking.

From the viewpoint of decomposing the network decision, (Bach et al. 2015) proposed several kinds of Layer-wise Relevance Propagation (LRP) rules with the concept of relevance and conservation. As a theoretical foundation, (Montavon et al. 2017) proposed Deep Taylor decomposition by utilizing Taylor expansion among neurons of intermediate layers. (Selvaraju et al. 2017) proposed Grad-CAM to generate class-discriminative activation maps by computing gradients with respect to the last convolutional units of the feature extraction stage. Guided BackProp (Springenberg et al. 2014) is based on gradient backpropagation that considers only positive values. Integrated Gradients (Sundararajan, Taly, and Yan 2017) addressed the gradient saturation problem by computing the average partial derivatives of the output. DeepLIFT (Shrikumar, Greenside, and Kundaje 2017) decomposes the differences in relevance scores between the activation and its reference. (Ancona et al. 2018) approached with the theoretical perspective to attributions and formally proved the conditions of equivalence of the previous methods. (Lundberg and Lee 2017) unifies some explaining methods and approximate with the shapley value. (Zhang et al. 2018) proposed Excitation Backprop (EB) by modeling the probabilistic winner-take-all process and addressed the class-discriminative issue with contrastive top-down attention. (Nam et al. 2019) pointed out the overlapping phenomenon of positive and negative relevance and utilized the influential perspective to separate relevant and irrelevant attributions. (Lapuschkin et al. 2019) discussed the spurious correlations among the objects in the input, (such as tags in pictures) and present the necessity of comprehending the network decision to unmask “Clever Hans” phenomenon.

We mainly focus on attribution methods based on backpropagation. Although there are many studies related to this, the interpretation of DNNs remains a trade-off between the objectives of each attribution method. Therefore, our method aims to overcome the main issues: class-discriminativeness, the details of neuron activations with full decomposition, and objectness which can separate the main objects and background.

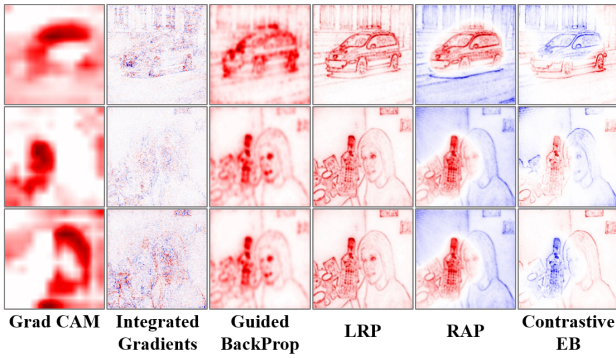


Figure 2: The shortcomings of some attribution methods. Results of Grad-CAM and Integrated Gradients are lack of detail. Gui, LRP, and RAP are not class discriminative. Some attributions of CEB are distributed in unrelated parts.

Revisiting Attribution Methods

In a multi-classification task, it is clear that objects in the input are not learned antagonistically during the training procedure, because most of the networks consider the correct predictions among output logits simultaneously, not competitively. When we investigate the saliency maps of intermediate layers, highly activated neurons always have the lion’s share of relevance.

Fig. 2 illustrates the motivational examples of trade-off among attribution methods (the same images in Fig. 1) and presents the result of Grad-CAM, LRP, CEP, and RAP, which are based on a modified backpropagation algorithm with their individual purposes. As is widely known, Grad-CAM shows impressive localization performance for finding attributions of DNN predictions. However, because it utilizes the end layer of the feature extraction stage and interpolation, there are many losses in describing the details of pixel-level granularity.

The other methods in Fig. 2 could fully decompose DNNs in a backward manner, including the feature extraction stage. The attributions from Integrated gradient provide the responsibility for a target label by computing the gradient w.r.t the features of the image. However, it is difficult to intuitively judge the quality of interpretation in a human-view due to the scattered and overlapped positive and negative attributions. Guided Backprop and LRP represent the output logit as the relevance scores in a pixel-level. However, although there is a minor difference in values, there is no visual difference in attributions between classes. The role of the CNN feature extractor stage goes from low-level features (edge or color) to high-level features (object or texture) as the layer becomes deeper. It is inevitable that the low-level features activated in the front stage, which turns out to be irrelevant in the latter part, are assigned the positive relevance in the backward layer-wise propagation. In conjunction with the above problem, unrelated parts to the target class, e.g. corner and watermark, tend to be attributed as positive.

While RAP approaches with the influential perspective to separate relevant/irrelevant attributions and shows strong ad-

vantages of objectness, it is not distinguishable among the predicted classes. As a countermeasure, (Zhang et al. 2018) proposed a contrastive perspective, which contrasts the relevance for one class with the ones of all the others. However, the positive or negative relevance is distributed in the background or other parts not related to its origin. As this method shows a relatively activated area compared to other classes, there is a probability of having positive or negative relevance scores in irrelevant parts.

In this paper, the concept of hostile refers to an element that could have a negative influence on finding the attribution corresponding to the target. For example, when we decompose the output of the right image in Fig. 1 from the bottle class, the hostile class would be a person. The relevance of the hostile class is represented as hostile attributions. By assigning the negative relevance scores to the hostile attributions, we thwart the characteristics of “winner always wins”, resulting in assigning bi-polar relevance scores among the target (positive), hostile (negative) attributions. We also present our method with a contrastive perspective for the hostile class. In this case, all other classes except the target are set hostile.

Relative Sectional Propagation

Inspired by the above problems, our method has two main streams: (i) Relative Gradient Activation Map and purging process, and (ii) Sectional propagation according to influence with gradient and uniform shifting.

Relative Gradient Activation Map Letter y denotes the value of the network output before passing through the final layer of the classification stage. $\{t, o_1, \dots, o_n, b\}$ is class notation where each notation represents a target, other predictions, and irrelevant classes, respectively. First, we obtain the gradient activation map G by backpropagating the gradient of y until the end convolution layer X of the feature extraction stage.

$$G_{ijk}^{(t)} = \lambda * ReLu \left(x_{ijk} * \frac{1}{Z} \sum_i \sum_j \frac{\partial y^t}{\partial x_{ijk}} \right) \quad (1)$$

$$F_{ijk}^{(t)} = n * G_{ijk}^{(t)} - \sum_{q=1}^n G_{ijk}^{(h_q)}$$

Here, x_{ijk} denotes the neuron in the k^{th} feature map of layer X , indexed by width i and height j dimensions, and λ is a normalization factor to keep a maximum value as 1. The equation for computing $G_{ijk}^{(t)}$ is the same process in Grad-CAM except for the last linear combination between the feature map X and partial linearization. h represents the hostile classes $\{o_1, \dots, o_n\}$.

The comparative gradients of the target against hostile class are contained in $F_{ijk}^{(t)}$. The attributions that are conflict to the channel attribution map (Fong, Patrick, and Vedaldi 2019), generated by summing over the channel dimension k , still exist along the channel. Conflicting attributions refer to the units in $F_{ijk}^{(t)}$ with opposite sign to the channel attribu-

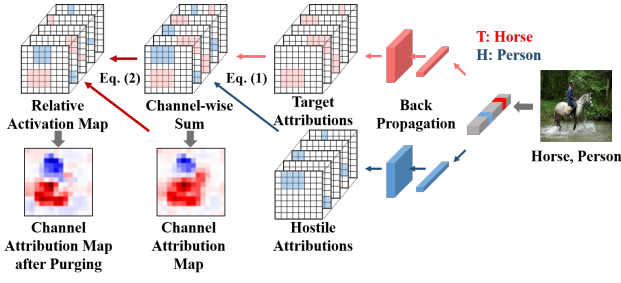


Figure 3: An illustration of generating the relative gradient activation map. The elements marked with red and blue color represent the target: *Horse* and hostile: *Person* attributions, respectively. Eq. (1) computes the comparative gradients of the target. Eq. (2) is the purging process to eliminate the conflicting attributions to the channel attribution map.

tion. To preserve the gap between target and hostile attributions during propagation, it is necessary to eliminate these conflicting attributions as follows.

$$F_{ijk}^{(t)} = \begin{cases} F_{ijk}^{(t)} & , \text{sign} \left(F_{ijk}^{(t)} \right) = \text{sign} \left(\sum_k F_{ijk}^{(t)} \right) \\ 0 & , \text{sign} \left(F_{ijk}^{(t)} \right) \neq \text{sign} \left(\sum_k F_{ijk}^{(t)} \right) \end{cases} \quad (2)$$

Fig. 3 and Fig. 4 show the overview of generating the relative gradient activation map and the effect of the purging process to attributions in the intermediate layer of ResNet-50, respectively. In Fig. 4, the first row illustrates the channel activation maps of intermediate layers, and most activations are concentrated in the dog region. After applying Eq. (1) in case $\{t=\text{person}, h=\text{dog}\}$, the channel attribution map is shown as the first column. Without the purging process, although the dog regions are negative in the channel attribution map, there are still conflicting units (positive) along the channel. When we backpropagate each attribution in this state, it is inevitable to have an adverse effect on the backward step due to the nature of “winner always wins”. Thus, they are canceled out and the gap between bi-polar relevance is decreased during the propagation procedure. As shown in the second row, instead of exact positions of hostile attributions, irrelevant parts, such as corners, are emphasized as negative. The phenomenon that the attribution is visualized as tiles is due to skip-connection operations in ResNet.

After, we have non-overlapping positive and negative attributions along the channel in $F_{ijk}^{(t)}$ through channel-wise purging. Here, we set the positive and negative sections as $\mathcal{P}_{ijk}^{(t)} = \{i, j, k | F_{ijk}^{(t)} > 0\}$ and $\mathcal{N}_{ijk}^{(t)} = \{i, j, k | F_{ijk}^{(t)} < 0\}$. We normalize these values to make the sum of positive values have twice as many values as the sum of negative values: $\mathcal{P}_{ijk}^{(t)} \leftarrow 2 * \mathcal{P}_{ijk}^{(t)}, R_{ijk}^{(t)} = \mathcal{P}_{ijk}^{(t)} \cup \mathcal{N}_{ijk}^{(t)}$. This normalization is necessary to preserve the conservation rule and prevent degeneracy problems.

We report the difference between the two cases of obtaining the relative activation map $R_{ijk}^{(t)}$, (i) from Eq. (1), and (ii) applying the contrastive perspective instead of Eq. (1). For the latter case, we obtain $F_{ijk}^{(t)}$ by applying the contrastive

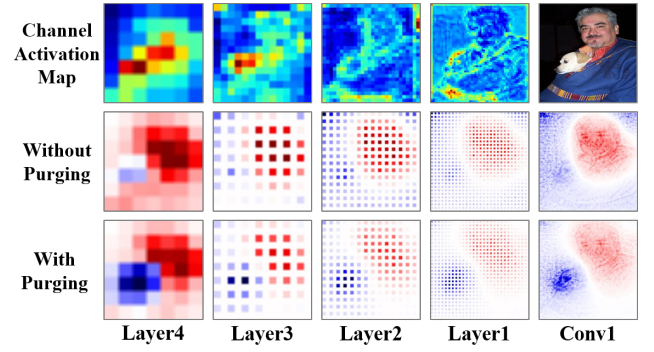


Figure 4: The difference between the channel attributions of intermediate layers with/without the purging process.

excitation backprop (Zhang et al. 2018) until the same convolution layer X without channel-wise sum and ReLu. This is conceptually the same in terms of computing the relative gradient map for target class with the ones of all other (hostile) classes. The purging process and normalization are equally applied. Since the configuration and working mechanism are different according to the type of network, we report the results of both perspectives. The detailed equation of the latter case is described in the supplementary material.

Sectional Relevance Propagation with Gradient The forward process between the current layer $l + 1$ and the layer l to which attributions are propagated is denoted as $f(x, w^{(l,l+1)})$. Also, we notate the boolean masks of $\mathcal{P}_{ijk \in (l+1)}^{(t)}$ and $\mathcal{N}_{ijk \in (l+1)}^{(t)}$ as $\mathcal{B}_{ijk \in (l+1)}^{+(t)}$, $\mathcal{B}_{ijk \in (l+1)}^{-(t)}$, respectively. Since $w^{(l,l+1)}$ are not directly influenced to $R_{ijk \in (l+1)}^{(t)}$, it is necessary to compute the gradient between $R_{ijk \in (l+1)}^{(t)}$ and the sectional influence of individual neurons $f(x, w^{(l,l+1)}) * \mathcal{B}_{ijk \in (l+1)}^{\pm(t)}$ with respect to weight $w^{(l,l+1)}$. This gradient contains the correlation between the individual contributions of each neuron in a forward pass and the attributions in $R_{ijk \in (l+1)}^{(t)}$.

$$\begin{aligned} \nu^+ &= \frac{f(x, w^{(l,l+1)}) * \mathcal{B}_{ijk \in (l+1)}^{+(t)}}{\partial w^{(l,l+1)}} \mathcal{P}_{ijk \in (l+1)}^{(t)} \\ \nu^- &= \frac{f(x, w^{(l,l+1)}) * \mathcal{B}_{ijk \in (l+1)}^{-(t)}}{\partial w^{(l,l+1)}} \mathcal{N}_{ijk \in (l+1)}^{(t)} \end{aligned} \quad (3)$$

Here, Eq. (3) could be represented as easily readable format: $\nu = f^*(f(x, w) * B, w, R)$ (vector-Jacobian product), which is implemented in many deep learning libraries and highly optimized. Through this gradient ν , we backpropagate the attributions in $R_{ijk \in (l+1)}^{(t)}$ to the previous layer l with an influential perspective (Nam et al. 2019) of individual neurons.

$$\begin{aligned} \hat{R}_{ijk \in (l)}^{\{t\}} &= x \odot f^*(f(x, \nu^+), \nu^+, \mathcal{P}_{ijk \in (l+1)}^{(t)} \odot f(x, \nu^+)) \\ &+ x \odot f^*(f(x, \nu^-), \nu^-, \mathcal{N}_{ijk \in (l+1)}^{(t)} \odot f(x, \nu^-)) \end{aligned} \quad (4)$$

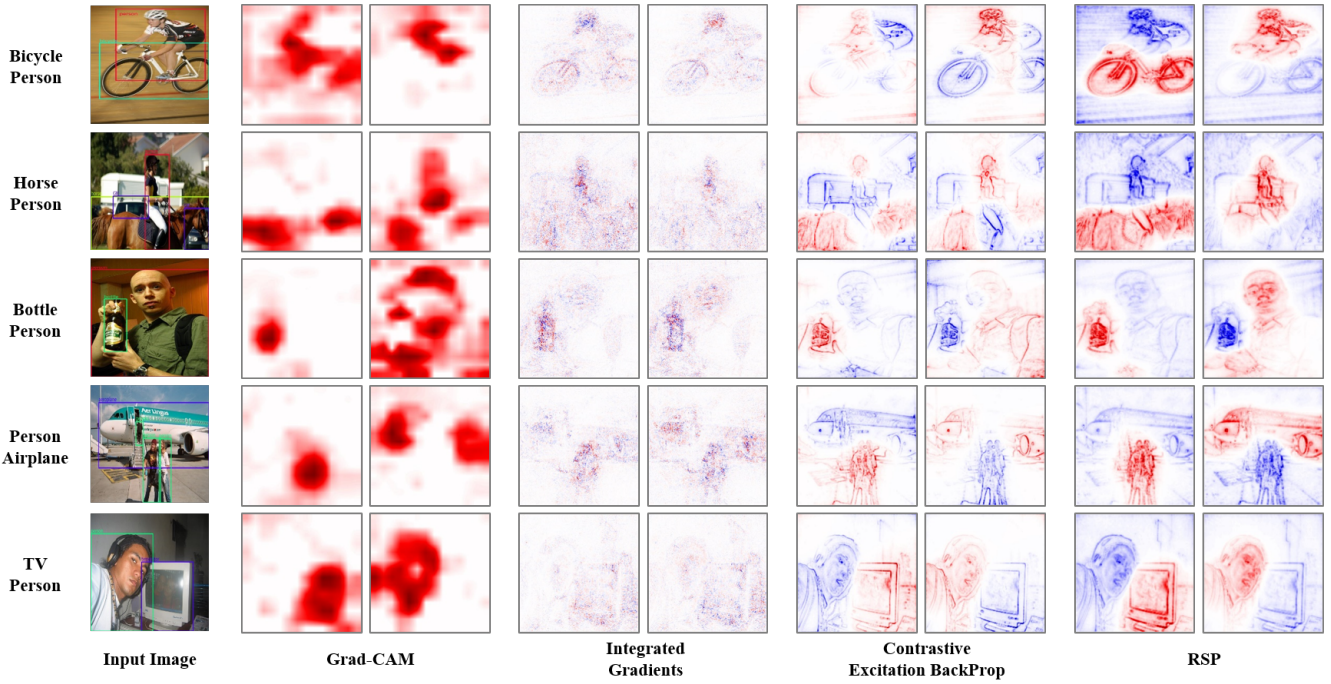


Figure 5: Comparison of the conventional attribution methods and RSP applied to VGG-16. The class names on the left side represent the predictions of DNNs among labels in input image. Each method shows the attribution results per predicted class.

Here, \odot and \oslash denote the element-wise multiplication and division, respectively. The influence perspective represents that the importance of each neuron is an order of the absolute value, not the sign of the value. To be more specific, the attributions of the highest (positive) and lowest (negative) relevance scores denote a large amount of influence on the target and hostile classes, respectively. Attributions with a relevance score near-zero mean having a relatively small influence. From this influential perspective, it is possible to assign the relevance scores to activated neurons in order of importance among intermediate layers.

Now, the whole relevance sum of $\hat{R}_{ijk \in (l)}^{(t)}$ is the same as the sum of $R_{ijk \in (l+1)}^{(t)}$. We utilize uniform shifting (Nam et al. 2019) to change the irrelevant attributions, in which relevance scores are near zero, into negative. Let Γ be the number of activated neurons in l and the sum of $\hat{R}_{ijk \in (l)}^{(t)}$ is S , this sum value is evenly divided and subtracted. To preserve the entire relevance sum as S for each layer, we double the values of attributions in $\hat{R}_{ijk \in (l)}^{(t)}$.

$$\ddot{R}_{ijk \in (l)}^{(t)} = \begin{cases} 2 * \hat{R}_{ijk \in (l)}^{(t)} - S * \frac{1}{\Gamma} & , x_{ijk} > 0 \\ 2 * \hat{R}_{ijk \in (l)}^{(t)} & , x_{ijk} = 0 \end{cases} \quad (5)$$

In case that a neuron x_{ijk} is not activated, $\hat{R}_{ijk \in (l)}^{(t)}$ is equal to zero. Relatively unimportant attributions, which are near zero, would be converted into the negative during the propagation procedure, thereby the irrelevant attributions, e.g. background, have the negative relevance scores in the final output. $\ddot{R}_{ijk \in (l)}^{(t)}$ is the input attributions for the next previous




Figure 6: Comparison of CEB, RSP and RSP with contrastive perspective in ResNet-50.

layer $l - 1$ and propagated by repeating this process from the purging step. This procedure is repeated until the first layer $l = 1$ of the model. For the final propagation between the input and first layer, we adopt the Z^β rule (Bach et al. 2015) which is commonly used for propagating to the input layer, resulting in clear visualizations without distorting the priority of attributions. Detail expansion of each equation is described in the supplementary material.

Experimental Evaluations

Implementation Details

We utilize the popular CNN architectures: VGG-16 and ResNet-50. Each model is trained on Pascal VOC 2007 (Everingham et al. 2010) and MS COCO 2014 dataset (Lin et al. 2014), which are widely employed and easily accessible.

PASCAL VOC 2007										COCO 2014							
METHOD	T	VGG-16				ResNet-50				VGG-16				ResNet-50			
		ALL		DIF		ALL		DIF		ALL		DIF		ALL		DIF	
		PG	mIOU	PG	mIOU	PG	mIOU	PG	mIOU	PG	mIOU	PG	mIOU	PG	mIOU	PG	mIOU
Grad-CAM	L	.866	.43/.49	.740	.39/.48	.903	.56/.57	.823	.47/.57	.542	.35/.46	.490	.33/.43	.573	.44/.51	.523	.40/.48
	P	.945	.41/.50	.924	.33/.54	.953	.55/.58	.932	.44/.59	.727	.30/.49	.689	.25/.45	.705	.39/.52	.674	.32/.47
Gradient	L	.762	.00/.47	.568	.00/.41	.723	.00/.45	.568	.00/.40	.355	.00/.39	.289	.00/.37	.319	.00/.39	.262	.00/.37
	P	.858	.00/.49	.716	.00/.50	.734	.00/.44	.605	.00/.43	.547	.00/.44	.492	.00/.40	.455	.00/.42	.405	.00/.38
DeconvNet	L	.675	.00/.41	.441	.00/.31	.686	.00/.43	.447	.00/.33	.241	.00/.35	.164	.00/.32	.273	.00/.35	.192	.00/.33
	P	.802	.00/.46	.573	.00/.37	.789	.00/.44	.595	.00/.39	.469	.00/.36	.372	.00/.31	.429	.00/.36	.338	.00/.31
Guided BackProp	L	.758	.00/.49	.530	.00/.43	.771	.00/.51	.594	.00/.46	.365	.00/.41	.288	.00/.39	.410	.00/.43	.340	.00/.41
	P	.880	.00/.52	.784	.00/.54	.857	.00/.53	.756	.00/.53	.600	.00/.47	.536	.00/.43	.573	.00/.47	.519	.00/.44
Excitation BackProp	L	.735	.00/.46	.520	.00/.45	.785	.00/.46	.623	.00/.45	.377	.00/.42	.304	.00/.40	.437	.00/.43	.374	.00/.41
	P	.856	.00/.47	.742	.00/.53	.864	.00/.47	.768	.00/.50	.573	.00/.47	.505	.00/.45	.582	.00/.46	.533	.00/.44
c*Excitation BackProp	L	.766	.38/.45	.634	.34/.50	.857	.49/.49	.741	.45/.56	.472	.32/.46	.417	.30/.45	.536	.41/.49	.485	.37/.48
	P	.856	.40/.42	.784	.39/.55	.945	.52/.49	.887	.51/.62	.659	.37/.49	.620	.34/.50	.671	.47/.53	.636	.42/.53
RSP	L	.849	.51/.51	.712	.43/.54	.859	.49/.51	.749	.39/.49	.540	.43/.49	.479	.37/.47	.558	.39/.46	.504	.35/.43
	P	.946	.56/.51	.903	.54/.63	.909	.54/.53	.836	.44/.54	.725	.51/.56	.680	.45/.54	.688	.44/.51	.654	.38/.48
c*RSP	L	.785	.46/.47	.627	.42/.52	.891	.52/.52	.777	.46/.54	.475	.39/.47	.418	.36/.45	.545	.41/.47	.488	.37/.44
	P	.881	.49/.46	.791	.51/.60	.949	.56/.53	.893	.53/.61	.675	.46/.51	.634	.42/.52	.697	.47/.52	.659	.42/.49

Table 1: The performance of Pointing Game and mIOU over Pascal VOC 2007 test set and COCO 2014 validation set. T denotes the different circumstance of each step for testing: P : Only predicted classes, L : All labels. ALL and DIF represent the full data and the subset of difficult images, respectively. For the mIOU results, left/right values means the performance when applying without/with the threshold. The threshold is set as the mean value of positive attributions. All attribution methods, except Grad-CAM, is fully decomposed from the output to the first layer of each network.

For a fair comparison, the models that we used are available online with TorchRay package (Fong, Patrick, and Vedaldi 2019). We implement our method with PyTorch and visualize the attributions as a heatmap represented by seismic colors, where red and blue colors denote positive and negative values, respectively. We utilize the implementation introduced in (Fong, Patrick, and Vedaldi 2019) for the other attribution methods. All conditions in experiments are the same except for setting the saliency layer. For the fully decomposed environment, the saliency layer for visualizing the attributions is set as the first convolution layer of each model. Each type of evaluation is described in subsequent sections.

Qualitative Assessment

To qualitatively evaluate the attributions of each method, we report the visual differences to compare how the high-rated points are gathered in the bounding box. As the goal of attribution methods are the same for seeking the most important factors, we could assess the consistency of positive relevance among methods. Among many researches related to attribution methods that are described in related work, we compare the methods that can be visualized with class-discriminateness and contain detailed information of neuron activations. The other methods are not qualitatively compared due to the reasons related to Fig. 2 (shown in supplementary material). The comparison methods are Grad-CAM, Integrated Gradients, Contrastive Excitation Backprop, and RSP. Fig. 5 and Fig. 6 illustrate the heatmaps of each method for the output predictions of VGG-16 and ResNet-50, respectively. Compared to other methods, RSP

shows detailed descriptions of activated neurons and clear separations between the target object and other objects (including background), resulting in much clearer visualizations with strong objectness. More qualitative comparisons are given in the supplementary material.




Figure 7: Model weights of VGG-16 are progressively initialized from the end to beginning.

Sanity Check (Adebayo et al. 2018) addresses the non-sensitivity problem of some saliency methods when the parameters of the model are randomly initialized in a cascading fashion from the end layer. It is crucial to support the basis that our explanation and the model decision are dependent on each other. Fig. 7 illustrates the variations of our method when applying weight randomization progressively. Attributions from each label are extremely distorted compared to the original explanations.

Evaluating Quality of Attributions

Pointing Game Pointing Game (Zhang et al. 2018) assesses the attribution methods by computing the matching scores of localization between the highest relevance point and the semantic annotations of object categories in the im-

Method	mIOU
Grad-CAM (threshold: mean) + CRF	52.14
Segmentation Prop (Guillaumin, Küttel, and Ferrari 2014a)	57.30
DeepMask (Pinheiro, Collobert, and Dollár 2015)	58.69
RAP (Nam et al. 2019)	59.46
DeepSaliency (Li et al. 2016)	62.12
Pixel Objectness (Xiong, Jain, and Grauman 2018)	64.22
RSP	60.81
RSP + CRF	64.51

Table 2: Segmentation mIoU results on the ImageNet Segmentation task. Our method is highly comparable to those methods without using the additional supervision.

age. However, in previous literature, this metric does not consider the performance of the model itself because the assessed attributions are decomposed from the label, not the predicted classes. It is necessary to consider the context of the predictions because performing the decomposition on a label that has not been identified by DNNs is likely to lead to incorrect interpretation. Therefore, we report both cases of the decomposition from (i) P: only predicted labels and (ii) L: all labels in Table 1.

To compare the attribution methods in completely decomposed circumstances, all methods except Grad-CAM do backward propagation until the first convolution layer of each network. Methods with notation c^* denote a contrastive perspective to set comparative classes to target. In our case, c^* represents the case of computing the relative gradient activation map from the contrastive perspective. As shown in the results, our method shows the superior performance of localization in both cases compared to the fully decomposed attribution methods. In VGG-16, the relative gradient maps between the target and the rest of the predicted classes show much superior performance compared to the contrastive initialization. On the contrary, ResNet-50 shows much better results in the class setting of contrasting.

Objectness and Weakly Supervised Segmentation Interpretation of the attribution methods and objectness is closely related in terms of aiming to find the pixels corresponding to the target object. Based on this concept, many studies in weakly supervised segmentation (image-label level supervision) (Ahn, Cho, and Kwak 2019; Lee et al. 2019; Huang et al. 2018) start from the seeds extracted from the attribution methods. We report the mean Intersection of Union (mIoU) to measure the false positives of attributions that are distributed in irrelevant parts (other objects or background). Instead of comparing it with the segmentation mask, we computed the accuracy of mIOU between the bounding box and positive attributions to allow the error range in both datasets. Since some methods: Guided Backprop, DeconvNet, and Excitation Backprop, have all positive values in output results, the mean value of attributions is taken as the threshold. In Table 1, the left (right) values for each method represent the without (with) the threshold. Our method outperforms in almost all cases, especially when the target is in the set of correct predictions.

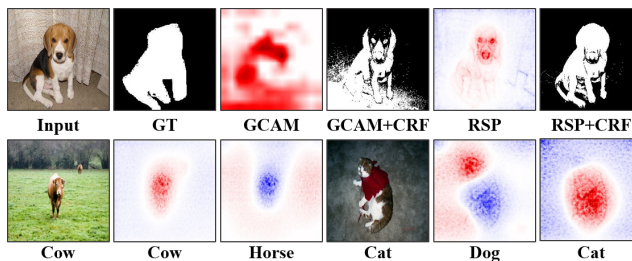


Figure 8: The first and second rows demonstrate the weakly supervised segmentation results with CRF and Misconception of ResNet-50 about a single object, respectively.

Furthermore, we represent the segmentation performance on the ImageNet segmentation dataset (Guillaumin, Küttel, and Ferrari 2014b), which consists of 4,276 images with segmentation mask. Table 2 compares other saliency based objectness methods and RSP. The results of Grad-CAM, RAP, and RSP are from the ImageNet pretrained model in PyTorch. Although some methods use additional supervision, (i.e. bounding box, optical flow), RSP shows comparable performance with only image-label level supervision. First row in Fig. 8 shows the results of Grad-CAM (mean threshold) and RSP with Dense-CRF (Krähenbühl and Koltun 2011) to refine the attributions, resulting in superior performance compared to other methods.

Discussions

To the best of our knowledge, there is still no exact elucidation of the internal mechanisms except for the structural and conceptual differences between VGG-Net and ResNet. Some differences could be inferred by investigating the failure attributions seen only in ResNet. ResNet tends to classify objects independently between classes. For example in Fig. 8 second row, a single object is misclassified as two classes by focusing on different features, resulting in overlap of the relative gradient activation map from Eq. (1). Furthermore, there is a clear effect of unlabeled objects in ResNet. In these cases, the contrastive hostile setting shows better performance to find the target attributions. More discussion and analysis are described in the supplementary materials.

Conclusion

In this paper, we propose a new attribution method for decomposing the output of DNNs by assigning the bi-polar relevance score among the target and hostile classes. From the antagonistic perspective among objects, it is possible to allocate the bi-polar relevance scores to neuron activations, resulting in distinguishable and attentive decomposition. We assess our methods in quantitative and qualitative ways with i) Pointing Game, ii) mIOU, and iii) Model randomization to confirm the quality of attributions. The results demonstrate that the attributions from RSP have the properties of strong objectness, class-discriminateness, and detailed descriptions of the neuron activations.

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 framework for explainable artificial intelligence & No. 2019-0-01371, Development of brain-inspired AI with human-like intelligence & No. 2019-0-00079, Artificial Intelligence Graduate School Program, Korea University).

References

- Adebayo, J.; Gilmer, J.; Muelly, M.; Goodfellow, I.; Hardt, M.; and Kim, B. 2018. Sanity checks for saliency maps. In *Advances in Neural Information Processing Systems*, 9505–9515.
- Ahn, J.; Cho, S.; and Kwak, S. 2019. Weakly supervised learning of instance segmentation with inter-pixel relations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2209–2218.
- 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.
- Bulthoff, H. H.; Lee, S.-W.; Poggio, T.; and Wallraven, C. 2003. *Biologically motivated computer vision*. Springer-Verlag.
- 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.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *International journal of computer vision* 88(2): 303–338.
- Fong, R.; Patrick, M.; and Vedaldi, A. 2019. Understanding deep networks via extremal perturbations and smooth masks. In *Proceedings of the IEEE International Conference on Computer Vision*, 2950–2958.
- Gu, J.; Yang, Y.; and Tresp, V. 2018. Understanding individual decisions of cnns via contrastive backpropagation. In *Asian Conference on Computer Vision*, 119–134. Springer.
- Guillaumin, M.; Küttel, D.; and Ferrari, V. 2014a. Imagenet auto-annotation with segmentation propagation. *International Journal of Computer Vision* 110(3): 328–348.
- Guillaumin, M.; Küttel, D.; and Ferrari, V. 2014b. ImageNet Auto-Annotation with Segmentation Propagation. *International Journal of Computer Vision* 110: 328–348.
- Huang, Z.; Wang, X.; Wang, J.; Liu, W.; and Wang, J. 2018. Weakly-supervised semantic segmentation network with deep seeded region growing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7014–7023.
- Kindermans, P.-J.; Schütt, K. T.; Alber, M.; Müller, K.-R.; and Dähne, S. 2017. PatternNet and PatternLRP—Improving the interpretability of neural networks. *arXiv preprint arXiv:1705.05598*.
- Krähenbühl, P.; and Koltun, V. 2011. Efficient inference in fully connected crfs with gaussian edge potentials. In *Advances in neural information processing systems*, 109–117.
- 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. doi:10.1038/s41467-019-08987-4. URL <http://dx.doi.org/10.1038/s41467-019-08987-4>.
- Lee, J.; Kim, E.; Lee, S.; Lee, J.; and Yoon, S. 2019. FickleNet: Weakly and Semi-Supervised Semantic Image Segmentation Using Stochastic Inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- 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.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- 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.
- Nam, W.-J.; Gur, S.; Choi, J.; Wolf, L.; and Lee, S.-W. 2019. Relative Attributing Propagation: Interpreting the Comparative Contributions of Individual Units in Deep Neural Networks. *arXiv preprint arXiv:1904.00605*.
- Petsiuk, V.; Das, A.; and Saenko, K. 2018. Rise: Randomized input sampling for explanation of black-box models. *arXiv preprint arXiv:1806.07421*.

- Pinheiro, P. O.; Collobert, R.; and Dollár, P. 2015. Learning to segment object candidates. In *Advances in Neural Information Processing Systems*, 1990–1998.
- Rebuffi, S.-A.; Fong, R.; Ji, X.; and Vedaldi, A. 2020. There and Back Again: Revisiting Backpropagation Saliency Methods. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8839–8848.
- 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.
- Roh, M.-C.; Kim, T.-Y.; Park, J.; and Lee, S.-W. 2007. Accurate object contour tracking based on boundary edge selection. *Pattern Recognition* 40(3): 931–943.
- Roh, M.-C.; Shin, H.-K.; and Lee, S.-W. 2010. View-independent human action recognition with volume motion template on single stereo camera. *Pattern Recognition Letters* 31(7): 639–647.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, 618–626.
- 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. *IEEE transactions on pattern analysis and machine intelligence* 41(11): 2677–2692.
- Yang, H.-D.; and Lee, S.-W. 2007. Reconstruction of 3D human body pose from stereo image sequences based on top-down learning. *Pattern Recognition* 40(11): 3120–3131.
- Zeiler, M. D.; and Fergus, R. 2014. Visualizing and understanding convolutional networks. In *European conference on computer vision*, 818–833. Springer.
- Zhang, J.; Bargal, S. A.; Lin, Z.; Brandt, J.; Shen, X.; and Sclaroff, S. 2018. Top-down neural attention by excitation backprop. *International Journal of Computer Vision* 126(10): 1084–1102.
- 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* .
- 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.