

# Towards Better Visualizing the Decision Basis of Networks via Unfold and Conquer Attribution Guidance

Jung-Ho Hong<sup>1†</sup>, Woo-Jeong Nam<sup>2‡</sup>, Kyu-Sung Jeon<sup>1</sup>, Seong-Whan Lee<sup>1\*</sup>

<sup>1</sup>Department of Artificial Intelligence, Korea University, Seoul, Republic of Korea

<sup>2</sup>School of Computer Science and Engineering, Kyungpook National University, Daegu, Republic of Korea  
{jungho-hong, ksjeon, sw.lee}@korea.ac.kr, nwj0612@knu.ac.kr

## Abstract

Revealing the transparency of Deep Neural Networks (DNNs) has been widely studied to describe the decision mechanisms of network inner structures. In this paper, we propose a novel post-hoc framework, Unfold and Conquer Attribution Guidance (UCAG), which enhances the explainability of the network decision by spatially scrutinizing the input features with respect to the model confidence. Addressing the phenomenon of missing detailed descriptions, UCAG sequentially complies with the confidence of slices of the image, leading to providing an abundant and clear interpretation. Therefore, it is possible to enhance the representation ability of explanation by preserving the detailed descriptions of assistant input features, which are commonly overwhelmed by the main meaningful regions. We conduct numerous evaluations to validate the performance in several metrics: i) deletion and insertion, ii) (energy-based) pointing games, and iii) positive and negative density maps. Experimental results, including qualitative comparisons, demonstrate that our method outperforms the existing methods with the nature of clear and detailed explanations and applicability.

## Introduction

Although Deep Neural Networks (DNNs) show remarkable performance in various computer science tasks due to their applicability, the issues of complexity and transparency derived from the non-linearity of the network hinder the clear interpretation of the decision basis. Visual explanations of network decisions alleviate these problems and provide analytical guidance for model development with intuitive explanations. Inspired by this, several approaches have been introduced, such as relevance or gradient-based generation, to describe the image regions that the network currently references to make decisions.

The approaches behind the relevance score generate the explanation based on the decomposition methods, Layer-wise Relevance Propagation (LRP) (Bach et al. 2015). LRP constructs the propagation rule according to the predicted scores of the specific class. Despite the nature of employing relevance score is robust to the shattered gradient problem and

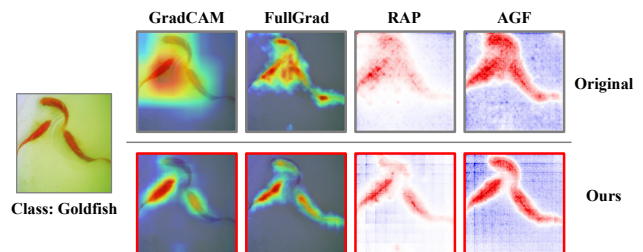


Figure 1: The qualitative results of our approach. The first and second rows represent the existing visualization methods and ours, respectively. Our approach enhances the representation ability of explanation maps, which is applicable to any explanation technique.

success in producing the input-level explanation map, the class-discriminateness within their representation is still limited. Recently, although Attribution Guided Factorization (AGF) enhances the class discriminability of the explanation by integrating both gradient/relevance-based methodologies, it still provides biased attribution toward the class.

GradCAM (Selvaraju et al. 2017) computes the coefficients for each map which are averaged over the gradients of the activation maps, and generates a class activation map by linearly combining the maps with the corresponding coefficients to provide a visual explanation. Although the explanation map generated from GradCAM successfully provides the class-discriminative explanation, the provided explanations are still coarse and contain numerous non-saliency regions due to the mechanism of utilizing the source map from the deepest layer, which keeps the lowest resolution activation map.

The compressive nature of DNNs limits the representation ability of relevance/gradient-based approaches. As the features in the deep layer of the network hold the abbreviated information, it restricts the propagated regions to represent abundant saliencies. Thus, the resulting saliency maps produced from the gradient or relevance approach are represented with limited quality, with obscured class-discriminative regions. In this paper, we propose Unfold-and-Conquer Attribution Guidance (UCAG) framework to guide the explaining methods to generate better saliencies. It improves interpretability by investigating the detailed spa-

\*Corresponding Author, † Equal Contribution

tial regions with a set of subregions. Figure 1 illustrates the intuitive examples of our method, representing the strong advantages of detailed and abundant descriptions agnostic to models and methods. To fulfill this, UCAG is divided into three sections: spatial unfoldment, partial saliency generation, and conquer with geometrical aggregation. First, the unfoldment procedure provides a series of patches delivered by deploying the original data from a spatial perspective. The attribute method independently computes the partial explanation map. During this procedure, UCAG accepts any specific type of methodology as an explanation module, including explaining mechanisms and model types. Finally, the conquering procedure aggregates the partial saliency maps to produce the final explanatory map. Therefore, UCAG encourages an explanation method that scrutinizes an image through the spatial axis.

The key contributions of this paper are summarized as:

- We propose a new post-hoc framework: Unfold and Conquer Attribution Guidance (UCAG), to improve the interpretability of explanation methods by spatially scrutinizing the input data. Spatial scrutinization and integration provide multiple partial semantic views, leading to exploring the obscured saliencies due to the compressing nature of DNNs.
- Post-hoc manner of UCAG has the advantage of being easily applicable to various relevance/gradient-based approaches and DNN models. We present the applicability of UCAG that successfully provides a function of post-hoc refinement with intuitive and detailed descriptions of explanation.
- We validate the quality of the saliency map in density, localization, and causality perspectives. We conduct the verified experiments: i) positive and negative map density (Jalwana et al. 2021), ii) energy-based and original pointing game (Wang et al. 2020), and iii) insertion and deletion (Petsiuk, Das, and Saenko 2018) to confirm the perspectives, respectively. The results demonstrate that our method outperforms the existing explaining methods with a sizable gap in terms of indicating the better points in the input correlated with the output decision.

## Related Work

Gradient-based approaches (Simonyan, Vedaldi, and Zisserman 2013; Zeiler and Fergus 2014; Springenberg et al. 2014; Selvaraju et al. 2017; Chattopadhyay et al. 2018; Srinivas and Fleuret 2019; Rebuffi et al. 2019; Fu et al. 2020; Rebuffi et al. 2020; Zhang, Rao, and Yang 2021) construct the saliency map derived from the gradients. In the manner of employing class activation maps, which employs the activation map to generate the saliency map, GradCAM (Selvaraju et al. 2017) alleviates the limitations of CAM that require a global average pooling and re-training process. They construct an explanatory map by linearly combining the activation map and the coefficients with the coefficients computed by taking the average of the gradients corresponding to the activation map. XGradCAM (Fu et al. 2020) introduces the way to generate the class activation map while obeying two theoretical axioms: sensitivity and conservation. WGradCAM (Zhang,

Rao, and Yang 2021) enhances the representation ability of the explanation map in terms of localization.

Several approaches (Fong, Patrick, and Vedaldi 2019; Jiang et al. 2021; Jalwana et al. 2021; Zhang, Rao, and Yang 2021) turned the literature to amplifying the quality of saliency maps to produce a more precise saliency map. Extremal perturbation (Fong, Patrick, and Vedaldi 2019) generates the saliency map following the searching nature of perturbation-based approaches. LayerCAM (Jiang et al. 2021) aggregates the information from earlier layers and introduces the way of deciding the coefficients of each source. CAMERAS (Jalwana et al. 2021) fuses the multi-resolution images to generate a high-resolution explanation while obeying the model fidelity. SigCAM (Zhang, Rao, and Yang 2021) combines perturbation/gradient-based approaches to produce more sparse and precise explanations using learnable external parameters.

Relevance-based approaches (Bach et al. 2015; Zhang et al. 2018; Gu, Yang, and Tresp 2018; Iwana, Kuroki, and Uchida 2019; Nam et al. 2020; Nam, Choi, and Lee 2021; Gur, Ali, and Wolf 2021) produce the granular explanation by decomposing the output predictions in a backward manner. LRP (Bach et al. 2015) introduces the propagation rules that decompose the output logits into the relevance scores while maintaining the conservation rule. As the recent research, RAP (Nam et al. 2020) aims to separate the relevant and irrelevant parts of input by allocating the bi-polar relevance scores during propagation. RSP (Nam, Choi, and Lee 2021) extends the rule of RAP to overcome the limitations of suppressing nature, resulting in solving the class-specific issues. AGF (Gur, Ali, and Wolf 2021) combines the accumulated gradients and relevance scores across the layers to correct the saliency bias.

Despite the superior performance of the previous approach, the generated explanation maps represent insufficient saliency and diminished class discrimination. To increase the fidelity of the explanation maps, we consider the schemes of spatial scrutinization and integration by providing multiple partial information from a single input image, leading to discovering the obscured regions while maintaining class-discriminative saliencies.

## Saliency Shedding

Propagation-based approaches provide a saliency map with the same size as the input image by propagating the relevance score with respect to the network decision. However, it is difficult to clarify the correlations of neurons among the whole layers, resulting in inevitably assigning positive attributions to unrelated pixels. On another side, since the CAM-based approach provides the explanation based on the activation map of the deepest feature extraction layer and gradients, the computed saliency map represents a coarse explanation and lacks details due to the upsampling. To provide more detailed saliency maps, upscaling techniques of the input image are introduced in (Jalwana et al. 2021) to increase the resolution of the class activation map.

However, as shown in our empirical experiments of Figure 2, employing the upsampled images both decreases the deletion and insertion scores proportional to the resolution,

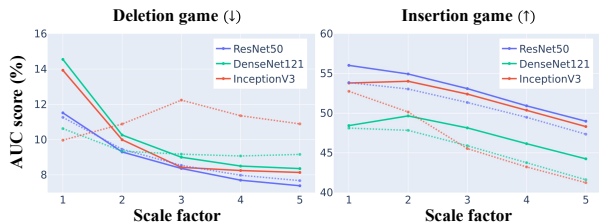


Figure 2: The motivations of our work. The scores of deletion and insertion games among various resolutions with Grad-CAM on the ImageNet dataset. We scaled the size of image from (224, 224) to (1120, 1120) utilizing the scale factor  $\alpha$ . The performance of the original and our methods are represented as dots and solid lines, respectively. We aim to present a resolution-agnostic method with the enhancement of interpretability.

which are originally in the trade-off relationships. Deletion and insertion scores are widely used metrics in the field of explainability (Petsiuk, Das, and Saenko 2018; Wang et al. 2020; Zhang, Rao, and Yang 2021; Jung and Oh 2021), to evaluate whether the explanation map successfully covers the more fine-grained representation and localizes the actual object better. Decreasing insertion scores according to the upscaling states that the visualized explanation loses saliency areas due to the distortion of neuron activation and corresponding gradients.

Motivated by this issue, we address the method to improve the quality of saliency maps by investigating the correlations among the input unfolding and confidence scores, leading to maintaining the salient regions while increasing the granularity of the explanation map.

### Unfold and Conquer Attribution Guidance

To enhance the granularity of the explanation without losing saliencies, our work starts from the assumption that if the highlighted regions obviously contribute to the decision, their saliency should remain regardless of the geometrical translation. To fulfill this, we extend the aspect of the convolution process generally used in the DNNs framework. Hence, we unfold the image into a series of segments and let the visualization method infer each fragment independently to yield partial explanation maps. By folding partial explanations weighted with corresponding confidence scores, the unified explanation contains the saliencies referring to the various perspectives of spatial views. Figure 3 illustrates the overview of our method: UCAG. A detailed description of each stage is described in each section.

### Spatial Unfoldment

In this section, we demonstrate a strategy for providing the unfolded segments along the spatial axis to construct the set of input materials for the explanation method. To accomplish this, we first declare the hyper-parameter  $n$ : the number of segmented patches. Then, we construct an unfolding function:  $\mathbb{U}$ , which striding the input image  $I \in \mathbb{R}^{C \times H \times W}$  into the series of  $n^2$  patches  $I_p$  by walking along the

spatial (weight and height) axis. The set of patches includes the partial image  $I_p^j \in \mathbb{R}^{C \times H' \times W'}$ , indexed by  $j \in \{0, 1, \dots, n^2 - 1\}$ . The size of patch ( $H', W'$ ) is computed with the specific resizing factor  $\rho \in [0, 1]$ , described as  $(\rho H, \rho W)$ .

However, from the perspective of the network, utilizing the downscaled input leads to representing coarse output features after passing through the network feature extraction. To relieve this, we enhance the resolution of each patch with the scaling function  $\Psi(I_p, \alpha)$ , which scales the set of patches into the size of  $(\alpha H', \alpha W')$ . In summary, the unfoldment procedure provides partial explanations as follows:

$$I_s = \Psi(\mathbb{U}(I, \rho, n), \alpha). \quad (1)$$

Hence, the input image  $I$  is unfolded into the set of patches  $I_s$  with a size of  $(n^2, C, \alpha H', \alpha W')$  and each patch shares the overlapped areas with the striding mechanisms following the nature of unfoldment. From the spatial perspective, a set of patches holds different subregions from the single image and enables a network to review the various views.

### Patch-wise Saliency Generation

In this section, we carefully revisit the general explaining methods:  $\mathbb{E}$  with the two streams of gradient/attribution and clarify the applying process to our post-hoc frameworks.

**Gradient-based saliency** To generate the saliency map revealing network decisions, gradient-based methods utilize the activation map in the defined layers and their corresponding neuron importance derived from the gradients. Generally, the activation map is obtained from the last layer of DNN before the linear or global average pooling layer to investigate the compressed feature importance. Several CAM approaches introduce their intrinsic way of computing the class-specific coefficient by utilizing the gradient of each pixel in the activations. Generic gradient-based saliency map (Selvaraju et al. 2017) is computed as follows:

$$\mathbb{E}_{grad}(I, f, k) = \text{ReLU} \left( \sum_c w_c^k A_c \right). \quad (2)$$

Here,  $c$ ,  $f$ , and  $k$  denote the index of the channel axis, the kind of DNNs, and the index of the target class, respectively.  $w_c^k$  is the neuron importance of each feature map derived from the gradients with respect to the target class. According to the provided assumption of existing approaches, e.g. (Selvaraju et al. 2017; Chattopadhyay et al. 2018; Fu et al. 2020; Zhang, Rao, and Yang 2021), there is a slight difference in deriving the  $w_c^k$ , but mostly dependent on the chain rule  $\partial f(I)^k / \partial A$ . Detailed expansion of each method is described in supplementary materials. In our post-hoc manners: using the unfolded segments  $I_s$  as input, partial explanations are generated independently with the activations computed from each patch as  $\mathbb{E}_{grad}(I_s, f, k)$ .

**Propagation-based saliency** Propagation-based methods are capable of providing the input-level explanation map, allowing for the representation of pixel-level attribution. They compute the relevance score by layer-wisely propagating the

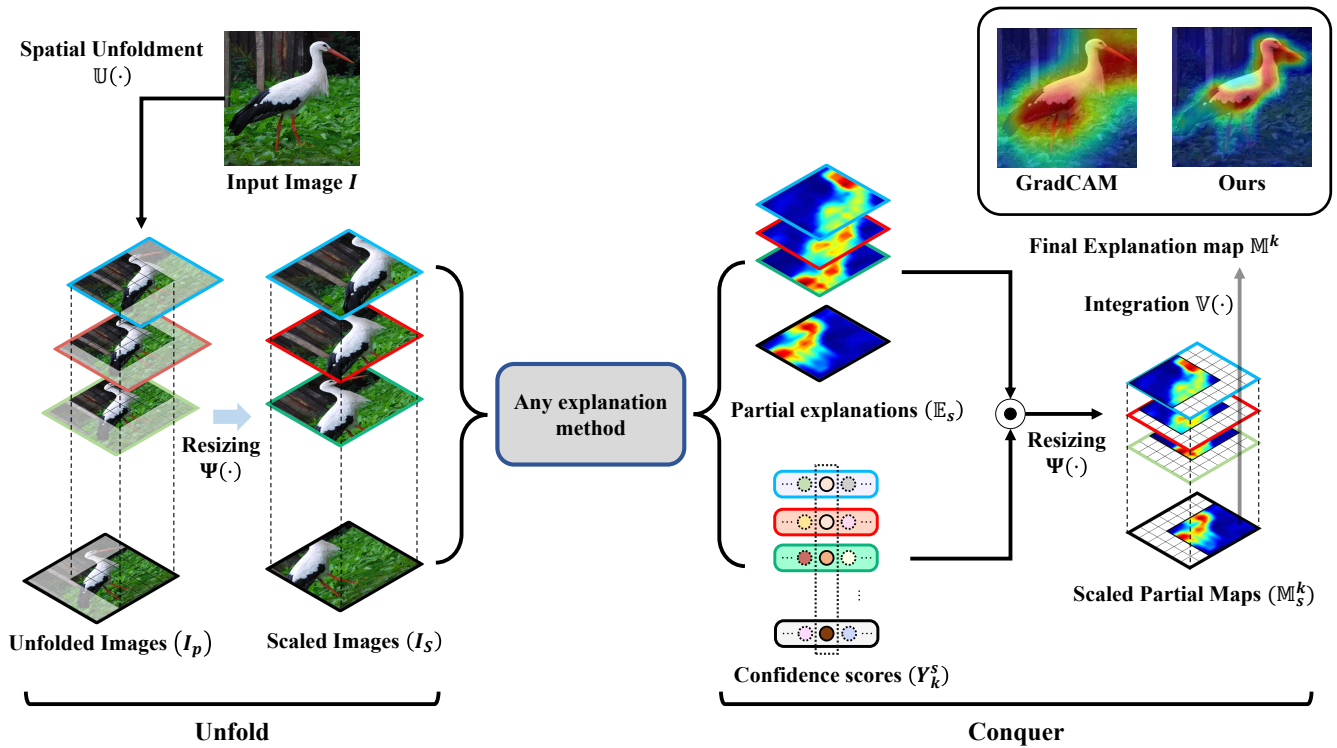


Figure 3: An overview of UCAG. The entire course is divided into three streams: i) unfoldment, ii) patch-wise saliency generation, and iii) conquer with aggregation. Through the unfolding process, the unfolded patches are individually delivered to the gradient/propagation-based visualization methods. Conquering is a procedure for integrating segmented explanations into a single explanation map with the corresponding confidence scores.

output logits to input in a backward manner. During the propagation, the whole relevance sum  $R$  is maintained with the conservation rule (Bach et al. 2015), acting as the evidence of the decision. The generic propagation rule (Montavon et al. 2017) is defined as follows:

$$\begin{aligned} \mathbb{E}_{att}(I, f, k) &= \mathcal{G}(I, \theta, R^{(l-1)}) \\ &= \sum_i A_j \frac{\partial f_i^{(l)}(A, \theta)}{\partial A_j} \frac{R_i^{(l-1)}}{\sum_{i'} f_{i'}^{(l)}(A, \theta)}. \end{aligned} \quad (3)$$

Here,  $l$  indexes the middle layer of DNNs, and  $A$  represents the activation of corresponding layers.  $i$  and  $j$  denote the neuron index in the  $l$  and  $l-1$  layers, respectively. During the propagation, the whole relevance sum is maintained as  $\sum_j R_j^{(l)} = \sum_i R_i^{(l-1)}$ , starting from the output one-hot vector:  $f(I)^k$  to the input layer.

Several kinds of propagation-based methods, e.g., SGLRP, c\*LRP, RSP, and AGF are mainly based on Eq. 3, but have some techniques, such as contrastive perspective, uniform shifting, and factorization, to enhance the quality of attribution maps. Our method is independent of each of these unique techniques as  $\mathbb{E}_{att}(I_s, f, k)$  and could be utilized for intuitive and clear visualizations.

### Conquer with Geometrical Aggregation

As the unfolded images do not guarantee to contain the actual object, it is necessary to discriminate the patches according to the proportion of the actual class included in the provided subregion. To fulfill this, we employ the confidence scores predicted by the model as the source of judging the validity of the partial explanations. The score for the target class is defined as follows:

$$\begin{aligned} Y_s^k &= \xi(\mathcal{L}_s, k) \\ &= \exp\left(\frac{\exp(\mathcal{L}_s^{k,j})}{\sum_{i=1}^N \exp(\mathcal{L}_s^{i,j})}\right), \forall j. \end{aligned} \quad (4)$$

Here,  $j \in [0, n^2 - 1]$  indexes the logit,  $\mathcal{L}_s$  is the output logit:  $f(I_s)$  and  $\xi(\cdot)$  returns the refined score with respect to the class  $k$  by individually taking the exponential function to the softmax probability of  $\mathcal{L}_s^j$ , leading to expanding the gap of the partial confidence scores.  $N$  represents the number of classes. The final explanation is aggregated by applying the element-wise product among each partial explanation and the confidence scores. The weighted partial explanations are scaled to the original patch size:  $(n^2, 1, H', W')$  by the scaling factor  $\beta$  to return to the original areas of the input image. In summary, the partial explanation before the folding function is computed as follows:

$$M_s^k = \Psi(\mathbb{E}_s^k \odot Y_s^k, \beta) \quad (5)$$

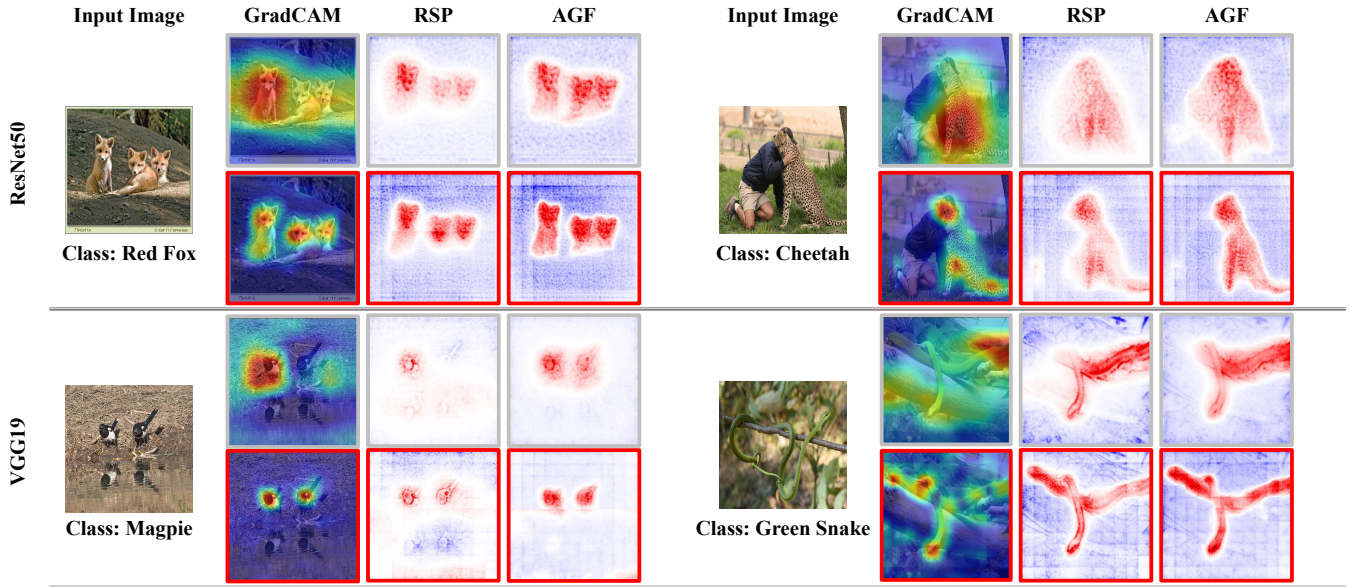


Figure 4: Qualitative comparisons in VGG-19 and ResNet-50. UCAG improves the quality of the saliency map and has an advantage of agnostic to model and method. First and second rows in each group represent the original and our (marked as red) results, respectively. Additional results are illustrated in supplementary materials.

Then, the scaling function  $\Psi(\cdot)$  is applied to interpolate the series of explanations with the size of  $(H', W')$ . Note that the map  $\mathbb{E}_s^k$  is not normalized or clipped (eliminating the negative scores) to gather a more abundant range of explanations. Furthermore, since the partial explanations share the overlapped portions from the single image, the frequency of the pixels in the saliency map is not accumulated evenly, destabilizing the resulting saliency map. For this reason, we devise the duplication matrix  $\Gamma$ , which counts the duplication of each pixel. By utilizing this count, we average each pixel in the saliency map according to its frequency of occurrence.

Intuitively, the integration procedure is similar to the reverse of the unfolding process except for the average with duplication matrix. Including the frequency normalization with the duplication matrix  $K$ , we notate the integration procedure to generate the final saliency as follows:

$$\mathbb{M}^k = \mathbb{V}(\mathbb{M}_s^k, \Gamma_s). \quad (6)$$

The final saliency map has the same size as the input image, representing the relatively relevant scores in each pixel with respect to the output predictions of DNNs. The overall streams of our method are described in Algorithm.1.

## Experimental Evaluation

### Implementation Details

We utilize the public accessible datasets: ImageNet-1k (Deng et al. 2009), Pascal VOC 2002 (Everingham et al. 2010), and MS COCO (Lin et al. 2014). For the localization and segmentation assessments, we utilize the ImageNet segmentation dataset (Guillaumin, Küttel, and Ferrari 2014). For fair comparisons in evaluations, we employ the online available models: VGG19 (Simonyan and Zisserman 2014),

---

### Algorithm 1: Saliency generation with UCAG

---

**Require:** Image:  $I$ , Network:  $f(\cdot)$ , Explanation method:  $\mathbb{E}$   
Scaling factor:  $\{\rho, \alpha, \beta\}$ , Target class:  $k$   
**Unfolding:**  $\triangleright$  Deploy a set of segments  
1:  $I_p = \mathbb{U}(I, \rho, n)$   $\triangleright$  Striding pass  
2:  $I_s = \Psi(I_p, \alpha)$   $\triangleright$  Scaling pass  
3:  $\mathcal{L}_s \leftarrow f(I_s)$   $\triangleright$  Forward pass  
**Patch-wise Explanations:**  $\triangleright$  Generating saliencies  
4: **if**  $\mathbb{E}$  is Gradient-based **then**  
5:  $\mathbb{E}_s \leftarrow \mathbb{E}_{grad}(I_s, f, k)$   $\triangleright$  Utilize gradients  
6: **else if**  $\mathbb{E}$  is Propagation-based **then**  
7:  $\mathbb{E}_s \leftarrow \mathbb{E}_{att}(I_s, f, k)$   $\triangleright$  Utilize propagations  
8: **end if**  
**Conquer:**  $\triangleright$  Integrating partial explanations  
9:  $Y_s^k \leftarrow \xi(\mathcal{L}_s, k)$   $\triangleright$  Compute confidence scores  
10:  $\mathbb{M}_s^k \leftarrow \Psi(\mathbb{E}_s^k \odot Y_s^k, \beta)$   $\triangleright$  Aggregation and Scaling  
11:  $\mathbb{M}^k \leftarrow \mathbb{V}(\mathbb{M}_s^k, \Gamma_s)$   $\triangleright$  Generating final explanation

---

DenseNet121 (Huang et al. 2017), ResNet50 (He et al. 2016), and InceptionV3 (Szegedy et al. 2016).

We utilized several metrics: i) insertion and deletion games, ii) (energy-based) pointing game, iii) mean Average Precision (mAP), and iv) positive and negative map density to measure the degree of localization, causality, and density. A detailed comparison is described in each section. We set the hyperparameter for the number of segments and rescaling factor as  $(\alpha = 2.6, n = 6, \rho = 0.555)$ . Comparisons according to the hyper-parameter are provided in the supplementary materials.

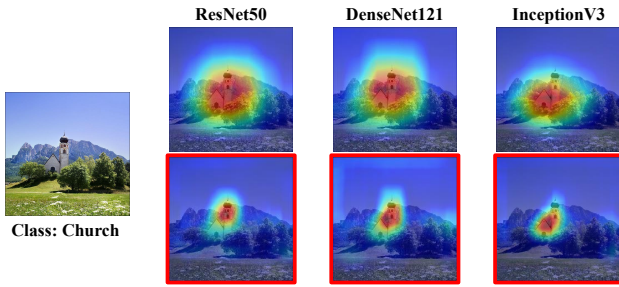


Figure 5: Qualitative results between multiple networks including ResNet50, DenseNet121, and InceptionV3. The top and bottom rows represent visualized heatmaps generated by GradCAM with and without our approach, respectively. More results are contained in the supplementary material.

## Qualitative Results

In terms of explainability, it is crucial to provide intuitive and convincing visualizations from a human view. To demonstrate how our method improves the quality of generated explanations, we compare the results with the representative explanation methods in Figure 4. The first and second rows represent the results of the original and our methods, respectively. Red and blue colors denote the positive and negative relevance scores, respectively. While existing methods have the advantage of depicting the highly relevant regions to the output predictions, they still lack providing fine-grained descriptions with the limited granularity of representing areas of contribution.

In the manner of the post-hoc framework, UCAG has the strong advantages of fine-grained visualization, easily applicable and agnostic to model variations. In visualization, it improves the quality of existing maps by revising the incorrect negative or positive relevance scores in the (ir)relevant area without losing salient regions, resulting in increasing the density of the explanation map. Figure 5 shows the improvement when UCAG is applied to various models: ResNet, DenseNet, and InceptionNet. These visual enhancements also lead to improved quantitative performance. We report the additional visual explanations and evaluations on various methods and models, including GradCAM, FullGrad, GradCAM++, WGradCAM, RAP, RSP, AGF, SGLRP, and transformer interpretability (Chefer, Gur, and Wolf 2021).

## Quantitative Results

**Causality assessment** The deletion and insertion game measures whether the explanation map correctly highlights areas contributing to the decision-making process. Deletion measures the model’s response by repeatedly deleting the high-priority pixels with respect to the output. When deleted pixels contribute significantly to predictions, model confidence and deletion scores decrease dramatically. In contrast, the insertion game measures how the confidence score would be increased. We set the baseline image blurred with kernel size: 51 and  $\sigma$ : 50 for the insertion game and images with zero pixels for the deletion game.

Table 1 displays the scores of class activation maps with-

Methods	ResNet50	DenseNet121	InceptionV3
GCAM	11.3/53.9	10.6/48.1	10.0/52.8
GCAM++	11.6/52.7	10.9/47.2	10.1/52.0
WGCAM	10.1/52.1	10.8/47.7	10.1/52.1
<b>GCAM*</b>	<b>8.6/53.4</b>	<b>8.9/49.0</b>	<b>8.59/53.4</b>
<b>GCAM++*</b>	<b>8.7/52.7</b>	<b>9.2/48.0</b>	<b>8.86/51.9</b>
<b>WGCAM*</b>	<b>9.1/52.8</b>	<b>9.3/48.6</b>	<b>9.08/52.5</b>

Table 1: The AUC scores regarding the deletion (lower is better)/insertion (higher is better) games on the ImageNet dataset. Mark \* represents the performance of applying our methods.

Methods	VOC07 Test		COCO14 Val	
	VGG16	ResNet50	VGG16	ResNet50
Center	69.6/42.4	69.6/42.4	27.8/19.5	27.8/19.5
Gradient	76.3/56.9	72.3/56.8	37.7/31.4	35.0/29.4
DeConv	67.5/44.2	68.6/44.7	30.7/23.0	30.0/21.9
Guid	75.9/53.0	77.2/59.4	39.1/31.4	42.1/35.3
MWP	77.1/56.6	84.4/70.8	39.8/32.8	49.6/43.9
cMWP	79.9/66.5	90.7/82.1	49.7/44.3	58.5/53.6
RISE	86.9/75.1	86.4/78.8	50.8/45.3	54.7/50.0
GradCAM	86.6/74.0	90.4/82.3	54.2/49.0	57.3/52.3
Extremal	88.0/76.1	88.9/78.7	51.5/45.9	56.5/51.5
NormGrad	81.9/64.8	84.6/72.2	-	-
CAMERAS	86.2/76.2	94.2/88.8	55.4/50.7	69.6/66.4
<b>Ours</b>	<b>91.1/82.8</b>	<b>94.2/89.4</b>	<b>61.8/57.6</b>	<b>71.0/67.6</b>

Table 2: The performance of the pointing game among various methods. Our method (applied to the GradCAM) represents a sizeable increment in performance compared to the other method.

/without UCAG. The results show that the saliency map passed through the UCAG yields a decreased deletion score while the insertion score is maintained or increased. Specifically, the spatial scrutinization in UCAG successfully induces the outcomes to be more precise while preserving saliency.

**Localization assessment** Evaluating localization performance is employed to assess the quality of the explanation, assuming that the highlighted pixels contain the main object of the input image. (Zhang et al. 2018) introduced the Pointing Game, widely used in many related researches, that computes whether the maximum point of the saliency map is contained in the label. We report our results (GradCAM with UCAG) in Table 2, and it outperforms the state-of-art explanatory approaches.

ScoreCAM (Wang et al. 2020) raised the issue that the pointing game does not provide a detailed measure of localization performance since it only considers the maximum saliency point inside the explanation. As an alternative, ScoreCAM introduces an energy-based pointing game to evaluate the quality with a more detailed measurement. The energy-based pointing game is based on the summation of the regions where saliency overlaps the inside label. We also report the comparisons of the energy-based pointing

Model	Dataset	GCAM	GCAM++	XGCAM	WGCAM
VGG	VOC07 Test	56.6(40.1)/ <b>59.7(43.6)</b>	53.7(36.0)/ <b>61.4(47.6)</b>	55.6(39.6)/ <b>59.3(43.2)</b>	55.6(37.4)/ <b>56.5(38.5)</b>
	COCO14 Val	25.5(19.6)/ <b>29.9(23.8)</b>	22.4(16.5)/ <b>29.2(23.3)</b>	25.0(19.2)/ <b>28.6(22.6)</b>	24.4(18.1)/ <b>26.8(20.3)</b>
ResNet	VOC07 Test	61.3(45.3)/ <b>74.0(66.8)</b>	52.9(31.8)/ <b>63.9(47.6)</b>	61.3(45.3)/ <b>74.0(66.8)</b>	55.9(36.8)/ <b>64.2(44.9)</b>
	COCO14 Val	27.5(21.1)/ <b>42.7(37.4)</b>	21.9(15.2)/ <b>31.4(25.0)</b>	27.5(21.1)/ <b>42.7(37.4)</b>	23.1(16.5)/ <b>35.4(29.6)</b>

Table 3: The performance of the energy-based pointing game represented as original / ours. (–) denotes the result in the difficult set of each dataset.

Model	Metric	FullGrad	LRP	c*LRP	SGLRP	RAP	RSP	AGF
VGG	mAP	<b>0.789</b> /0.783	<b>0.700</b> /0.682	0.547/ <b>0.559</b>	0.551/ <b>0.570</b>	0.735/ <b>0.743</b>	0.743/ <b>0.752</b>	0.740/ <b>0.752</b>
	Pixel Acc.	<b>0.742</b> /0.737	<b>0.750</b> /0.730	0.534/ <b>0.556</b>	0.543/ <b>0.587</b>	0.786/ <b>0.789</b>	0.793/ <b>0.796</b>	0.779/ <b>0.794</b>
ResNet	mAP	0.748/ <b>0.774</b>	0.717/ <b>0.722</b>	0.671/ <b>0.694</b>	0.588/ <b>0.617</b>	0.711/ <b>0.761</b>	0.719/ <b>0.766</b>	0.758/ <b>0.783</b>
	Pixel Acc.	0.798/ <b>0.821</b>	0.770/ <b>0.774</b>	0.716/ <b>0.736</b>	0.608/ <b>0.648</b>	0.763/ <b>0.809</b>	0.775/ <b>0.813</b>	0.804/ <b>0.826</b>

Table 4: Performance of mAP and outside–inside ratio over ImageNet segmentation dataset (Original / Ours). Propagation-based methods that have the class-discriminative ability are compared. A Low Out/In ratio denotes attributions are intensively distributed in the semantic mask.

Model	GCAM	CAMERAS	Ours
Positive map density ( $\mathbb{D}_{map}^+ \uparrow$ )			
ResNet50	2.33	3.20	<b>3.89</b>
DenseNet	2.35	3.23	<b>3.45</b>
Inceptionv3	2.18	3.15	<b>3.83</b>
Negative map density ( $\mathbb{D}_{map}^- \downarrow$ )			
ResNet50	0.86	0.81	<b>0.81</b>
DenseNet	0.94	0.83	<b>0.82</b>
Inceptionv3	1.04	0.93	<b>0.85</b>

Table 5: Evaluated results of positive (higher is better) and negative (lower is better) map density.

game in Table 3 with the results of various CAM approaches. The increased scores indicate that UCAG consistently enhances the ability of various CAM approaches regardless of the type of network and method.

Evaluating the weakly supervised segmentation is closely related to interpretability in terms of finding the salient objects in the input image. In recent works (Nam et al. 2020; Nam, Choi, and Lee 2021; Gur, Ali, and Wolf 2021), mean average precision (mAP) is utilized to assess the quality of interpretability. We follow the experiential settings with the ImageNet segmentation dataset (Guillaumin, Küttel, and Ferrari 2014). Table 4 shows the comparisons of existing methods and ours. In most cases, our method improves the quality of interpretations in both VGG and ResNet models.

**Density assessment** Several previous works (Jalwana et al. 2021; Poppi et al. 2021) pointed out that only reporting the ratio of (de)increase in confidence score does not fully reflect the interpretability of explaining methods. As an alternative, CAMERAS (Jalwana et al. 2021) introduces the positive (negative) map density, which simultaneously evaluates the variations of confidence scores and saliency density during

the perturbation. Metrics for positive (negative) map density are defined as  $\mathbb{D}_{map}^+ = f(I \odot \mathbb{M}^k) / \sum_i \sum_j \mathbb{M}_{(i,j)}^k \times (h \times w)$  and  $\mathbb{D}_{map}^- = f(I \odot 1 - \mathbb{M}^k) / \sum_i \sum_j (1 - \mathbb{M}_{(i,j)}^k) \times (h \times w)$ . Therefore, it is possible to measure the quality of the explanation in a more precise way.

The results in Table 5 show that UCAG enhances the density of saliency maps more than the existing multi-resolution fusion approach. This implies that our approach has the advantage of strong objectness with concentrated scores in the saliency regions.

## Ablation and Discussion

As an ablation study, the gap of effectiveness between upscaling and spatial scrutinization is described in Figure 2. UCAG provides a more precise explanation with the tendency to preserve the areas contributing to the decision-making process.

Here, we note that the computational cost is inevitably increased due to the manner of post-hoc framework. However, it is more efficient to find various spatial saliency with the scrutinization rather than the cost of calculating the exponentially increasing image resolution. In addition, we implement the overall frameworks batch-wisely, leading to enhancing computational efficiency.

## Conclusion

In this paper, we propose a post-hoc framework: UCAG, to enhance the interpretability of network decisions by guiding the explaining methods to carefully investigate the saliencies with spatial scrutinization and aggregation. We demonstrate the validity of UCAG from qualitative and quantitative results. With the advantage of being agnostic to the model and method, the increments of causality, localization, and density demonstrate the effectiveness of UCAG for producing a better visual representation ability of network decisions.

## Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2022-0-00984, Development of Artificial Intelligence Technology for Personalized Plug-and-Play Explanation and Verification of Explanation & No. 2019-0-00079, Artificial Intelligence Graduate School Program, Korea University)

## References

- 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.
- Chattopadhyay, A.; Sarkar, A.; Howlader, P.; and Balasubramanian, V. N. 2018. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In *2018 IEEE winter conference on applications of computer vision (WACV)*, 839–847. IEEE.
- Chefer, H.; Gur, S.; and Wolf, L. 2021. Transformer interpretability beyond attention visualization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 782–791.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- 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/CVF international conference on computer vision*, 2950–2958.
- Fu, R.; Hu, Q.; Dong, X.; Guo, Y.; Gao, Y.; and Li, B. 2020. Axiom-based grad-cam: Towards accurate visualization and explanation of cnns. *arXiv preprint arXiv:2008.02312*.
- 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. 2014. Imagenet auto-annotation with segmentation propagation. *International Journal of Computer Vision*, 110(3): 328–348.
- Gur, S.; Ali, A.; and Wolf, L. 2021. Visualization of supervised and self-supervised neural networks via attribution guided factorization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 11545–11554.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Huang, G.; Liu, Z.; Van Der Maaten, L.; and Weinberger, K. Q. 2017. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4700–4708.
- Iwana, B. K.; Kuroki, R.; and Uchida, S. 2019. Explaining convolutional neural networks using softmax gradient layer-wise relevance propagation. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, 4176–4185. IEEE.
- Jalwana, M. A.; Akhtar, N.; Bennamoun, M.; and Mian, A. 2021. CAMERAS: Enhanced resolution and sanity preserving class activation mapping for image saliency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16327–16336.
- Jiang, P.-T.; Zhang, C.-B.; Hou, Q.; Cheng, M.-M.; and Wei, Y. 2021. Layercam: Exploring hierarchical class activation maps for localization. *IEEE Transactions on Image Processing*, 30: 5875–5888.
- Jung, H.; and Oh, Y. 2021. Towards better explanations of class activation mapping. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1336–1344.
- 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.
- 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.; Choi, J.; and Lee, S.-W. 2021. Interpreting deep neural networks with relative sectional propagation by analyzing comparative gradients and hostile activations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 11604–11612.
- Nam, W.-J.; Gur, S.; Choi, J.; Wolf, L.; and Lee, S.-W. 2020. Relative attributing propagation: Interpreting the comparative contributions of individual units in deep neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 2501–2508.
- Petsiuk, V.; Das, A.; and Saenko, K. 2018. Rise: Randomized input sampling for explanation of black-box models. *arXiv preprint arXiv:1806.07421*.
- Poppi, S.; Cornia, M.; Baraldi, L.; and Cucchiara, R. 2021. Revisiting the evaluation of class activation mapping for explainability: A novel metric and experimental analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2299–2304.
- Rebuffi, S.-A.; Fong, R.; Ji, X.; Bilen, H.; and Vedaldi, A. 2019. NormGrad: Finding the pixels that matter for training. *arXiv preprint arXiv:1910.08823*.
- 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.
- 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.



- Simonyan, K.; Vedaldi, A.; and Zisserman, A. 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Springenberg, J. T.; Dosovitskiy, A.; Brox, T.; and Riedmiller, M. 2014. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*.
- Srinivas, S.; and Fleuret, F. 2019. Full-gradient representation for neural network visualization. *Advances in neural information processing systems*, 32.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2818–2826.
- Wang, H.; Wang, Z.; Du, M.; Yang, F.; Zhang, Z.; Ding, S.; Mardziel, P.; and Hu, X. 2020. Score-CAM: Score-weighted visual explanations for convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 24–25.
- 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.
- Zhang, Q.; Rao, L.; and Yang, Y. 2021. A novel visual interpretability for deep neural networks by optimizing activation maps with perturbation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 3377–3384.