

Visualizing Color-wise Saliency of Black-Box Image Classification Models

Yuhki Hatakeyama^{1*}, Hiroki Sakuma¹, Yoshinori Konishi¹, and Kohei Suenaga²

SenseTime Japan, 4F, Oike Koto Building, 324 Oikeno-cho, Nakagyo-ku, Kyoto,
Japan

{hatakeyama,sakuma,konishi}@sensetime.jp

Graduate School of Informatics, Kyoto University, 36-1 Yoshida-Honmachi, Sakyo-ku,
Kyoto, Japan
ksuenaga@gmail.com

Abstract. Image classification based on machine learning is being commonly used. However, a classification result given by an advanced method, including deep learning, is often hard to interpret. This problem of interpretability is one of the major obstacles in deploying a trained model in safety-critical systems. Several techniques have been proposed to address this problem; one of which is RISE, which explains a classification result by a heatmap, called a *saliency map*, that explains the significance of each pixel. We propose *MC-RISE* (Multi-Color RISE), which is an enhancement of RISE to take color information into account in an explanation. Our method not only shows the saliency of each pixel in a given image as the original RISE does, but the significance of *color components* of each pixel; a saliency map with color information is useful especially in the domain where the color information matters (e.g., traffic-sign recognition). We implemented MC-RISE and evaluate them using two datasets (GT-SRB and ImageNet) to demonstrate the effectiveness of our methods in comparison with existing techniques for interpreting image classification results.

1 Introduction

As machine learning is widely applied to image classification, there is a surging demand for the methods to explain classification results and visualize it. One can use such an explanation to check whether a trained model classifies images based on a rational and acceptable criterion, by which he or she can convince various stakeholders that the model is readily deployed.

One of the most popular ways of the visualization is by a *saliency map*—a heatmap overlayed on the original image that indicates which part of the image contributes to the classification result [1–14]. Fig. 1 shows an example of a saliency map generated by a method called RISE [4]. Given the image (a) of a road sign “STOP” and a model that indeed classifies this image as a stop sign, RISE generates the saliency map (b), which indicates the part of (a) that




Fig. 1. Visual comparison between the saliency map of RISE [4] and the color sensitivity maps generated by our proposed method (MC-RISE). Best viewed in color.

contributes to the classification result by a heatmap. From this saliency map, we can figure out that the part of the sign surrounding the text “STOP” contributes positively to the classification result.

Among the explanation methods proposed so far, *model-agnostic* techniques such as LIME [1], SHAP [2], and RISE [4] generate a saliency map without accessing the internal information of a model, treating it as a black box. Although these procedures differ in their detail, they all share the following central idea: They compute a saliency map for a classification result by perturbing the given image and observing how the output of the perturbed image changes from the original. Concretely, given an image I , a model M , and its classification result, they compute a saliency map by (1) generating perturbed image I_1, \dots, I_N from I by masking a part of it, (2) computing the classification result $M(I_1), \dots, M(I_N)$ for each perturbed image, and (3) comparing $M(I)$ with each $M(I_i)$. If the output of a perturbed image with a specific part unmasked tends to be the same as $M(I)$, then this part is considered to be important.

Although these model-agnostic explanation techniques give valuable insights about a classification result of an image, there is important information that has been overlooked by these methods: *colors*. They compute a saliency map indicating each pixel’s existence, but they do not take the color information of a pixel into account. A *color-aware* explanation is effective especially in the application domains in which the colors in an image convey important information (e.g., traffic-sign recognition).

This paper proposes an extension of RISE [4] so that it generates a color-aware saliency map. We extend the original RISE so that it can compute a saliency map for each color of a given color set; we call our extended method *MC-RISE (Multi-Color RISE)*. Fig. 1 (c) shows the saliency maps generated by MC-RISE. It consists of five heatmaps, each of which corresponds to the significance of a certain color of each pixel; the associated color is shown above each heatmap. We can figure out that the red color at the peripheral part of the stop sign contributes much to the classification result, whereas the other colors do not contribute as much as red.

We also propose an enhancement to *debias* the original RISE. In our extension, the saliency of the pixels in an input image is close to 0 if it is irrelevant for the classification result. Notice that the heatmaps in Fig. 1 (c) have value close to 0 for the irrelevant pixels, whereas the heatmap generated by the original RISE in Fig. 1 (b) does not contain a pixel of value 0; it is hard to figure out which part of the image is irrelevant to the classification result from the latter saliency map.

The contributions of the paper are summarized as follows. (i) We propose a new model-agnostic explanation method for an image classifier, *MC-RISE*, which generates *color-aware* saliency maps for the classifier’s decision. (ii) To improve the interpretability of the saliency map, we propose a method to *debias* the saliency map of RISE and incorporate it into MC-RISE. (iii) We qualitatively and quantitatively compare our method to existing model-agnostic explanation methods (LIME [1] and RISE [4]) in GTSRB [15] dataset and ImageNet [16] dataset, and showed that our method can extract additional information which can not be obtained by the existing methods, such as color-sensitivity.

Due to the space limitation, the proofs of theorems and several experimental results are in a separate supplementary material.

2 Related Work

Explanation-generating methods for an image classifier proposed so far can be categorized into the following two groups: (1) ones that treat a model as a *black box* and (2) ones that treat a model as a *white box*. The former observes the input-output relation in generating an explanation without using the internal information of a model; whereas the methods in the latter group assume that the internal information is accessible. As far as we know, our method is the first explanation-generating method for black-box models that considers color information in generating an explanation.

Many white-box methods compute a saliency map utilizing the classifier’s gradient information. Grad-CAM [5] and its extension Grad-CAM++ [6] use the gradient with respect to an intermediate feature map to obtain class-specific weights for the feature map. Integrated gradient [7] and SmoothGrad [8] accumulate the gradients with respect to modified input images in order to get a more interpretable sensitivity map than a single gradient map. In [9–13], the relevance map for an intermediate layer is back-propagated layer-by-layer from the output layer to the input layer with a modified back-propagation rule.

However, it is pointed out that the explanation generated by gradient-based methods are not necessarily faithful to the classifier’s decision process. Adebayo et al. [17] shows that some gradient-based methods are nearly independent of the classifier’s weight, and act like a model-ignorant edge detector rather than an explanation of the classifier. [18–20] adversarially attack gradient-based methods and can manipulate a saliency map without regard to the classifier’s output. From these results, we expect that the methods based on the input-output

relation are more faithful to the classifier’s actual behavior than gradient-based methods.

Other white-box methods include optimization-based methods and attention-based methods. In Meaningful perturbation [21], Extremal Perturbations [22], and FGVis [23], the image region to add perturbation (e.g. blurring, masking) are optimized by gradient descent with respect to an input image, and saliency information is extracted from the perturbed region. In [24], a visual question answering model with the attention mechanism is proposed and the attention maps can be interpreted as the relevant parts in an image.

LIME (Local Interpretable Model-agnostic Explanation) [1] and SHAP [2] are popular explanation-generating methods for a black-box image-classifier. Both perturb the given image preprocessed into a set of superpixels, observe how the output to the perturbed image, and generate an explanation based on the change in the output. Although their merits are widely appreciated, it is known that a generated explanation is not robust to how an image is decomposed into superpixels [25]. Our extension proposed in this paper is based on RISE [4] instead, which does not require a prior preprocessing of an input image. We will explain RISE in detail in §3.

Some black-box methods are based on these black-box techniques. IASSA [14] is an extension of RISE, which refines a saliency map by iteratively adapting the mask sampling process based on the previous saliency map and attention map. As extensions of the LIME framework, LORE [26] generates sample data with a genetic algorithm and fits decision trees instead of the linear regression model in LIME, and [27] incorporates the effect of higher-order interactions between input features.

Other black-box methods include Anchors [28], which searches for the minimal feature set which is sufficient for a correct prediction, and CXPlain [29], which trains the causal explanation model for the classifier’s behavior when some input feature is removed.

3 RISE

This section explains an explanation-generating method RISE [4], which is the basis of our method. For a detailed exposition, see Petsiuk et al. [4].

We first designate several definitions to define RISE. An *image* is a mapping from a finite set Λ of *pixels* to \mathbb{R}^3 . For an image i and a pixel λ , the tuple $i(\lambda) \in \mathbb{R}^3$ is the RGB value of λ in i . We write $\mathcal{I} = \{i \mid i : \Lambda \rightarrow \mathbb{R}^3\}$ for the set of images. We also designate the finite set of labels \mathcal{L} . We fix an image-classification model $M : \mathcal{I} \times \mathcal{L} \rightarrow [0, 1]$ throughout this paper; $M(i, l) \in [0, 1]$ is the model confidence in classifying i as l . For any i , we assume that $\sum_{l \in \mathcal{L}} M(i, l) = 1$.

A *mask* is an element of $\Lambda \rightarrow \{0, 1\}$. A mask m represents an image transformation that sets the RGB value of a pixel λ to $(0, 0, 0)$ if $m(\lambda) = 0$; the transformation keeps the original value of λ if $m(\lambda) = 1$. Therefore, the image $i \odot m$ that is obtained by applying the image transformation m to an image i is defined by $\lambda \mapsto i(\lambda) \times m(\lambda)$.

Given an image $i \in \mathcal{I}$ and a label $l \in \mathcal{L}$, we define the *saliency map* $S_{i,l}$ as $\lambda \mapsto \mathbb{E}_m[M(i \odot m, l) \mid m(\lambda) = 1]$, where the expectation is taken over all the possible masks. The intuition behind this definition is that λ can be considered to positively contribute to the classification of i to l if the model confidence of classifying the image to l remains high when λ is *not masked*.

Let X be a random variable whose value is a mask. Then, the definition of $S_{i,l}(\lambda)$ is equal to

$$\frac{1}{P[X(\lambda) = 1]} \sum_m M(i \odot m, l) \times m(\lambda) \times P[X = m] . \quad (1)$$

This expression shows that we can compute $S_{i,l}(\lambda)$ by computing the model confidence $M(i \odot m, l)$ for every mask m that does not mask λ and taking the weighted sum of these values.

Remark 1. In the actual implementation of RISE, the following optimizations are often applied.

- It is prohibitively expensive to precisely computing the value of the expression (1). Therefore, $S_{i,l}$ is approximated by the Monte-Carlo method. We randomly sample sufficiently many masks m_1, \dots, m_N and compute

$$S_{i,l}(\lambda) \approx \frac{1}{N} \sum_n \frac{m_n(\lambda)}{p} M(i \odot m_n, l) . \quad (2)$$

Here, p is $P[X(\lambda) = 1]$, m_n is the n -th mask, and N is the number of samples.

- RISE preprocesses a mask by bilinear interpolation before it applies the mask to an image. This is based on the intuition that (1) the saliency of a pixel is not binary and (2) a pixel that is close to an important pixel is often important. After this optimization, a mask is an element of $\Lambda \rightarrow [0, 1]$ instead of $\Lambda \rightarrow \{0, 1\}$.

Our extensions proposed in this paper also uses these optimizations.

4 Debiased RISE and MC-RISE

This section describes our extension to RISE. As we mentioned in §1, our extension consists of two enhancements of RISE: (1) MC-RISE, which generates *color-aware* saliency maps, and (2) *debiasing* saliency maps, which is a tweak to the RISE procedure so that the value of an irrelevant pixel in a saliency map is close to 0. We first present (2) in §4.1 and then (1) in §4.2.

4.1 Removing Bias from Saliency Maps

The range of the saliency values of a heatmap generated by RISE is highly variable depending on an input image. For example, the values in the saliency map at the first row in Fig. 4 (c) ranges from 0.016 to 0.023, whereas the one

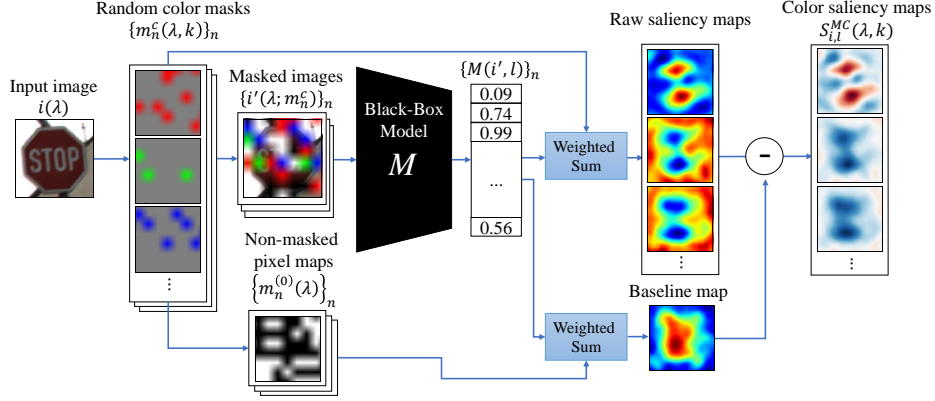


Fig. 2. Overview of our proposed method MC-RISE. An input image $i(\lambda)$ is randomly masked by given colors and the masked images $i'(\lambda; m_n^c)$ are fed to the black-box model M . Raw saliency maps are obtained by the weighted sum of the color masks $m_n^c(\lambda, k)$, where the weights are the output probabilities $M(i', l)$. Then a baseline map obtained from non-masked pixel maps $m_n^{(0)}(\lambda)$ is subtracted from them so that the saliency value in image parts irrelevant to classification is aligned to zero.

at the third row ranges from 0.375 to 0.550. This variability of saliency values degrades interpretability because it makes *thresholding* saliency value difficult: There is no clear threshold to decide which pixel is positively important, which is negatively important, and which is irrelevant.

To address this problem, we tweak the definition of RISE so that the saliency value of an irrelevant pixel to be close to 0, a positively-important pixel to be positive, and a negatively-important pixel to be negative; in other words, our method generates a *debiased* saliency map. Concretely, we generate a saliency map using the following definition instead of Equation (1):

$$S_{i,l}^{\text{PN}}(\lambda) := \mathbb{E}_m[M(i \odot m, l) | m(\lambda) = 1] - \mathbb{E}_m[M(i \odot m, l) | m(\lambda) = 0]. \quad (3)$$

Intuition behind Equation (3) is that we also need to use *negative* saliency $\mathbb{E}_m[M(i \odot m, l) | m(\lambda) = 0]$ of a pixel λ for the classification result l in computing a saliency map, not only using the positive saliency $\mathbb{E}_m[M(i \odot m, l) | m(\lambda) = 1]$ as RISE does. The expression $\mathbb{E}_m[M(i \odot m, l) | m(\lambda) = 0]$ indeed expresses the negative saliency of λ because its value becomes high if *masking out* the pixel λ increases the model confidence on average over the masks. To give the saliency as a single number, we calculate the difference of the positive saliency and the negative saliency. As a result, the sign of $S_{i,l}^{\text{PN}}(\lambda)$ indicates whether positive or negative saliency is dominant at a given pixel λ .

Calculating the difference of the positive saliency and the negative saliency also makes the saliency value of an irrelevant pixel to be 0. We can prove that a saliency map generated from Equation (3) is equipped with the following property (for the proof, see Section C in the supplementary material):

Proposition 1 (Irrelevant pixels). *If, for all binary mask samples, the output probability of class l from a black-box classifier M does not vary no matter whether a pixel λ in the mask m is retained or masked (i.e. $M(i \odot m, l)_{m(\lambda)=1} = M(i \odot m, l)_{m(\lambda)=0}$), then $S_{i,l}^{\text{PN}}(\lambda) = 0$.*

This property justifies the following interpretation of $S_{i,l}^{\text{PN}}$: A pixel whose value of $S_{i,l}^{\text{PN}}$ is close to 0 has almost no effect on the classification result. It is worth noting that (1) the sign of the value of $S_{i,l}^{\text{PN}}(\lambda)$ carries information on whether the pixel λ positively/negatively significant for the classification of i to l and (2) its absolute value expresses how significant λ is for the classification. In the original RISE, we can observe only that λ is more/less positively significant than another pixel λ' in the image i if the value of $S_{i,l}(\lambda)$ is larger/smaller than $S_{i,l}(\lambda')$. In §5.2, we empirically justify the above interpretation and argue that these properties render the interpretation of our saliency map easier than the one of the original RISE.

Equation (3) can be rewritten to the following equation that computes $S_{i,l}^{\text{PN}}$ by a weighted sum over all masks m (see Section A in the supplementary material for the detailed derivation):

$$S_{i,l}^{\text{PN}}(\lambda) = \sum_m \frac{m(\lambda) - p}{p(1-p)} M(i \odot m, l) P[X = m] , \quad (4)$$

where $p = P[X(\lambda) = 1]$. We approximate $S_{i,l}^{\text{PN}}$ by Monte Carlo sampling and express it by the following equation:

$$S_{i,l}^{\text{PN}}(\lambda) \approx \frac{1}{N} \sum_n \frac{m_n(\lambda) - p}{p(1-p)} M(i \odot m_n, l) , \quad (5)$$

where N is the number of mask samples and $\{m_n\}_{n=1}^N$ are randomly-sampled masks. We remark that Equation (5) is obtained from Equation (2) simply by replacing $\frac{m_n(\lambda)}{p}$ with $\frac{m_n(\lambda) - p}{p(1-p)}$; therefore, we can easily implement our debiasing method by modifying the implementation of RISE.

4.2 Multi-Colored RISE (MC-RISE)

We next present MC-RISE, our extension to RISE that explains the color-wise saliency of each pixel. The main idea of MC-RISE is to use a colored mask instead of a binary mask that RISE uses. Then, we can compute the positive/negative saliency of the given color component of a pixel from the response of the model to the masked image. The pseudocode for MC-RISE is presented in Section F in the supplementary material.

Color Masks We explain the method to generate color masks and color-masked images in MC-RISE. We fix a set of colors $c_1, \dots, c_K \in \mathbb{R}^3$ where each c_i is a 3-tuple of color values in the RGB colorspace. We are to generate a saliency map

for each c_i . To this end, MC-RISE generates *colored masks* instead of the binary masks that RISE uses. The saliency map is computed from the change in the average model confidence for the color-masked image in a similar way to RISE. We will explain the details of the saliency-map computation later.

To generate a colored mask, MC-RISE first generates a low-resolution color mask $m_{\text{low}}(\lambda, k)$ of the size $h \times w$ (which is smaller than the input image's size), where λ is the position of a pixel and k is the index for masking color in the color set $\{c_k\}_k$. The low-resolution color mask $m_{\text{low}}(\lambda, k)$ is generated as follows:

1. For each pixel, MC-RISE randomly decides whether it is masked or not with a masking probability p_{mask} .
2. For each pixel that is decided to be masked, MC-RISE chooses the color used to mask it from the color set $\{c_k\}_k$ with the uniform probabilities.
3. Then, the color mask $m_{\text{low}}(\lambda, k)$ is defined as follows:

$$m_{\text{low}}(\lambda, k) = \begin{cases} 1 & \text{if } \lambda \text{ is masked with color } k. \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

Then, $m_{\text{low}}(\lambda, k)$ is converted to a high-resolution color mask $m^c(\lambda, k)$ by (1) resizing $m_{\text{low}}(\lambda, k)$ to the size $H \times W$ of input images using bilinear interpolation and (2) shifting the resized mask by a random number of pixels from $(0, 0)$ up to $(\lfloor H/h \rfloor, \lfloor W/w \rfloor)$ (i.e., the size of a low-resolution pixel). We remark that the computation of $m^c(\lambda, k)$ corresponds to the preprocessing of a mask in RISE mentioned in Remark 1.

We also compute a non-masked pixel map $m^{(0)}(\lambda)$ by

$$m^{(0)}(\lambda; m^c) = 1 - \sum_{k=1}^K m^c(\lambda, k). \quad (7)$$

$m^{(0)}(\lambda; m^c)$ is 1 for non-masked pixels and 0 for masked pixels.

From a color mask m^c and an input image i , the color-masked image $i'(\lambda; m^c)$ is computed by $i'(\lambda; m^c) = i(\lambda)m^{(0)}(\lambda; m^c) + \sum_{k=1}^K c_k m^c(\lambda, k)$. Because of Equations (6) and (7), this expression can be seen as the alpha blending of an input image and the images uniformly filled with the masking colors $\{c_k\}_k$.

In Fig. 2, an example of a color mask $m^c(\lambda, k)$, a non-masked pixel map $m^{(0)}(\lambda)$, and a color-masked image $i'(\lambda; m^c)$ are presented.

Color-aware Saliency Map The saliency map of MC-RISE is defined by

$$S_{i,l}^{\text{MC}}(\lambda, k) := \mathbb{E}_{m^c \sim \mathcal{M}_c} [M(i'(\lambda; m^c)) | m^c(\lambda, k) = 1] - \mathbb{E}_{m^c \sim \mathcal{M}_c} [M(i'(\lambda; m^c)) | m^{(0)}(\lambda) = 1], \quad (8)$$

where \mathcal{M}_c denotes the probability distribution of color masks induced by the above mask-generation procedure. $S_{i,l}^{\text{MC}}(\lambda, k)$ represents how sensitively the model's confidence in classifying an image i to class l responds when the pixel λ in an input image is masked by the color c_k .

The debiasing method of a saliency map discussed in Section 4.1 is incorporated also into the definition of color saliency maps (Equation (8)), where the baseline saliency map for non-masked pixels ($\mathbb{E}_{m^c \sim \mathcal{M}_c}[M(i'(\lambda; m^c)) | m^{(0)}(\lambda) = 1]$, corresponding to the effect of retaining the pixel λ) is subtracted from the raw saliency maps ($\mathbb{E}_{m^c \sim \mathcal{M}_c}[M(i'(\lambda; m^c)) | m^c(\lambda, k) = 1]$, corresponding to the effect of masking the pixel λ).

The formula to estimate a color saliency map with Monte Carlo sampling can be derived in a similar way to §4.1 (see Section B in the supplementary material for the detailed derivation);

$$\begin{aligned} S_{i,l}^{\text{MC}}(\lambda, k) &= \mathbb{E}_{m^c \sim \mathcal{M}_c} \left[\left(\frac{m^c(\lambda, k)}{p_{\text{mask}}/K} - \frac{m^{(0)}(\lambda)}{1 - p_{\text{mask}}} \right) M(i'(\lambda; m^c)) \right] \\ &\approx \frac{1}{N} \sum_n \left(\frac{K m_n^c(\lambda, k)}{p_{\text{mask}}} - \frac{m_n^{(0)}(\lambda)}{1 - p_{\text{mask}}} \right) M(i'(\lambda; m_n^c)), \end{aligned} \quad (9)$$

where N is the number of mask sampling and $\{m_n^c\}_{n=1}^N$ are color masks sampled from the distribution \mathcal{M}_c . In our implementation, $\frac{K m_n^c(\lambda, k)}{p_{\text{mask}}}$ and $\frac{m_n^{(0)}(\lambda)}{1 - p_{\text{mask}}}$ in Equation (9) are separately accumulated using the same samples of color masks $\{m_n^c\}_n$; $S_{i,l}^{\text{MC}}(\lambda, k)$ is computed by subtracting the weighted-sum of the second term (baseline map) from that of the first term (raw saliency maps), as illustrated in Fig. 2.

Interpretation of Color Saliency Maps Since the debiasing method of §4.1 is incorporated into the definition of color saliency maps (Equation (8)), we can interpret the sign of a color saliency map as a positive/negative effect on the decision given by the model as we discussed in §4.1. However, because the sign of a color saliency map represents the expected change in the model confidence by *color-masking* a pixel λ , the interpretation of the sign of a saliency value is different from the debiased RISE in §4.1, in which a saliency value represents the effect of *retaining* original pixels. Therefore, we can interpret the color saliency map $S_{i,l}^{\text{MC}}(\lambda, k)$ as follows:

- If $S_{i,l}^{\text{MC}}(\lambda, k)$ has a positive or negative value *for all* k , the model’s output probability of class l should increase or decrease by any kind of color mask at a pixel λ , respectively. This suggests that the original *texture* at pixel λ is an obstacle or an important feature to the model, respectively, and its overall color is not relevant to the model’s decision.
- If $S_{i,l}^{\text{MC}}(\lambda, k) \sim 0$ *for all* k , any kind of masking at the pixel λ should not affect the model’s output. It means that the pixel λ is irrelevant to the model’s decision.
- If $S_{i,l}^{\text{MC}}(\lambda, k) > 0$ *for some* k , the model’s output probability of class l should increase by masking the pixel λ with the color c_k . This indicates that the presence of the color c_k at the pixel λ is an important feature for the model to classify an image as class l , but the original image lacks this feature.

- If $S_{i,l}^{\text{MC}}(\lambda, k) < 0$ for some k , the model’s output probability of class l should decrease by masking the pixel λ with the color c_k . This suggests that the pixel λ ’s original color is an influential feature for the model to classify an image as class l and masking it with the color c_k degrade the model’s confidence for class l .
- If $S_{i,l}^{\text{MC}}(\lambda, k) \sim 0$ for some k , color-masking the pixel λ with the color c_k should not affect the model’s output. It suggests that the color c_k is similar to the original color of the pixel λ from the viewpoint of the model, therefore masking λ with other colors would affect the model’s output.

5 Experiments

5.1 Experimental Settings

We use LIME [1] and RISE [4] as the baseline for the evaluation of our methods. We designated two kinds of models trained with different datasets for the evaluation.

- Models trained with GTSRB [15] dataset: We trained (1) VGG-16 [30] model with batch normalization and (2) ResNet-50 [31] model. We trained the models for 90 epochs with momentum SGD using cross-entropy loss. We set the learning rate to 10^{-1} , the momentum to 0.9, and the weight decay to 5×10^{-4} ; we decayed the learning rate by a factor of 0.1 in every 30 epochs. We resized images to $H \times W = 96 \times 96$ pixels for training and evaluation.
- Models trained with ImageNet[16] dataset: We used the pretrained ResNet-50 model provided by PyTorch [32] for the evaluation. We cropped and resized images to $H \times W = 224 \times 224$ pixels for the evaluation.

For MC-RISE and RISE, the number of samples N was set to 8000 and the masking probability was set to 0.5. For MC-RISE, we used the following five colors: red ($c_k = (255, 0, 0)$ in the RGB colorspace), green ($c_k = (0, 255, 0)$), blue ($c_k = (0, 0, 255)$), white ($c_k = (255, 255, 255)$), and black ($c_k = (0, 0, 0)$). For LIME, the parameters that we used were the same as the original paper [1].

We also applied MC-RISE to a person re-identification model with metric learning. See Section D in the supplementary material for details.

5.2 Results

Qualitative Evaluation of the Debiasing Method We qualitatively evaluated the debiasing method presented in §4.1. Fig. 3 compares the saliency maps of RISE and our debiased saliency map; we used ImageNet dataset to generate these maps. In the saliency map generated by the original RISE (Fig. 3(b)), the saliency values of the background regions—which should be irrelevant to the decision of the model—largely differs depending on the input images both in their absolute values and in their relative magnitude (i.e., pseudo-color in the plots). In our debiased maps (Fig. 3(c)), the background regions of all the images

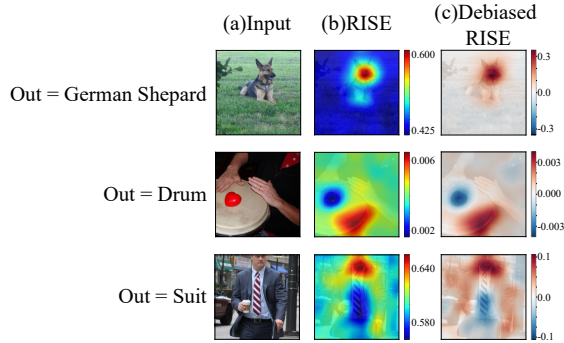


Fig. 3. Visual comparison between (b)raw saliency maps of RISE[4] and (c)debiased saliency maps calculated by our method. The parts of the image contributing positively/negatively to the output class have positive/negative intensity in (c), respectively. The background regions irrelevant to the model’s output have near-zero intensity in (c). Best viewed in color.

have the saliency values that are close to zero. These results demonstrate that the debiased RISE computes a saliency map in which the saliency values of the irrelevant regions are close to zero as we intend.

The saliency maps on the second and third rows in Fig. 3(c) shows that the saliency values near the objects of the output class (i.e., Drum and Suit) are positive and that the saliency values near the “noisy” objects (i.e., heart-shaped weight and necktie) are negative. These behaviors are also in accordance with the interpretation we presented in §4.1.

Qualitative Evaluation of MC-RISE Fig. 4 compares saliency maps generated by LIME [1], RISE [4], and our method MC-RISE; we used GTSRB dataset to generate these saliency maps. The saliency values of MC-RISE on the third row are not sensitive to the colors (i.e. the sign and the values of $S_{i,l}^{MC}(\lambda, k)$ is similar for all k). As we mentioned in §4.2, this suggests that the texture information is important for the model than the color information.

The saliency map of the image on the first row in Fig. 6 is sensitive to colors. In this image, the saliency near the white arrow is negative for all the colors; this suggests that the texture of this region (i.e., the arrow shape) is important for the model. In contrast, the values near the peripheral region of the sign in the green and the blue saliency maps are positive; this suggests that the blue-like color of this region is important for the model’s decision. Notice that LIME nor RISE can detect the sensitivity to colors; in fact, these parts identified as salient in the green and the blue saliency maps are not necessarily identifiable in the saliency maps generated by LIME and RISE.

Fig. 5 shows the comparison of the saliency maps in ImageNet dataset. Compared to maps in the GTSRB dataset, many of the maps in the ImageNet dataset have low color-dependency among the color saliency maps. This suggests that

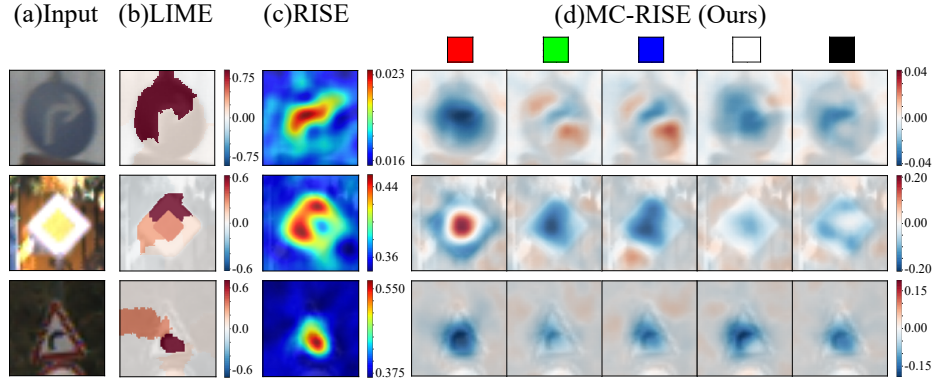


Fig. 4. Visual comparison among the saliency maps of (b) LIME [1], (c) RISE [4], and (d) the color saliency maps of MC-RISE in the GTSRB dataset. The boxes in the top row in (d) present the masking color used to obtain the maps in the same column. All samples are correctly classified by the model. Best viewed in color.

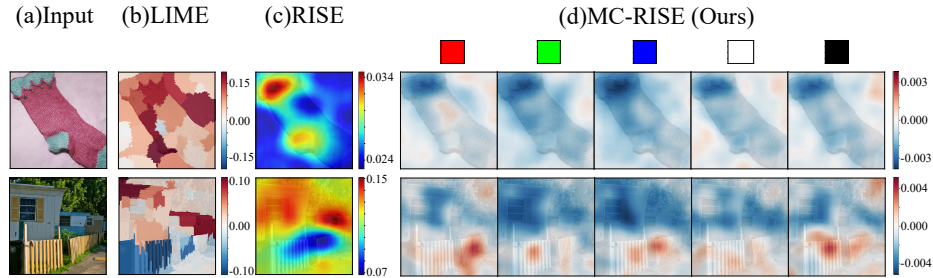


Fig. 5. Visual comparison among the saliency maps of (b) LIME [1], (c) RISE [4], and (d) the color saliency maps of MC-RISE in the ImageNet dataset. All samples are correctly classified by the model. Best viewed in color.

the model’s decision is more dependent on the detailed texture of objects than on the color of image regions. In such cases, we can extract information about positive or negative saliency from the color saliency maps. For example, in the third-row samples in Fig. 5, color saliency maps are interpreted to mean that the upper part of the image is important for the classification output (“Mobile home” class), but the fences in the lower part of the image are reducing the model’s confidence for the output class. From these observations, it is expected that MC-RISE can extract more useful information in domains where the color in image regions is highly important for classification, such as the GTSRB dataset, than domains where detailed texture is important such as the ImageNet dataset.

Quantitative Evaluation of MC-RISE For the quantitative evaluation of MC-RISE, we first would like to establish the evaluation metrics we use. We adapt the *deletion* metric proposed by Petsiuk et al. [4]—which measures how

Table 1. Evaluations of CA-deletion metric (lower is better) in LIME [1], RISE [4], and MC-RISE (ours).

Dataset	Model	LIME	RISE	MC-RISE
GTSRB	VGG-16	0.1324	0.0664	0.0270
	ResNet50	0.1294	0.0627	0.0204
ImageNet	ResNet50	0.1146	0.1046	0.0980

well a saliency-map-based explanation of an image classification result localizes the important pixels—to our color-aware setting.

To explain the deletion metric, let i be an image and M be a model. The deletion metric is the AUC value of the plot of f where $f(j)$ is the model confidence of the image in which the 1st to the j -th important pixels are removed. The lower the deletion metric is, the better the saliency map localizes the important region; a lower metric value means that the model confidence drops rapidly because the saliency map successfully points out the important pixels.

Our metric, called *CA-deletion (color-aware deletion)* is derived from the deletion metric as follows. In the CA-deletion metric, pixels in an input image are removed by masking them with *the most sensitive color* obtained from the color saliency maps. Concretely, pixel λ is removed in the ascending order of $\min_k [S_{i,l}^{\text{MC}}(\lambda, k)]$ and the removed pixel is filled by the color c_k where $k = \text{argmin}_k [S_{i,l}^{\text{MC}}(\lambda, k)]$. The model confidence for the pixel-removed images are plotted against the pixel-removal fractions; the CA-deletion value is the area-under-curve (AUC) of the plot. If a color saliency map correctly captures the model’s sensitivity to color masking, this removal method must be effective to reduce the model confidence, resulting in a low value of the CA-deletion metric.¹

Table 1 shows the evaluation result of MC-RISE in comparison with LIME and RISE. MC-RISE consistently outperforms RISE and LIME especially by a large margin in the GTSRB dataset. This indicates that MC-RISE can correctly capture the color-dependent sensitivity of the model, and effectively localize important pixels in an input image by using this information. The large margin in the GTSRB dataset is in accordance with our expectation that MC-RISE is especially effective in a domain, such as traffic sign classification, where color information conveys important information.

Use Case: Error Analysis Using MC-RISE We demonstrate a use case of MC-RISE for analyzing misclassified samples. Fig. 6 shows the color saliency maps generated by MC-RISE for a sample in GTSRB dataset that is misclassified by our model. The input is an image of a “Priority road” sign, whereas our model misclassified it as a “No vehicles” sign.

¹ Notice that other common metrics (e.g., pointing game) are not appropriate for evaluating MC-RISE because these metrics focus on the *positional* saliency; they do not consider the color saliency.




Fig. 6. Error analysis example in the GTSRB classification model with MC-RISE. Sensitivities for a red color mask at the center of the sign are opposite for the output class (No Vehicle) and the GT class (Priority road). It suggests that the absence of red color at the sign center is the cause of the error. Best viewed in color.

The saliency maps for the output “No vehicles” class show that the colors other than white at the center part of the sign have a negative effect in classifying this image as “No vehicles”. This suggests that, if the color of the center part were not white, then this image would not be classified as “No vehicles” class.

In contrast, in the saliency maps of the GT class (i.e., “Priority road” sign), the saliency map for red color has a strong positive value at the center of the sign, suggesting that the confidence of “Priority Road” class would largely increase if the color of the central part were red. From the above observations, we can hypothesize that the cause of the error is the absence of red color at the center of the sign (possibly due to overexposure). As one can observe, the information obtained by MC-RISE in addition to that obtained by the other methods, such as color-sensitivity or negative saliency, is indeed useful to explore the cause of the misclassification, especially when there is no objects which clearly hinder the correct decision of the model in an input image.

6 Conclusion

We proposed two extensions of a model-agnostic interpretation method RISE [4]. One is debiased RISE; it computes a saliency map in which the saliency of an unimportant pixel is close to 0; this makes the interpretation of a saliency map easier than RISE. Another is MC-RISE, which computes color-wise saliency; a saliency map for a color c computed by MC-RISE explains the saliency of the color c in each pixel. Such color-wise saliency is useful in application domains such as road sign recognition, in which colors convey significant information. We empirically demonstrated the effectiveness of our extensions.

We plan to extend our idea to use other features than colors; for example, by using masks that blur a part of an image, we expect that we can generate a saliency map that explains the significance of the clarity of the region.

Acknowledgment We thank the reviewers for their fruitful comments.

References

1. Ribeiro, M.T., Singh, S., Guestrin, C.: "why should I trust you?": Explaining the predictions of any classifier. In: SIGKDD 2016. (2016)
2. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions. In: NIPS 2017. (2017)
3. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: ICCV 2017. (2017)
4. Petsiuk, V., Das, A., Saenko, K.: RISE: randomized input sampling for explanation of black-box models. In: BMVC 2018. (2018)
5. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. In: 2017 IEEE International Conference on Computer Vision (ICCV). (2017) 618–626 ISSN: 2380-7504.
6. Chattopadhyay, A., Sarkar, A., Howlader, P., Balasubramanian, V.N.: Grad-CAM++: Generalized Gradient-Based Visual Explanations for Deep Convolutional Networks. In: 2018 IEEE Winter Conference on Applications of Computer Vision (WACV). (2018) 839–847
7. Sundararajan, M., Taly, A., Yan, Q.: Axiomatic attribution for deep networks. In: Proceedings of the 34th International Conference on Machine Learning - Volume 70. ICML'17, Sydney, NSW, Australia, JMLR.org (2017) 3319–3328
8. Smilkov, D., Thorat, N., Kim, B., Viégas, F., Wattenberg, M.: SmoothGrad: removing noise by adding noise. arXiv preprint (2017) arXiv:1706.03825 arXiv: 1706.03825.
9. Zeiler, M.D., Fergus, R.: Visualizing and Understanding Convolutional Networks. In Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T., eds.: Computer Vision – ECCV 2014. Lecture Notes in Computer Science, Cham, Springer International Publishing (2014) 818–833
10. Springenberg, J., Dosovitskiy, A., Brox, T., Riedmiller, M.: Striving for Simplicity: The All Convolutional Net. In: International Conference on Learning Representations. (2015)
11. Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.R., Samek, W.: On Pixel-Wise Explanations for Non-Linear Classifier Decisions by Layer-Wise Relevance Propagation. PLOS ONE **10** (2015) e0130140 Publisher: Public Library of Science.
12. Shrikumar, A., Greenside, P., Kundaje, A.: Learning Important Features Through Propagating Activation Differences. In Precup, D., Teh, Y.W., eds.: Proceedings of the 34th International Conference on Machine Learning. Volume 70 of Proceedings of Machine Learning Research., International Convention Centre, Sydney, Australia, PMLR (2017) 3145–3153
13. Zhang, J., Bargal, S.A., Lin, Z., Brandt, J., Shen, X., Sclaroff, S.: Top-Down Neural Attention by Excitation Backprop. Int J Comput Vis **126** (2018) 1084–1102
14. Vasu, B., Long, C.: Iterative and Adaptive Sampling with Spatial Attention for Black-Box Model Explanations. In: The IEEE Winter Conference on Applications of Computer Vision. (2020) 2960–2969
15. Stallkamp, J., Schlipsing, M., Salmen, J., Igel, C.: Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition. Neural Networks **32** (2012)

16. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. *Int J Comput Vis* **115** (2015) 211–252
17. Adebayo, J., Gilmer, J., Muelly, M., Goodfellow, I., Hardt, M., Kim, B.: Sanity Checks for Saliency Maps. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R., eds.: *Advances in Neural Information Processing Systems* 31. Curran Associates, Inc. (2018) 9505–9515
18. Heo, J., Joo, S., Moon, T.: Fooling Neural Network Interpretations via Adversarial Model Manipulation. In Wallach, H., Larochelle, H., Beygelzimer, A., Alché-Buc, F.d., Fox, E., Garnett, R., eds.: *Advances in Neural Information Processing Systems* 32. Curran Associates, Inc. (2019) 2925–2936
19. Dombrowski, A.K., Alber, M., Anders, C., Ackermann, M., Müller, K.R., Kessel, P.: Explanations can be manipulated and geometry is to blame. In Wallach, H., Larochelle, H., Beygelzimer, A., Alché-Buc, F.d., Fox, E., Garnett, R., eds.: *Advances in Neural Information Processing Systems* 32. Curran Associates, Inc. (2019) 13589–13600
20. Subramanya, A., Pillai, V., Pirsiavash, H.: Fooling Network Interpretation in Image Classification. In: *Proceedings of the IEEE International Conference on Computer Vision*. (2019) 2020–2029
21. Fong, R.C., Vedaldi, A.: Interpretable Explanations of Black Boxes by Meaningful Perturbation. In: *Proceedings of the IEEE International Conference on Computer Vision*. (2017) 3429–3437
22. Fong, R., Patrick, M., Vedaldi, A.: Understanding deep networks via extremal perturbations and smooth masks. In: *ICCV 2019*. (2019)
23. Wagner, J., Kohler, J.M., Gindele, T., Hetzel, L., Wiedemer, J.T., Behnke, S.: Interpretable and Fine-Grained Visual Explanations for Convolutional Neural Networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2019) 9097–9107
24. Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., Bengio, Y.: Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. In: *International Conference on Machine Learning*. (2015) 2048–2057 ISSN: 1938-7228 Section: Machine Learning.
25. Schallner, L., Rabold, J., Scholz, O., Schmid, U.: Effect of superpixel aggregation on explanations in LIME — A case study with biological data. *CoRR abs/1910.07856* (2019)
26. Guidotti, R., Monreale, A., Ruggieri, S., Pedreschi, D., Turini, F., Giannotti, F.: Local Rule-Based Explanations of Black Box Decision Systems. *arXiv preprint* (2018) arXiv:1805.10820 arXiv: 1805.10820.
27. Tsang, M., Cheng, D., Liu, H., Feng, X., Zhou, E., Liu, Y.: Feature Interaction Interpretability: A Case for Explaining Ad-Recommendation Systems via Neural Interaction Detection. In: *International Conference on Learning Representations*. (2019)
28. Ribeiro, M.T., Singh, S., Guestrin, C.: Anchors: High-Precision Model-Agnostic Explanations. In: *Thirty-Second AAAI Conference on Artificial Intelligence*. (2018)
29. Schwab, P., Karlen, W.: CXPlain: Causal Explanations for Model Interpretation under Uncertainty. In Wallach, H., Larochelle, H., Beygelzimer, A., Alché-Buc, F.d., Fox, E., Garnett, R., eds.: *Advances in Neural Information Processing Systems* 32. Curran Associates, Inc. (2019) 10220–10230
30. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014)

31. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR 2016. (2016)
32. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: PyTorch: An Imperative Style, High-Performance Deep Learning Library. In Wallach, H., Larochelle, H., Beygelzimer, A., Alché-Buc, F.d., Fox, E., Garnett, R., eds.: Advances in Neural Information Processing Systems 32. Curran Associates, Inc. (2019) 8026–8037