

## SIDU

### *Similarity Difference And Uniqueness Method for Explainable AI*

Muddamsetty, Satya Mahesh; Jahromi, Mohammad Naser Sabet; Moeslund, Thomas B.

*Published in:*  
2020 IEEE International Conference on Image Processing (ICIP)

*DOI (link to publication from Publisher):*  
[10.1109/ICIP40778.2020.9190952](https://doi.org/10.1109/ICIP40778.2020.9190952)

*Publication date:*  
2020

*Document Version*  
Accepted author manuscript, peer reviewed version

[Link to publication from Aalborg University](#)

*Citation for published version (APA):*  
Muddamsetty, S. M., Jahromi, M. N. S., & Moeslund, T. B. (2020). SIDU: Similarity Difference And Uniqueness Method for Explainable AI. In *2020 IEEE International Conference on Image Processing (ICIP)* Article 9190952 IEEE (Institute of Electrical and Electronics Engineers). <https://doi.org/10.1109/ICIP40778.2020.9190952>

#### **General rights**

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

#### **Take down policy**

If you believe that this document breaches copyright please contact us at [vbn@aub.aau.dk](mailto:vbn@aub.aau.dk) providing details, and we will remove access to the work immediately and investigate your claim.

# SIDU: SIMILARITY DIFFERENCE AND UNIQUENESS METHOD FOR EXPLAINABLE AI

Satya M. Muddamsetty

Mohammad N. S. Jahromi

Thomas B. Moeslund

Visual Analysis of People Laboratory (VAP), Aalborg University, Aalborg, Denmark

## ABSTRACT

A new brand of technical artificial intelligence ( Explainable AI ) research has focused on trying to open up the 'black box' and provide some explainability. This paper presents a novel visual explanation method for deep learning networks in the form of a saliency map that can effectively localize entire object regions. In contrast to the current state-of-the art methods, the proposed method shows quite promising visual explanations that can gain greater trust of human expert. Both quantitative and qualitative evaluations are carried out on both general and clinical data sets to confirm the effectiveness of the proposed method.

**Index Terms**— Explainable AI, Visual explanation, Clinical application, CNN

## 1. INTRODUCTION


Deep Learning has resulted in a significant performance breakthrough in a wide variety of areas in computer vision tasks such as object detection [1], image classification [2], image captioning [3] and many problems in machine learning (ML) [4, 5, 6, 7]. Despite all the popularity and high performance of deep learning, it is difficult to clearly understand or visualize the inner stacked layers of its architecture and ultimately, interpreting the output decision of such a network with millions of free parameters. Without principled understanding of how the 'black box' achieves its result, it is difficult to trust and deploy such AI models in, for example, domain like medical diagnosis or criminal justice where, the final decision may have serious consequences. Therefore, Explainable AI (XAI) is a newly emerging discipline of AI that attempts to shed light on the 'black box' by providing visual explanation or analysis of feature representations, hidden inside deep learning models. In this way, a particular deep learning model can be further evaluated by a human user/expert to establish trust in final predictions or help fix any classification errors. For example, in spite of a justifiably high skepticism rate within the medical community in supporting clinical decisions made by powerful machine learning models, XAI provides clinicians a useful tool by which to better audit the model's predictions and reason about external factors influencing prediction such as bias in the data [8].

Although there have been a number of early studies focusing on generating explanation schemes for deep network

models, considering the complexity of such a challenging task there is still much more effort needed to establish both reliable quantitative and qualitative methods in this new field. The majority of the proposed methods are based on generating visual feature explanations known as **saliency maps**. For example, in *Gradient-based algorithms* such as Grad-Cam [9], DeCovNet [10] or LRP [11], such a map can be obtained by backtracking the network's activations from the output back to the input via backpropagation [12], in order to highlight the input regions that are most important in realizing the final prediction. Although the visual explanation of this class results in a well-detailed activation heat-map, computing the gradient for certain architecture models is not very straightforward [13].

Alternatively, more recent approaches, like the state-of-the-art method of RISE [14], attempts to find the effect of selectively inserting or deleting parts of the input (*Perturbation-based*) in the model's output prediction. Despite more accurate saliency maps and higher classification scores (after a selective deletion process of the input) of Perturbation-based methods compared to those of the gradient-based methods, it is not yet possible to visualize all the perturbations and determine which one characterizes the desired explanation best. Moreover, for both the gradient based and perturbation explanation methods, the generated visual explanation suffers from localizing the entire salient regions of an object required for higher classification scores. An example of this drawback is illustrated in Fig. 1 where the estimated saliency maps were unable to localize entire object class properly. With more general interpretation approaches (*Approximation-based*), as in the case of LIME [15] that creates its saliency map based on random superpixels, the problem becomes even more deteriorated. This issue plays a significant role in certain classification tasks such as those found in medical domain where generating highly accurate visual explanations are equivalent to highlighting complete regions of interest.

In this paper, motivated by studies in both GRAD-CAM and RISE, we propose a new visual explanation approach for estimating pixel saliency by extracting the last convolutional layer of the deep CNN model and creating the similarity difference mask which is eventually combined to form a final map for generating the visual explanation for the prediction. We refer to the proposed method as *similarity difference and uniqueness method (SIDU)*. The SIDU method



**Fig. 1:** An example of failure of saliency maps to capture entire object class in (b):GRAD-CAM and (c): RISE.

is gradient-free (as opposed to GRAD-CAM) and unlike the random mask mechanism of the RISE algorithm, the final combined mask of our proposed methods comes from the last activation masks of the CNN model. The algorithm provides much better localization of the object class in question (see, for example, Fig. 1 (d)). This results in gaining greater trust of human expert level to rely on the deep model. Quantitative and qualitative tests on both general and clinical datasets further demonstrate the effectiveness of our proposed approach compared to the state-of-the-art RISE method.

$F$  with last convolution layers of size  $n \times n \times N$  where ' $n$ ' is size of that layer and ' $N$ ' is the number of features in the activation maps  $\mathbf{f}$  of class  $c$ , i.e.,  $\mathbf{f}^c = [f_1^c, \dots, f_N^c]$ . Each feature activation map  $f_i^c$  is then converted into a binary mask  $M_i^c$  corresponding to the feature activation map  $f_i^c$  in the convolution layer. Next, a bi-linear interpolation is applied to up-sample the binary mask for a given input image  $I$  of size  $W \times H$ . After interpolation, the binary mask  $M_i^c$  will be no longer binary and the values range between  $[0, 1]$ . Point-wise multiplication is performed between interpolated binary mask  $M_i^c$  and input image  $I$  and is represented as

$$A_i^c = F(I \odot M_i^c) \quad (1)$$

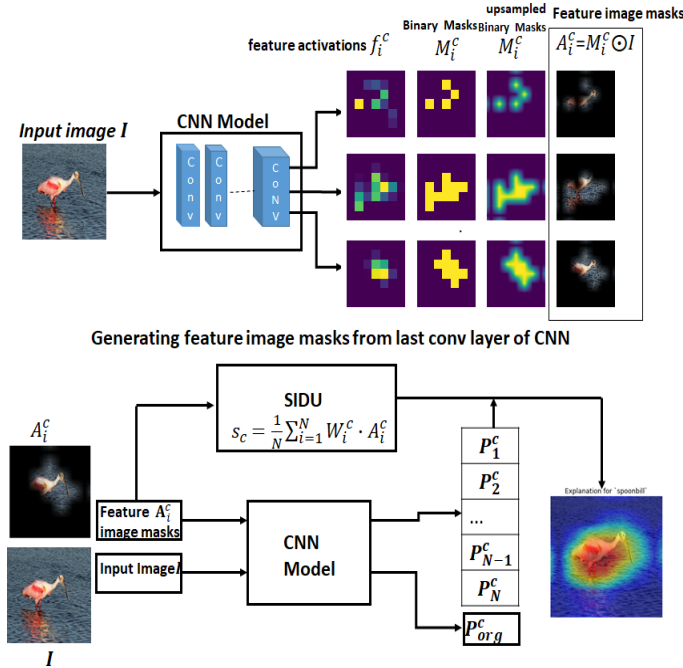
where  $F$  is an CNN model,  $A_i^c$  is the feature activation image mask of feature map  $f_i^c$  and  $i = 1, \dots, N$ .

## 2.2. Computing Similarity Differences and Uniqueness

In the next step, we compute probability prediction scores for all the feature activation image masks  $\mathbf{A}^c$  of class  $c$ , i.e.,  $\mathbf{A}^c = [A_1^c, \dots, A_N^c]$  using the CNN model  $F$ , respectively. Let the probability prediction score of the feature activation image mask  $A_i^c$  be  $P_i^c$  and the probability prediction score for the original image  $I$  be  $P_{org}^c$ . The similarity differences are then computed between each input feature activation image mask prediction score  $P_i^c$  and prediction score  $P_{org}^c$  of the original image  $I$ . The main idea is that the relevance of a feature map is estimated by measuring how the prediction changes if the feature is unknown, i.e., the similarity difference between prediction scores. The relevance value of the feature activation image mask increases if it is similar to the predicted class and decreases otherwise. The similarity difference of set of feature activation maps  $SD_i^c$  is given by

$$SD_i^c = \exp\left(\frac{-1}{2\sigma^2} \|P_{org}^c - P_i^c\|\right) \quad (2)$$

We also compute a uniqueness measure  $U$  which implements the commonly employed assumption that image regions which stand out from other regions in certain aspects catch our attention and hence should be labeled more salient. We therefore evaluate how different each respective feature mask is from all other feature masks constituting an image. The intuition behind this is to suppress the false regions with low weights and highlight the actual regions which are



**Fig. 2:** Block diagram of the proposed explanation method.

## 2. PROPOSED METHOD

Overview of the proposed explanation method is illustrated in Fig. 2 In the following subsections, we describe each step.

### 2.1. Generating Masks

We first generate masks from the last convolution layers of the deep CNN model. Let us consider any deep CNN model

responsible for predictions with higher weights. The uniqueness measure  $U_i^c$  is given in Eq.3.

$$U_i^c = \sum_{j=1}^N \|P_i^c - P_j^c\|, \quad (3)$$

The final weight of feature importance  $W_i^c$  is the dot product of the similarity difference  $SD_i^c$  and uniqueness measure  $U_i^c$  and is given

$$W_i^c = SD_i^c \cdot U_i^c, \quad (4)$$

where  $SD_i^c, U_i^c$  are the similarity difference and uniqueness values for the feature activation image mask  $A_i^c$  of the object class  $c$ .

### 2.3. Explanations for the prediction

The final visual explanation map  $S_c$ , also known as the class discriminative localization map can be computed as a weighted sum of image masks  $A_i^c$ , where the weights are computed by Eq. 4. Thus the visual explanation of the predicted class  $c$  is given by

$$S_c = \frac{1}{N} \sum_i^N W_i^c \cdot A_i^c \quad (5)$$

In summary, to explain the decision of the predicted class  $c$  visually, we first extract the last convolution layer from the deep CNN model  $F$  which has  $N$  number of features activation maps of size  $n \times n$ . We then generate  $N$  binary masks and point wise multiplication is performed between each generated binary mask  $M_i$  and the input image  $I$ . The similarity difference  $SD$  between probability scores of the predicted class and each point-wise multiplied image mask  $A_i$  and uniqueness measure  $U$  between the image masks are computed. Weights  $W_i$  of each image mask  $A_i$  computed by the dot product of  $SD$  and  $U$ . Finally, the visual explanation  $S_c$  is a weighted sum of feature activation image masks  $A_i$  given in Eq. 5

## 3. EXPERIMENTS AND RESULTS

A good explanation must be consistent with the CNN model's prediction and visualize the results in a manner that is intuitive for a human. In order to evaluate the performance of the SIDU explanation method we choose two datasets with different characteristics. The ImageNet [16] dataset of Natural Images with 1000 classes. We use 2000 images randomly collected from the ImageNet validation dataset. The other is a Retinal Fundus Image Quality Assessment (RFIQA) dataset from the medical domain. The dataset consists of 9,945 images with two levels of quality, 'Good' and 'Bad'. The retinal images were collected from a large number of patients with retinal diseases [17].

There is no consensus about what interpretability is in machine learning. Nor it is clear how to measure it. Initial research attempts to formulate approaches for evaluation [18]

**Table 1:** Quantitative Comparison XAI methods for ResNet50 on ImageNet validation set.

METHODS	Insertion $\uparrow$	Deletion $\downarrow$
RISE [14]	0.63571	0.13505
SIDU	<b>0.65801</b>	<b>0.13424</b>

**Table 2:** Quantitative Comparison XAI methods on RFIQA dataset for our trained ResNet-50 model.

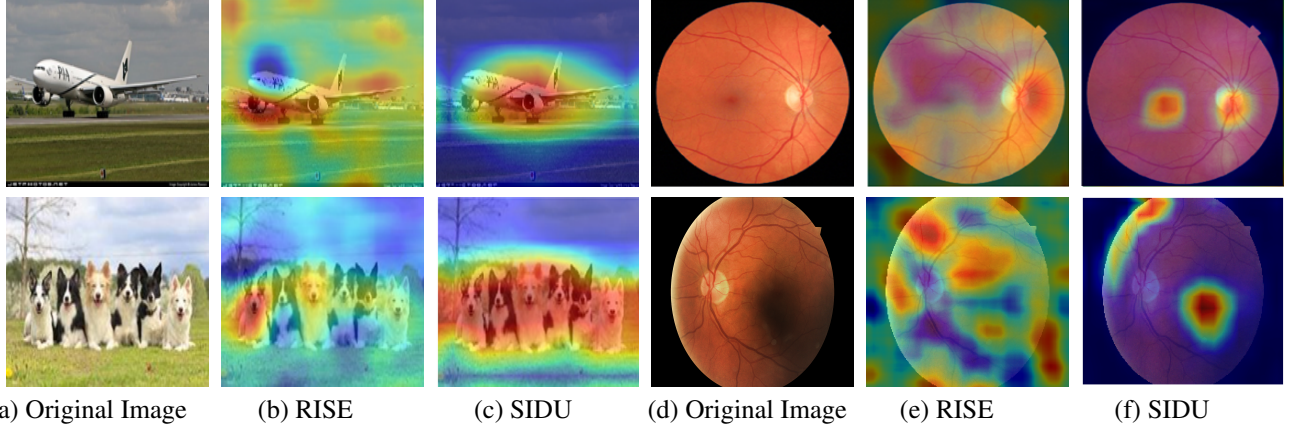
METHODS	Insertion $\uparrow$	Deletion $\downarrow$
RISE [14]	0.75231	0.59632
SIDU	<b>0.87883</b>	<b>0.47818</b>

which suggest faithfulness and human trust. We first evaluate the faithfulness of our model by studying the correlation between the visual explanation and prediction. Second, an expert level evaluation is performed, where human domain experts are involved in evaluating the human interpretability or human trust of the model. The experimental evaluation of faithfulness and human trust of the SIDU model is described in Section 3.1 and Section 3.2, respectively.

### 3.1. Evaluating Faithfulness

The faithfulness of the proposed method is evaluated using two automatic causal metrics *insertion* and *deletion* proposed by [14]. We choose these metrics to compare with the state-of-art methods. The *deletion metric* removes the saliency region responsible for predicting the object class in the image and forces the base model to change its decision. As more pixels are removed from the saliency region, this metric measures a decrease in the probability of the predicted score. Good explanation shows a sharp drop in the predicted score and thus a low area under the probability curve. On the other hand the *insertion metric* is a complementary approach. It measures the increase in probability of predicted score as more and more pixels are included with higher Area Under Curve (AUC) as an indication of a good explanation.

To perform the experimental evaluation of the proposed explanation method we conduct two experiments. In the first experiment we choose one of the existing standard CNN models, ResNet-50 [19] pre-trained on ImageNet dataset [16] and evaluate the faithfulness of the explainable model on the ImageNet validation dataset. Table 1 summarizes the results obtained by the proposed method and compares it to most recent work [14]. We can observe that the proposed method achieved better performance for both metrics, i.e., outperforming RISE [14]. This can be explained by the fact that the RISE method generates  $N$  number of random masks and the weights predicted for these masks give higher weights to false regions which makes the final map of RISE noisy. An example is shown in Fig. 3. In our proposed method, however, the generated masks come from the last feature activation maps



**Fig. 3:** Saliency maps generated for the natural images class of ImageNet <(a), (b), (c)> and Good / Bad quality eye fundus images <(d), (e), (f)> from RFIQA by RISE and the SIDU method with ResNet50 as the base network. In practice, the doctors verify the visibility of the optical disc and macular regions in a good quality image (4<sup>th</sup> image, 1<sup>st</sup> row) corresponding to the highlighted regions in the heatmap of the proposed method. Similarly, the bad quality (4<sup>th</sup> image, 2<sup>nd</sup> row) is due to the shadow just above the center of the image, i.e., exactly the region highlighted by the proposed method.

of the CNN model. Due to this the final explanation map will localize the entire region of interest (object class). The visual comparison of the explanations of the proposed method on ImageNet are shown in Fig. 3 (a), (b), (c).

To justify the above statement, we conducted a second experiment. We trained the existing ResNet-50 [19] with an additional two FC layers and softmax layer on the RFIQA dataset. The CNN model achieves 94% accuracy. The proposed explanation method uses the trained model for explaining the prediction of the RFIQA test subset with 1028 images. The evaluated results are summarized in Table 2. We can observe that the proposed method achieves higher AUC for *insertion* and lower AUC for *deletion* compared to RISE [14]. The visual explanations of the proposed and the RISE methods on the RFIQA test dataset are shown in Figure. 3 (d), (e), (f).

### 3.2. Qualitative Evaluation

Human expert level is an essential criterion for those end-users who have less trust in the results of prediction models (e.g. clinician). As in our medical diagnosis case, to show the effectiveness of the proposed method in terms of capturing the correct region with respect to the state-of-art method, we ask two ophthalmologist experts at the hospital to evaluate which visual representation invokes more trust and hence matches with actual examination results performed in the clinic. Here, the generated heat-maps in the RISE algorithm were treated as the baseline for the comparison.

Next, we follow the same setting as discussed in [9] and generate explanation heat-maps of 100 fundus images for two classes of ‘Good’ and ‘Bad’ quality using both the proposed method and the RISE algorithm. The exact nature of each algorithm in the test remains unknown to the ophthalmologists. Indeed, they are labeled as either ‘model I’ or ‘model II’ to the

test participants. Once the ophthalmologist determined which model better represents the regions of interest (good/bad quality regions) for each image, we next calculate the relative frequency of each outcome per total fundus image. Note that each participant had the option to select “both” models if they feel both the generated explanation maps were rather similar. In such cases we may have three different possibilities for each test image. In the case of the first ophthalmologist, the RISE explanation map was selected with the relative frequency of 0.02, the proposed algorithm with 0.84 and 0.14 being the same. For the second ophthalmologist, these numbers are 0.05, 0.93 and 0.02, respectively. This demonstrates that both ophthalmologists significantly favor the visual explanation generated by the proposed method over the RISE method. This can be clearly observed by visual examples of these explanation maps for the fundus image in Fig. 3. It is visually evident that the proposed algorithm is capable of properly localizing the region of interest and hence gaining greater trust by the expert.

## 4. CONCLUSION

In this paper we proposed a novel method called SIDU for explanation of black box models in a heat-map form via feature maps of the last convolution layers in the model. The proposed method is a gradient-independent method that can effectively localize entire object classes in an image. The quantitative and qualitative (human trust) experiments show that for both general and critical medical data, the proposed method outperforms state-of-the-art. The new explanation approach can provide further insight and helps in gaining greater trust in ML-based prediction results for the end-user in a sensitive-domain.

## 5. REFERENCES

- [1] R Girshick, J Donahue, T Darrell, and J Malik, “Rich feature hierarchies for accurate object detection and semantic segmentation,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014, pp. 580–587.
- [2] M. Oquab, L. Bottou, I. Laptev, and J. Sivic, “Is object localization for free?-weakly-supervised learning with convolutional neural networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 685–694.
- [3] H Fang, S Gupta, F Iandola, Rupesh K Srivastava, L Deng, Piotr Dollár, J Gao, X He, M Mitchell, J C Platt, et al., “From captions to visual concepts and back,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1473–1482.
- [4] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. Lawrence Z, and D. Parikh, “Vqa: Visual question answering,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 2425–2433.
- [5] A. Das, S. Datta, G. Gkioxari, S. Lee, D. Parikh, and D. Batra, “Embodied question answering,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018, pp. 2054–2063.
- [6] H. De Vries, F. Strub, S. Chandar, O. Pietquin, H. Larochelle, and A. Courville, “Guesswhat?! visual object discovery through multi-modal dialogue,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5503–5512.
- [7] Zachary C Lipton, “The mythos of model interpretability,” *Queue*, vol. 16, no. 3, pp. 31–57, 2018.
- [8] M. Wu, S. Parbhoo, M. C Hughes, V. Roth, and F. Doshi-Velez, “Optimizing for interpretability in deep neural networks with tree regularization,” *arXiv preprint arXiv:1908.05254*, 2019.
- [9] R. R. Selvaraju, A. Das, R. Vedantam, M. Cogswell, D. Parikh, and D. Batra, “Grad-cam: Why did you say that? visual explanations from deep networks via gradient-based localization,” *CoRR*, vol. abs/1610.02391, 2016.
- [10] S. Bach, A. Binder, G. Montavon, F. Klauschen, K. Müller, and W. Samek, “On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation,” *PloS one*, vol. 10, no. 7, 2015.
- [11] M. D. Zeiler and R. Fergus, “Visualizing and understanding convolutional networks,” in *European conference on computer vision*. Springer, 2014, pp. 818–833.
- [12] R. Fong, M. Patrick, and A. Vedaldi, “Understanding deep networks via extremal perturbations and smooth masks,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 2950–2958.
- [13] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, “Learning deep features for discriminative localization,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2921–2929.
- [14] V. Petsiuk, A. Das, and K. Saenko, “RISE: randomized input sampling for explanation of black-box models,” *CoRR*, vol. abs/1806.07421, 2018.
- [15] M. T. Ribeiro, S. Singh, and C. Guestrin, “‘’ 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*, 2016, pp. 1135–1144.
- [16] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, “ImageNet Large Scale Visual Recognition Challenge,” *International Journal of Computer Vision (IJCV)*, vol. 115, no. 3, pp. 211–252, 2015.
- [17] S. M Muddamsetty and T. B Moeslund, “Multi-level quality assessment of retinal fundus images using deep convolutional neural network,” *Submitted to IEEE-ICIP*, 2020.
- [18] F. Poursabzi-Sangdeh, D. G. Goldstein, Jake M. Hoffman, J. W. Vaughan, and H. Wallach, “Manipulating and measuring model interpretability,” *arXiv preprint arXiv:1802.07810*, 2018.
- [19] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” *CoRR*, vol. abs/1512.03385, 2015.