

Empowering CAM-Based Methods with Capability to Generate Fine-Grained and High-Faithfulness Explanations

Changqing Qiu¹, Fusheng Jin^{1*}, Yining Zhang²

¹Beijing Institute of Technology

²Peking University

changqing_qiu@bit.edu.cn, jfs21cn@bit.edu.cn, zhangyining@stu.pku.edu.cn

Abstract

Recently, the explanation of neural network models has garnered considerable research attention. In computer vision, CAM (Class Activation Map)-based methods and LRP (Layer-wise Relevance Propagation) method are two common explanation methods. However, since most CAM-based methods can only generate global weights, they can only generate coarse-grained explanations at a deep layer. LRP and its variants, on the other hand, can generate fine-grained explanations. But the faithfulness of the explanations is too low. To address these challenges, in this paper, we propose FG-CAM (Fine-Grained CAM), which extends CAM-based methods to enable generating fine-grained and high-faithfulness explanations. FG-CAM uses the relationship between two adjacent layers of feature maps with resolution differences to gradually increase the explanation resolution, while finding the contributing pixels and filtering out the pixels that do not contribute. Our method not only solves the shortcoming of CAM-based methods without changing their characteristics, but also generates fine-grained explanations that have higher faithfulness than LRP and its variants. We also present FG-CAM with denoising, which is a variant of FG-CAM and is able to generate less noisy explanations with almost no change in explanation faithfulness. Experimental results show that the performance of FG-CAM is almost unaffected by the explanation resolution. FG-CAM outperforms existing CAM-based methods significantly in both shallow and intermediate layers, and outperforms LRP and its variants significantly in the input layer. Our code is available at <https://github.com/dongmo-qc/FG-CAM>.

Introduction

In recent years, deep learning has been widely used in many areas. As a result, it is increasingly important to explain why the network makes a specific decision. There has been extensive research on the interpretability of deep learning, especially in some important areas such as medicine (2018; 2022; 2021; 2018; 2019) and autonomous driving (2021; 2021; 2022). In computer vision, a very large number of methods have been proposed, such as CAM-based methods (2018; 2020; 2021; 2021; 2021; 2021a; 2021b; 2022; 2019; 2020; 2017; 2020; 2021a; 2016; 2021b), LRP (2015)

*Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

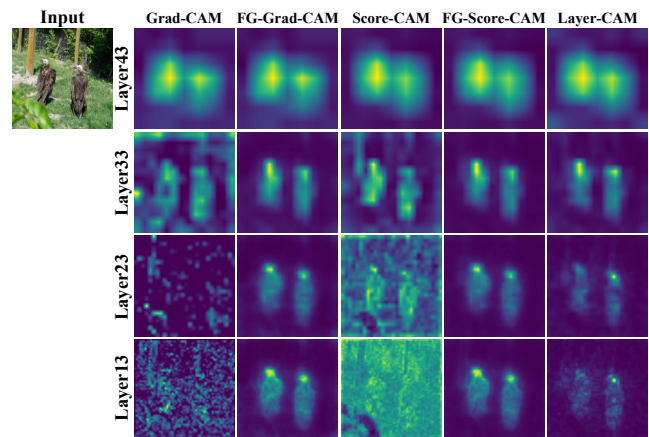


Figure 1: Visualization results of FG-CAM and other CAM-based methods on VGG-16 with batch normalization. Layer43 represents model.features[43]. Pixels with higher brightness are more important.

and its variants (2018; 2019), and other methods (2020; 2018; 2017; 2019; 2021; 2021).

CAM method (2016) generates an explanation through a linearly weighted combination of the feature maps of the last convolutional layer. Although CAM method can locate important regions in input image, it has great restrictions on the model structure. Therefore, as an extension of CAM, Grad-CAM (2017), Score-CAM (2020) and other CAM-based methods have been proposed. However, before Layer-CAM (2021), all CAM-based methods could only be applied to the final convolutional layer of the model. If the target layer became shallow, the faithfulness (measuring how accurately the explanation reveals the model decision) of the explanations would be significantly reduced. Thus, only coarse-grained explanations could be generated, providing only a rough location of the target object. This limitation restricted performance on tasks requiring pixel-accurate object locations (Jiang et al. 2021). Layer-CAM, on the other hand, although outperforming other CAM-based methods on shallow layers, suffers from shattered gradient problem (Balduzzi et al. 2017) as the model gets deep and the target layer becomes shallow, which can lead to problems such as mak-

ing explanation noisy and reduced faithfulness (not as dramatically as other CAM-based methods, but still not negligible). As shown in Figure 1, as the target layer becomes shallower, Grad-CAM and Score-CAM diminish the explanation accuracy significantly and even become unusable (unable to extract contributing pixels or filter out non-contributing pixels). The explanation of Layer-CAM has become more and more noisy and less accurate, while the FG-Grad-CAM and FG-Score-CAM still show good performance. LRP and its variants, which are based on the Deep Taylor Decomposition (Montavon et al. 2017), backpropagate the relevance score from output layer to input layer through specific rules to obtain the contribution of each pixel in the input image. They can generate explanations at the same resolution as the input image, but their faithfulness is low.

In order to empower CAM-based methods with capability to generate fine-grained and high-faithfulness explanations without changing their characteristics, this paper proposes a novel method: FG-CAM. Different from other methods which are directly applied to shallow layers to obtain high-resolution explanations, FG-CAM uses the relationship between two adjacent layers of feature maps with resolution differences to gradually increase the resolution of explanations produced on the last convolutional layer, while finding the contributing pixels and filtering out pixels that do not contribute. FG-CAM not only solves the problem that CAM-based methods can only provide rough explanations without changing their characteristics, but also generates fine-grained explanations that have higher faithfulness than LRP and its variants. This paper also presents FG-CAM with denoising, which is able to generate less noisy explanations with almost no change in explanation faithfulness.

To the best of our knowledge, our method achieves the SOTA in shallow and intermediate layers as well as in the input layer (different layers represent different explanation resolutions). The contributions of this paper are as follows.

1. This paper proposes FG-CAM, a novel explanation method, which can generate fine-grained explanations with high faithfulness in shallow and intermediate layers as well as in the input layer. FG-CAM solves the shortcoming of CAM-based methods without changing their characteristics, and achieves generality. Different from other CAM-based methods, the performance of FG-CAM is almost unaffected by the explanation resolution.
2. We also present FG-CAM with denoising, which is a variant of FG-CAM and can generate less noisy explanations with almost no change in explanation faithfulness.
3. We have verified qualitatively and quantitatively that FG-CAM performs significantly better than Layer-CAM (currently the best-performing CAM-based method for shallow layers), Grad-CAM, and Score-CAM in both shallow and intermediate layers. It is proved that whether FG-CAM is applied or not makes a very large difference.
4. Through experiments, we verify that FG-CAM can generate explanations with high faithfulness in the input layer, which is not possible for other CAM-based methods, and also outperforms LRP and its variants significantly.

Related Work

CAM-Based Methods

CAM-based methods generate an explanation through a linearly weighted combination of the feature maps of a certain layer (usually the last convolutional layer). Consider a layer l in a model f , A_l^i represents the feature map of the i th channel output by the layer, then the explanation of the CAM-based methods for class c can be defined as:

$$L_{CAM}^c = ReLU\left(\sum_i w_c^i A_l^i\right) \quad (1)$$

In original CAM, w_c^i is the connection weight between classification layer and global average pooling layer which is applied in the final convolution layer. It has significant limitations on the model structure.

Grad-CAM uses the average gradient of the output *w.r.t.* the feature map as the weight, solving the model structure dependency problem of CAM. Grad-CAM is equivalent to CAM when the target layer is followed by one and only one fully connected layer. Other gradient-based CAM methods, such as Grad-CAM++ (2018) and XGrad-CAM (2020), will not be discussed in this paper.

Different from gradient-based CAM methods, Score-CAM uses upsampled and normalized feature maps as masks of the image. The image perturbed by the masks is input into the model, and the weights of the feature maps are obtained according to the output. Score-CAM gets rid of the dependence on the gradient and shows very good performance. However, Score-CAM still can not be used at shallow layers. We use Grad-CAM and Score-CAM to verify that FG-CAM can be used on the CAM-based methods based on both score weight and gradient weight.

Unlike other CAM-based methods where one feature map corresponds to one global weight, Layer-CAM uses pixel-level weights and uses the positive gradient of the output *w.r.t.* the pixel as the weights (*i.e.* $w_c^i = ReLU\left(\frac{\partial y^c}{\partial A_l^i}\right)$) to address the problem that CAM-based methods can not be used at shallow layers. Layer-CAM uses fine-grained weights and can generate high-resolution explanations at shallow layers, which is not possible for the majority of CAM-based methods. However, as the network gets deeper, the gradients become noisy and discontinuous (Lee et al. 2021a).

LRP and Its Variants

LRP is another explanation method, which backpropagates the relevance score from the output layer to the input layer through specific rules to obtain the contribution of input image pixels to the model output.

The general propagation rule for LRP is the z^+ rule and z^β rule. For the output layer, LRP uses the model's output as relevance scores and sets the scores for non-target classes to 0. After propagating relevance score to the input layer, LRP determines the contribution of each pixel in the input image based on its relevance score. Therefore, LRP can generate explanations with the same resolution as input image. While LRP can generate fine-grained explanations, the results are not class-discriminative (Gu, Yang, and Tresp 2018). To address this problem, CLRP (2018) and SGLRP (2019) were

proposed, which differ from LRP only in their definition of the relevance score for the output layer. However, CLRP and SGLRP, which over-remove pixels that contribute to other classes, often result in important parts of the target class being removed (even removing the entire object). It can significantly affect the correctness of the explanation and may lead people in the wrong direction. Therefore, the explanations generated by LRP, CLRP, and SGLRP have low faithfulness.

Our Method

FG-CAM

CAM-based methods usually assign a global weight to a feature map (*i.e.* every pixel within the feature map has the same weight). However, in shallow layers, the contribution of each pixel in a feature map can be very different (*e.g.* contains pixels of both the target class and other classes simultaneously). The global weight cannot filter out these irrelevant pixels and highlight important pixels. Thus the generated explanation cannot focus attention on the important regions (*e.g.* Figure 1, explanations of Score-CAM and Grad-CAM on Layer13). Therefore, CAM-based methods can only be applied in the last convolutional layer to generate explanations that have very low resolution even though highlighting important regions (*e.g.* for VGG-16 with batch normalization, the explanation only has 7×7 resolution, which is much lower than 224×224 of the input image). And we need to be able to determine the importance of each pixel in order to generate a fine-grained explanation.

Layer-CAM proposes fine-grained weights so that it can be used directly at shallow layers to generate high-resolution explanations. However, the explanations generated by Layer-CAM are very noisy, and almost all CAM-based methods fail to produce fine-grained weights. Therefore, in order to achieve generality, we address the problem that CAM-based methods can only generate coarse-grained explanations from another perspective, which is to improve the explanation resolution, while finding the contributing pixels and filtering out the pixels that do not contribute (*i.e.* make the explanation fine-grained), rather than trying to use CAM-based methods directly on the shallow layers.

For the final explanation, it is difficult for us to make it fine-grained because we cannot easily obtain the relationship between the correct fine-grained explanation and the current explanation. In other words, we cannot accurately determine the importance of each pixel in the fine-grained explanation through the current explanation. However, for feature maps, this is very easy to implement. For two adjacent layers, the feature maps of previous layer usually have more details and higher resolution, and the feature maps of the later layer are calculated from the feature maps of the previous layer (*i.e.* we can easily obtain the relationship between the feature maps of adjacent layers). Therefore, if we have the importance of each pixel in the feature map of later layer, then we can determine the importance of the pixels in the previous layer’s feature map through the relationship between the two adjacent layer’s feature maps (*e.g.* based on computational relationships, we can allocate the importance of pixels from the later layer to pixels in the previous layer

using such as gradients, specific rules, and so on).

Through the above method, we obtain the feature maps with higher resolution and more detail, and filter out non-contributing pixels within them. By performing this layer by layer, we can obtain the importance of each pixel in the feature map with higher resolution of the target layer.

Therefore, we can first use CAM-based methods to obtain the importance of each pixel in the last convolutional layer (without combining them to get the final explanation), which can be viewed as the explanation components. Then we use the relationship between the two adjacent layers’ feature maps to improve the resolution of the explanation components. Finally, we sum the processed components to obtain fine-grained explanations. The pipeline of FG-CAM is shown in Figure 2.

Consider a model f , the last convolutional layer of f is L (the shape of output is $C \times W \times H$), the target layer is l (the shape of output is $c \times w \times h$). The FG-CAM explanation, written L_{FG-CAM} , can be calculated as follows.

Step 1: calculate the explanation components in L .

$$w^{(C \times 1 \times 1)} = CAM(A_L^{(C \times W \times H)}) \quad (2)$$

$$I_L^{(C \times W \times H)} = w^{(C \times 1 \times 1)} A_L^{(C \times W \times H)} \quad (3)$$

where A_L is the output of layer L , w denotes global weights of feature maps in layer L and can be obtained by any CAM-based methods. I_L denotes the explanation components in layer L .

Step 2: improve the resolution of I_L to obtain I_l .

$$I_l^{(c \times w \times h)} = IR(I_L^{(C \times W \times H)}) \quad (4)$$

where $IR()$ denotes improving the resolution of the explanation components layer by layer while finding the contributing pixels and filtering out the non-contributing pixels. In this paper, we implement IR based on the z^+ and z^β rule.

Consider a model, i represents the i th pixel in layer l , j represents the j th pixel in layer $l+1$, and $w_{ij}^{(l)}$ represents the weight of connection between i and j , and the importance of i and j is I_i^l and I_j^{l+1} respectively. Then, in this paper, the z^+ rule can be defined as:

$$I_{ij}^l = \frac{w_{ij}^{+(l)} x_i^l}{\sum_{i'} w_{i'j}^{+(l)} x_{i'}^l} I_j^{l+1} \quad (5)$$

$$I_i^l = \sum_j I_{ij}^l \quad (6)$$

where $w_{ij}^{+(l)} = \max(w_{ij}^{(l)}, 0)$, x_i^l is the pixel value of i in layer l . The z^β rule can be defined as:

$$I_{ij}^l = \frac{w_{ij}^{(l)} x_i^l - w_{ij}^{+(l)} b_i^l - w_{ij}^{-(l)} h_i^l}{\sum_{i'} w_{i'j}^{(l)} x_{i'}^l - w_{i'j}^{+(l)} b_{i'}^l - w_{i'j}^{-(l)} h_{i'}^l} I_j^{l+1} \quad (7)$$

where $w_{ij}^{-(l)} = \min(w_{ij}^{(l)}, 0)$, interval $[b_i^l, h_i^l]$ is the domain of the pixel value x_i^l .

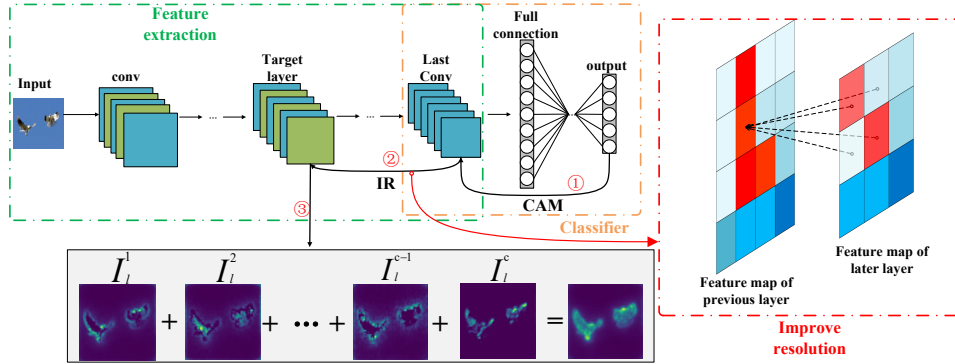


Figure 2: FG-CAM pipeline. In step 1, the explanation components are calculated using a CAM-based method in last convolutional layer. In step 2, improve resolution of the explanation components. Finally, in step 3, generate a fine-grained explanation.

Step 3: sum the I_i to obtain L_{FG-CAM} .

$$L_{FG-CAM} = ReLU\left(\sum_i^c I_i^i\right) \quad (8)$$

Therefore, different from trying to apply CAM-based method directly to shallow layers, FG-CAM solves the problem that CAM-based methods cannot generate fine-grained explanations while maintaining the good performance and characteristics of CAM-based methods, achieving good generality. Even if new CAM-based methods are proposed in the future, FG-CAM can still be used to generate fine-grained explanations based on these methods. In addition, it can be found that FG-CAM can generate explanations with the same resolution as the input image, which is impossible with CAM-based methods.

FG-CAM with Denoising

Due to the global weight, the explanations generated by the existing CAM-based methods are likely to be noisy. These noises will spread as the explanation resolution increases, resulting in many background pixels that have positive contributions (although their importance is very small compared to the target class pixels, they may also cause interference). To solve the problem, we propose FG-CAM with denoising, a variant of FG-CAM with one more step. After getting I_L , we first reshape I_L to $I_L^{(C \times K)}$, and for each row in matrix I , subtract the mean of that row. Secondly, we perform singular value decomposition on $I_L^{(C \times K)}$, and then select the top 10% singular values to reconstruct $I_L^{(C \times K)}$, finally reshape $I_L^{(C \times K)}$ to I_L . Then, we use the new I_L as the input to $IR()$. FG-CAM with denoising can generate explanations that focus more attention on the target class and have less noise.

Experiment

In this section, we compare FG-CAM with other methods through qualitative and quantitative experiments. FG-Score-CAM denotes FG-CAM that uses Score-CAM to obtain the global weights, and so for FG-Grad-CAM. We compare FG-Grad-CAM, FG-Score-CAM with Grad-CAM, Score-CAM,

and Layer-CAM to verify that FG-CAM does not significantly reduce the faithfulness of explanation as the target layer becomes shallower, as is the case with CAM-based methods, therefore solving the problem that CAM-based methods cannot generate fine-grained explanations, achieving SOTA. We also verify that FG-CAM has better performance than LRP, CLRP, and SGLRP in the input layer.

In the following experiments, we use pre-trained VGG-16 (2014) with batch normalization from PyTorch model zoo as a base model. ILSVRC2012 val (2015) is used as the data set. For each image, we resize it to $(224 \times 224 \times 3)$, convert it to range $[0, 1]$, then normalize it using mean vector $[0.485, 0.456, 0.406]$ and standard deviation vector $[0.229, 0.224, 0.225]$, and no further pre-processing beyond that. Layer13, Layer23, and Layer33 represent the 2nd, 3rd, and 4th max pooling layer, respectively. FG-CAM^D represents FG-CAM with denoising. For FG-CAM, LRP and its variants, we use the z^β rule in input layer and the z^+ rule in other layers. The blurred image is realized by Gaussian blur, that is, $\tilde{I} = gaussian_blur2d(I, ksize, sigma)$, where I is the original image, \tilde{I} is the blurred image. In this paper, we set $ksize = 51$, $sigma = 50$. The experimental results of FG-CAM have been bolded.

Comparison with CAM-Based Methods

In this section, we verify that FG-CAM can generate fine-grained explanations with high faithfulness and has better performance in both shallow and intermediate layers.

Class Discriminative Visualization. We compare the visual explanation of each method in shallow and intermediate layers. As shown in Figure 3, as the target layer becomes shallower, other methods diminish the explanation accuracy significantly, or even become unusable (can not filter out pixels that do not contribute), while FG-CAM still performs very well. At the same resolution, FG-CAM performs significantly better than other methods. The higher the explanation resolution, the bigger the gap between other methods and FG-CAM. Therefore, when generating high-resolution explanations, FG-CAM can still accurately extract the pixels with contribution and filter out non-contributing pixels. The

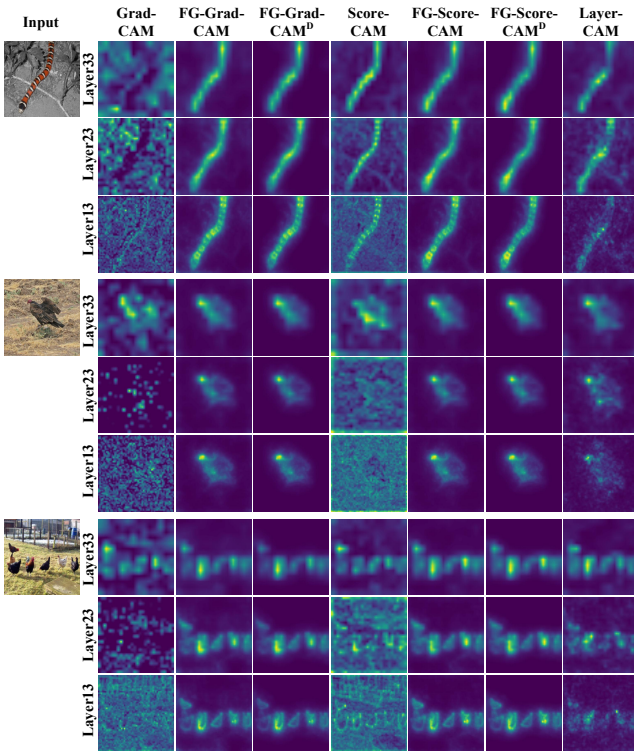


Figure 3: Qualitative comparisons of FG-CAM with other methods results. The results of FG-CAM are less noisy.

explanations generated by FG-CAM with denoising have the least noise.

Evaluation on Faithfulness. Faithfulness is an important characteristic to evaluate explanation methods, which measures whether the explanations correctly highlight regions related to the model’s decision. For the Average Drop/Increase metric, we retain only the most important 50% pixels of the input image, and then evaluate the faithfulness by Average Drop (A.D.) and Average Increase (A.I.) as Lee et al.; Ramaswamy et al.; Wang et al. The A.D. is expressed as $\sum_i^N \frac{\max(0, y_i^c - o_i^c)}{y_i^c} \times \frac{100}{N}$, and the A.I. is expressed as $\sum_i^N \frac{\text{sign}(y_i^c < o_i^c)}{N} \times 100$, where N is the number of images, y_i^c is the output softmax value for class c on image i , and o_i^c is the output softmax value for class c on the image after removing 50% of the pixels, $\text{sign}()$ indicates a function that returns 1 if the input is true and 0 otherwise. If the explanations correctly highlight important regions, Average Drop will have a low value and Average Increase will have a high value. We randomly selected 2000 images from the ILSVRC2012 val for evaluation as Wang et al.; Lee et al. To remove 50% of the pixels, we replace the least important 50% with blurred pixels. The experimental result is reported in Table 1.

As shown in Table 1, FG-CAM performs better than other methods in both shallow and intermediate layers. For example, in Layer13, FG-Grad-CAM achieves 19.65 A.D. and 32.35 A.I. respectively, which outperforms other methods by a large margin. Therefore, the fine-grained explanations

	Layer33		Layer23		Layer13	
Method	A.D.	A.I.	A.D.	A.I.	A.D.	A.I.
Grad-CAM	44.53	18.85	72.89	5.10	83.80	3.25
FG-Grad-CAM	18.69	34.90	19.55	33.45	19.65	32.35
FG-Grad-CAM ^D	18.00	33.85	17.75	33.55	17.90	34.00
Score-CAM	32.63	21.60	51.24	13.80	57.06	12.00
FG-Score-CAM	21.13	29.15	22.63	27.15	23.50	25.85
FG-Score-CAM ^D	19.19	31.70	18.99	32.00	19.24	30.50
Layer-CAM	24.25	26.80	27.32	24.50	30.76	24.40

Table 1: The results of the faithfulness evaluation (lower is better for A.D., higher is better for A.I.).

generated by FG-CAM can highlight important pixels in the image, and have higher faithfulness than explanations generated by other methods. From the experimental results, it can be found that as the target layer becomes progressively shallower, the results of Score-CAM and Grad-CAM become significantly worse, while the performance of Layer-CAM declined obviously, and there is a growing gap between other methods and FG-CAM. FG-CAM with denoising also performs better than FG-CAM in the evaluation.

To further verify the superiority of FG-CAM, we also conduct deletion and insertion tests proposed by Petsiuk, Das, and Saenko, and also adopted by Li et al.; Wang et al.; Zhang, Rao, and Yang. The deletion metric measures the decrease of the output softmax value for class c when important pixels are gradually removed, where the importance of pixels is given by explanation. If the explanation indicates the importance of pixels correctly, the output softmax value should drop sharply, and the area under the deletion curve (AUC) should be low. Similarly, the insertion metric measures the increase of the output softmax value when important pixels are gradually added, with higher AUC indicating a better explanation.

There are several methods of removing pixels (2017), such as setting the pixel value to zero or using blurred pixels instead, each with its own advantages and disadvantages. In this experiment, for the deletion metric, we replaced the original pixel with the blurred pixel, and for the insertion metric, we replaced the blurred pixel with the original pixel. We removed or added 448 pixels on each step. We use the *Over – all* score to comprehensively evaluate deletion and insertion results, as Zhang, Rao, and Yang. The *Over – all* can be calculated by $AUC(insertion) - AUC(deletion)$.

The average result over 2000 images is reported in Table 2. Similar to the results of the previous experiments, we find that as the target layer becomes progressively shallower, the *Over – all* scores of Score-CAM and Grad-CAM decrease significantly, while FG-CAM is barely affected and achieves the highest *Over – all* score. Whether our method is used or not makes a huge difference, e.g. in Layer13, FG-Grad-CAM and FG-Score-CAM have *Over – all* scores 0.42958 and 0.25348 higher than Grad-CAM and Score-CAM, respectively. FG-CAM with denoising does not have much impact on the final *Over – all* score.

Localization Evaluation. We also measure the quality

		Grad-CAM	FG-Grad-CAM	FG-Grad-CAM ^D	Score-CAM	FG-Score-CAM	FG-Score-CAM ^D	Layer-CAM
Layer33	Insertion	0.43920	0.60214	0.60247	0.51428	0.58202	0.59716	0.56566
	Deletion	0.15392	0.09640	0.11108	0.11988	0.10196	0.11198	0.09997
	Over-all	0.28528	0.50574	0.49139	0.39440	0.48006	0.48518	0.46569
Layer23	Insertion	0.30033	0.59217	0.60248	0.39734	0.57210	0.59676	0.54096
	Deletion	0.19113	0.09459	0.11378	0.14254	0.09985	0.11403	0.09333
	Over-all	0.10920	0.49758	0.48870	0.25480	0.47225	0.48273	0.44763
Layer13	Insertion	0.23266	0.58630	0.59976	0.36485	0.56570	0.59341	0.51348
	Deletion	0.16850	0.09256	0.11421	0.15037	0.09774	0.11435	0.08469
	Over-all	0.06416	0.49374	0.48555	0.21448	0.46796	0.47906	0.42879

Table 2: Results of deletion/insertion tests. The Over-all score (higher is better) shows that FG-CAM outperforms other methods.

	Grad-CAM	FG-Grad-CAM	FG-Grad-CAM ^D	Score-CAM	FG-Score-CAM	FG-Score-CAM ^D	Layer-CAM
Layer33	0.32417	0.51458	0.66736	0.35646	0.57161	0.68256	0.49272
Layer23	0.31645	0.50995	0.65744	0.28588	0.56192	0.67233	0.49854
Layer13	0.30760	0.50861	0.65472	0.29219	0.56001	0.66959	0.48238

Table 3: Comparative evaluation on localization ability (higher is better).

of explanation through localization ability. We did not use the pointing game proposed by Zhang et al. for this measure because it is too simple, requiring only the point with max value in the explanation to fall into the target object bounding box, which cannot distinguish each method significantly. Following Wang et al., we measure how many values in the explanation fall into the target object bounding box, that is, measure the ability of explanation to locate target object using the *Proportion* metric, where $Proportion = \frac{\sum L_{(i,j) \in bbox}^c}{\sum L_{(i,j) \in bbox}^c + \sum L_{(i,j) \notin bbox}^c}$. As Wang et al., we randomly select 500 images with only one bounding box and no more than 50% of the pixels occupied by the box for the experiment, and observe average *Proportion* value.

The result is reported in Table 3. Even though the evaluation measures only how much attention falls within the target object bounding box (it cannot accurately reflect whether the explanation fits the edge of objects), FG-CAM performs better than other methods. On average, over 50% of the values in the explanation generated by FG-CAM fall into the target object bounding box, indicating that the explanations generated by FG-CAM are less noisy and more focused on the target object. It can be found that FG-CAM with denoising performs significantly better than the original FG-CAM.

All experimental results show that there is a significant difference in whether our method is used or not. At the same resolution, FG-CAM outperforms other methods by a large margin. This fully verifies that FG-CAM solves the problem that CAM-based methods cannot generate fine-grained explanations while achieving generality, and has better results than Layer-CAM, which is designed for shallow layers. Moreover, it can be concluded from the experimental results that FG-CAM’s performance is almost unaffected by the explanation resolution. We also find that FG-Grad-CAM perform better than FG-Score-CAM on faithfulness evalua-

tion. We believe that this is because the global weights obtained by Grad-CAM are more accurate and using interpolation masks this advantage.

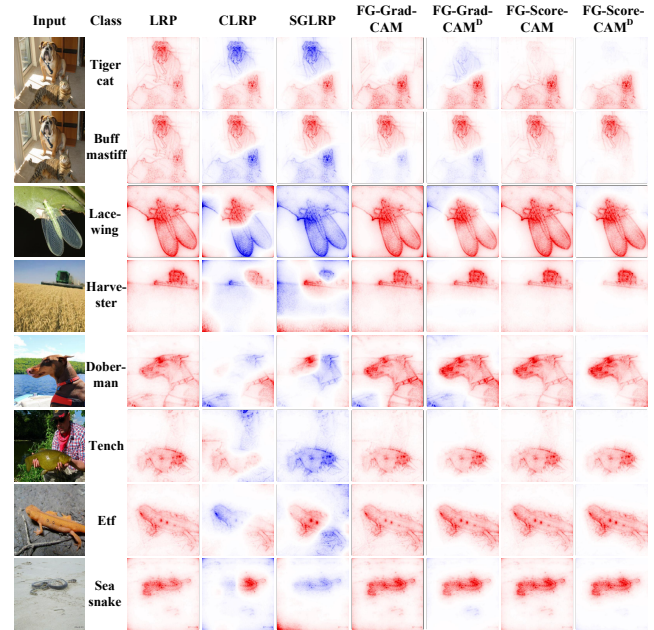


Figure 4: Visualization results of LRP, CLRP, SGLRP, and FG-CAM. Red indicates positive contribution pixels (the darker the color, the greater the contribution) and blue indicates negative contribution pixels.

Comparison with LRP and Its Variants

In this section, we compare LRP, SGLRP, CLRP with FG-CAM to verify that FG-CAM can generate explanations with

	LRP	CLRP	SGLRP	FG-Grad-CAM	FG-Grad-CAM ^D	FG-Score-CAM	FG-Score-CAM ^D
A.D.	25.48	30.64	40.62	16.81	16.13	20.28	16.84
A.I.	25.00	24.15	21.75	36.45	36.55	30.15	31.80

Table 4: The results of the faithfulness evaluation for LRP, CLRP, SGLRP and FG-CAM.

	LRP	CLRP	SGLRP	FG-Grad-CAM	FG-Grad-CAM ^D	FG-Score-CAM	FG-Score-CAM ^D
Insertion	0.55344	0.56740	0.51655	0.60757	0.61386	0.58608	0.60566
Deletion	0.09975	0.19168	0.22315	0.08668	0.11271	0.08980	0.11143
Over-all	0.45369	0.37572	0.29340	0.52089	0.50115	0.49628	0.49423

Table 5: The results of insertion and deletion tests for LRP, CLRP, SGLRP and FG-CAM.

	LRP	CLRP	SGLRP	FG-Grad-CAM	FG-Grad-CAM ^D	FG-Score-CAM	FG-Score-CAM ^D
Proportion	0.48385	0.59515	0.52136	0.49503	0.65496	0.54000	0.66983

Table 6: The results of evaluation on localization for LRP, CLRP, SGLRP and FG-CAM.

high faithfulness and same resolution as the input image in the input layer, and outperforms these methods significantly. We performed the same experiment. However, in the visualization and faithfulness evaluations, the explanations include negative contribution pixels to get more accurate results.

The visualization results are shown in Figure 4. It can be found that the explanations generated by LRP are not class-discriminative (*e.g.* on the first image, LRP even considers the pixels in the ‘dog’ region to be more important). CLRP and SGLRP often remove the important regions of the target class, and even regard pixels of important regions as negatively contributing. SGLRP, in particular, may regard all pixels of the target object as negatively contributing, which may make people mistakenly think that the model does not pay attention to the target object. FG-CAM correctly removed pixels of non-target classes without removing important parts, and focused attention on the target class. The explanations generated by FG-CAM with denoising focus more attention on the target class and remove a lot of noise.

For the Average Drop/Increase metric, if the proportion of positive-contributing pixels does not exceed 50%, then we will retain only the positive-contributing pixels. This is to ensure that only pixels with positive contribution are used. The result is reported in Table 4. The A.D. and A.I. values of CLRP and SGLRP are very poor, due to the fact that they often remove important regions of the target class. FG-CAM is significantly better than other methods in every metric, which indicates that FG-CAM still maintains high faithfulness of explanation in the input layer and can better highlight important pixels related to the model’s decision.

The results of the deletion and insertion tests are shown in Table 5. In this test, FG-CAM still outperforms other methods. For example, FG-Grad-CAM’s *Over – all* score is 0.0672 higher than LRP and 0.22749 higher than SGLRP., which is a huge lead. This indicates that the explanation with the same resolution as the input image generated by FG-CAM can fit the importance distribution of pixels well.

The localization evaluation results are shown in Table 6. It can be found that CLRP and SGLRP have higher *Proportion* value than LRP, but lower faithfulness in explanation than LRP, which further suggests that they only focus on some regions of target class and often remove important regions. The *Proportion* value of FG-CAM with denoising is significantly higher than that of others. FG-CAM is also better than LRP. This indicates that FG-CAM maintains high attention to the target class without over-removing pixels.

The experiments in this section show that FG-CAM can not only generate explanations with the same resolution as the input image, but also outperform LRP and its variants by a large margin. Therefore, we believe that FG-CAM greatly mitigates the shortcoming of CAM-based methods.

Conclusion

In this paper, we propose FG-CAM, a novel explanation method. FG-CAM uses the relationship between two adjacent layers of feature maps with resolution difference to gradually increase the explanation resolution, while finding the contributing pixels and filtering out the pixels that do not contribute. FG-CAM solves the problem that CAM-based methods cannot generate fine-grained explanations from another perspective. It not only maintains the characteristics of CAM-based methods, but also produces better results. Moreover, the performance of FG-CAM is almost unaffected by the explanation resolution. FG-CAM greatly outperforms CAM-based methods in both shallow and intermediate layers, and outperforms LRP and its variants in the input layer significantly. We also present FG-CAM with denoising, which is able to generate less noisy explanations with almost no change in explanation faithfulness. Our method provides a novel idea to solve the shortcoming of CAM methods, and achieves very good results. We believe that it will stimulate more research in this area.

Acknowledgments

This work was funded by National Natural Science Foundation of China (No. 62272045).

References

- Atakishiyev, S.; Salameh, M.; Yao, H.; and Goebel, R. 2021. Explainable Artificial Intelligence for Autonomous Driving: A Comprehensive Overview and Field Guide for Future Research Directions. *arXiv preprint arXiv:2112.11561*.
- 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.
- Balduzzi, D.; Frean, M.; Leary, L.; Lewis, J.; Ma, K. W.-D.; and McWilliams, B. 2017. The shattered gradients problem: If resnets are the answer, then what is the question? In *International Conference on Machine Learning*, 342–350. PMLR.
- Biffi, C.; Oktay, O.; Tarroni, G.; Bai, W.; Marvao, A. D.; Doumou, G.; Rajchl, M.; Bedair, R.; Prasad, S.; Cook, S.; et al. 2018. Learning interpretable anatomical features through deep generative models: Application to cardiac remodeling. In *International conference on medical image computing and computer-assisted intervention*, 464–471. Springer.
- 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.
- Chen, Y.; Zhang, J.; and Qin, X. 2022. Interpretable instance disease prediction based on causal feature selection and effect analysis. *BMC medical informatics and decision making*, 22(1): 1–14.
- Dabkowski, P.; and Gal, Y. 2017. Real time image saliency for black box classifiers. *Advances in neural information processing systems*, 30.
- 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*.
- Gao, W.; Wan, F.; Pan, X.; Peng, Z.; Tian, Q.; Han, Z.; Zhou, B.; and Ye, Q. 2021. Ts-cam: Token semantic coupled attention map for weakly supervised object localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2886–2895.
- Grabska-Barwińska, A. 2020. Measuring and improving the quality of visual explanations. *arXiv preprint arXiv:2003.08774*.
- 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.
- Ioannou, G.; Papagiannis, T.; Tagaris, T.; Alexandridis, G.; and Stafylopatis, A. 2021. Visual interpretability analysis of Deep CNNs using an Adaptive Threshold method on Diabetic Retinopathy images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 480–486.
- 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.
- 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.
- Lee, J. R.; Kim, S.; Park, I.; Eo, T.; and Hwang, D. 2021a. Relevance-cam: Your model already knows where to look. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14944–14953.
- Lee, K. H.; Park, C.; Oh, J.; and Kwak, N. 2021b. LFI-CAM: Learning feature importance for better visual explanation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1355–1363.
- Li, H.; Li, Z.; Ma, R.; and Wu, T. 2022. FD-CAM: Improving Faithfulness and Discriminability of Visual Explanation for CNNs. *arXiv preprint arXiv:2206.08792*.
- Molle, P. V.; Strooper, M. D.; Verbelen, T.; Vankeirsbilck, B.; Simoens, P.; and Dhoedt, B. 2018. Visualizing convolutional neural networks to improve decision support for skin lesion classification. In *Understanding and interpreting machine learning in medical image computing applications*, 115–123. 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.
- Omeiza, D.; Webb, H.; Jirotko, M.; and Kunze, L. 2021. Explanations in autonomous driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*.
- Patro, B. N.; Lunayach, M.; Patel, S.; and Nambodiri, V. P. 2019. U-cam: Visual explanation using uncertainty based class activation maps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7444–7453.
- Petsiuk, V.; Das, A.; and Saenko, K. 2018. Rise: Randomized input sampling for explanation of black-box models. *arXiv preprint arXiv:1806.07421*.
- Ramaswamy, H. G.; et al. 2020. Ablation-cam: Visual explanations for deep convolutional network via gradient-free localization. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 983–991.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3): 211–252.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, 618–626.

- Shrikumar, A.; Greenside, P.; and Kundaje, A. 2017. Learning important features through propagating activation differences. In *International conference on machine learning*, 3145–3153. PMLR.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Wagner, J.; Kohler, J. M.; Gindele, T.; Hetzel, L.; Wiedemer, J. T.; and Behnke, S. 2019. Interpretable and fine-grained visual explanations for convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9097–9107.
- 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.
- Xiang, A.; and Wang, F. 2019. Towards interpretable skin lesion classification with deep learning models. In *AMIA annual symposium proceedings*, volume 2019, 1246. American Medical Informatics Association.
- Yadu, A.; Suhas, P.; and Sinha, N. 2021. Class Specific Interpretability in CNN Using Causal Analysis. In *2021 IEEE International Conference on Image Processing (ICIP)*, 3702–3706. IEEE.
- Zablocki, É.; Ben-Younes, H.; Pérez, P.; and Cord, M. 2022. Explainability of Deep Vision-Based Autonomous Driving Systems: Review and Challenges. *International Journal of Computer Vision*, 1–28.
- 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. 2021a. Group-CAM: group score-weighted visual explanations for deep convolutional networks. *arXiv preprint arXiv:2103.13859*.
- Zhang, Q.; Rao, L.; and Yang, Y. 2021b. 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.
- Zhang, R.; Madumal, P.; Miller, T.; Ehinger, K. A.; and Rubinstein, B. I. 2021. Invertible concept-based explanations for cnn models with non-negative concept activation vectors. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 11682–11690.
- Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2921–2929.