

Towards Better Explanations of Class Activation Mapping

Hyungsik Jung Youngrock Oh
Samsung SDS

{hs89.jung, y52.oh}@samsung.com

Abstract

Increasing demands for understanding the internal behavior of convolutional neural networks (CNNs) have led to remarkable improvements in explanation methods. Particularly, several class activation mapping (CAM) based methods, which generate visual explanation maps by a linear combination of activation maps from CNNs, have been proposed. However, the majority of the methods lack a clear theoretical basis on how they assign the coefficients of the linear combination. In this paper, we revisit the intrinsic linearity of CAM with respect to the activation maps; we construct an explanation model of CNN as a linear function of binary variables that denote the existence of the corresponding activation maps. With this approach, the explanation model can be determined by additive feature attribution methods in an analytic manner. We then demonstrate the adequacy of SHAP values, which is a unique solution for the explanation model with a set of desirable properties, as the coefficients of CAM. Since the exact SHAP values are unattainable, we introduce an efficient approximation method, LIFT-CAM, based on DeepLIFT. Our proposed LIFT-CAM can estimate the SHAP values of the activation maps with high speed and accuracy. Furthermore, it greatly outperforms other previous CAM-based methods in both qualitative and quantitative aspects.

1. Introduction

Recently, convolutional neural networks (CNNs) have achieved excellent performance in various real-world vision tasks. However, it is difficult to explain their predictions due to a lack of understanding of their internal behavior. To grasp why a model makes a certain decision, numerous *saliency methods* have been proposed. The methods generate *visual explanation maps* that represent pixel-level importances for which regions in an input image are responsible for the model’s decision and which are not. Towards better comprehension of CNNs, *class activation mapping (CAM)* based methods, which utilize the responses of a convolutional layer for explanations, have been widely used.

CAM-based methods [3, 4, 7, 14, 18, 20] (abbreviated as CAMs in the remainder of this paper) *linearly* combine *activation maps* to produce visual explanation maps. Since the activation maps are fixed for a given input image and a model pair, the coefficients of a linear combination govern the performance of the methods. Therefore, it is critical to design a reasonable method of determining the coefficients. However, the majority of CAMs rely on heuristic conjectures for coefficient assignment without a clear theoretical basis. Specifically, the underlying linearity of CAM w.r.t. the activation maps is not fully taken into account. In addition, they do not set rigorous standards of which properties are expected to be satisfied in a good explanation model.

In this work, we leverage the linearity of CAM to analytically determine the coefficients beyond heuristics. Focusing on the fact that CAM defines an explanation map using a linear combination of activation maps, we formulate an explanation model as a linear function of the binary variables denoting the existence of the associated activation maps. Under this scheme, each activation map can be seen as an individual feature in *additive feature attribution methods* [2, 10, 12, 16]. Notably, *SHapley Additive exPlanations (SHAP)* [10] provides SHAP values as a unified measure of feature importance that satisfies three desirable properties (described in Sec. 2.2). Thus, the coefficients can be determined by the SHAP values of the corresponding activation maps. However, the exact SHAP values are not computable. To solve this, we propose a novel saliency method using *Deep Learning Important FeaTures (DeepLIFT)* [16], called *LIFT-CAM*, which efficiently approximates the SHAP values of the activation maps. Our contributions are summarized as follows:

- We propose a novel framework of determining a plausible visual explanation map of CAM, by reframing the problem as determining a reliable solution for the explanation model using additive feature attribution methods. The recent Ablation-CAM [4] can be reinterpreted by this framework.
- We formulate the SHAP values of the activation maps as a unified solution for the proposed framework and

verify their benefits in terms of generating faithful visual explanations.

- We introduce a new saliency method, LIFT-CAM, based on DeepLIFT. It effectively estimates the SHAP values of the activation maps with a single backward propagation and outperforms the other previous CAMs qualitatively and quantitatively.

2. Related Work

2.1. Class Activation Mapping

Visual explanation map. Let f be an original prediction model and c denote a target class of interest. CAMs [3, 4, 7, 14, 18, 20] aim to explain the target output of the model for a specific input image x (i.e., $f^c(x)$) through the visual explanation map, which can be generated by:

$$L_{\text{CAM}}^c(A) = \text{ReLU}\left(\sum_{k=1}^{N_l} \alpha_k A_k\right) \quad (1)$$

with $A = f^{[l]}(x)$, where $f^{[l]}(\cdot)$ denotes the output of the l -th layer¹. A_k is a k -th activation map of A and α_k is the coefficient (i.e., the importance) of A_k , respectively. N_l indicates the number of the activation maps of the l -th layer. This concept of linearly combining activation maps was firstly proposed by [20], leading to its variants.

Previous methods. Grad-CAM [14] decides the coefficient of a specific activation map by averaging the gradients over all activation neurons in that map. Grad-CAM++ [3], which is a modified version of Grad-CAM, focuses on positive influences of neurons considering higher-order derivatives. However, the gradients of deep neural networks tend to diminish due to the gradient saturation problem. Hence, using unmodified raw gradients induces failure of localization for relevant regions.

To overcome this limitation, gradient-free CAMs have been proposed. Score-CAM [18] overlaps normalized activation maps to an input image and makes predictions to acquire the coefficients. Ablation-CAM [4] defines a coefficient as the fraction of decline in the target output when the associated activation map is removed. They are free from the saturation issue, but time-consuming because they require N_l forward propagations to acquire the coefficients.

All the methods described above determine their coefficients in a heuristic way. XGrad-CAM [7] addresses this issue by suggesting two axioms. The authors derived the coefficients that satisfy the axioms as much as possible. However, their derivation is demonstrated only for ReLU-CNNs.

¹Conventionally, the last convolutional layer is used for the layer l because it is expected to provide the best compromise between high-level semantics and spatial information [14].

2.2. SHapley Additive exPlanations

Additive feature attribution method. SHAP [10] is a unified explanation framework for additive feature attribution methods. The methods follow:

$$g(z') = \phi_0 + \sum_{i=1}^M \phi_i z'_i \quad (2)$$

where g is an explanation model of an original prediction model f for a specific input x and a target class c . M is the number of input features and $z' \in \{0, 1\}^M$ indicates a binary vector in which each entry represents the existence of the corresponding original input feature; 1 for *presence* and 0 for *absence*. ϕ_i denotes an importance of the i -th feature and ϕ_0 is a baseline explanation. The methods are designed to ensure $g(z') \approx f^c(h_x(z'))$ whenever $z' \approx x'$, with a mapping function h_x that satisfies $x = h_x(x')$. While several existing attribution methods [2, 10, 12, 16] match Eq. (2), only one explanation model satisfies three desirable properties: *local accuracy*, *missingness*, and *consistency* [10].

SHAP values. A feature attribution of the explanation model which obeys Eq. (2) while adhering to the above three properties is defined as SHAP values [10] and can be formulated by:

$$\phi_i = \sum_{z' \subset x'} \frac{(M - |z'|)!(|z'| - 1)!}{M!} [f^c(h_x(z')) - f^c(h_x(z' \setminus i))] \quad (3)$$

where $|z'|$ denotes the number of non-zero entries in z' and $z' \subset x'$ indicates all z' vectors, where the non-zero entries are a subset of the non-zero entries in x' . In addition, $z' \setminus i$ means setting $z'_i = 0$. This definition of the SHAP values intimately aligns with the classic Shapley values [15].

2.3. Deep Learning Important Features

DeepLIFT [16] focuses on the difference between an *original* activation and a *reference* activation. It propagates the difference through a network to assign the contribution score to each input feature by linearizing non-linear components in the network. Through this technique, the gradient saturation problem is alleviated.

Let o represent the output of the target neuron and $x = (x_1, \dots, x_n)$ be inputs whose reference values are $r = (r_1, \dots, r_n)$. The contribution score of the i -th input feature $C_{\Delta x_i \Delta o}$ quantifies the influence of $\Delta x_i = x_i - r_i$ on $\Delta o = f^c(x) - f^c(r)$. In addition, DeepLIFT satisfies the *summation-to-delta* property as below:

$$\sum_{i=1}^n C_{\Delta x_i \Delta o} = \Delta o. \quad (4)$$

Note that if we set $C_{\Delta x_i \Delta o} = \phi_i$ and $f^c(r) = \phi_0$, then Eq. (4) matches Eq. (2). Therefore, DeepLIFT is also an addi-

tive feature attribution method. It *approximates* the SHAP values efficiently, satisfying the local accuracy and missingness [10].

3. Methodology

In this section, we clarify the problem formulation of CAM and suggest an approach to solve it analytically. First, we propose a framework that defines a linear explanation model and determines the coefficients of CAM based on the model. Then, we formulate the SHAP values of the activation maps as a unified solution for the framework. Finally, we introduce a fast approximation method for the SHAP values of the activation maps: LIFT-CAM.

3.1. Problem formulation of CAM

As identified in Eq. (1), CAM produces a visual explanation map L_{CAM}^c linearly w.r.t. the activation maps A_1, \dots, A_{N_l} except for ReLU, which is applied for the purpose of only considering the positive influence on the target class c . In addition, the complete activation maps A does not change for a given pair of the model f and the input image x . Thus, the quality of L_{CAM}^c is controlled by the coefficients $\alpha = (\alpha_1, \dots, \alpha_{N_l})$, which represent the importance scores of the associated activation maps. In sum, the purpose of CAM is to find α for a linear combination in order to generate L_{CAM}^c , which can reliably explain the target output $f^c(x)$.

3.2. Proposed framework

How can we acquire the desirable α in an analytic way? To this end, we first consider each activation map as an individual feature (i.e., we have N_l features) and define a binary vector $a' \in \{0, 1\}^{N_l}$ of the features. In the vector, an entry a'_k of 1 indicates that the corresponding A_k maintains its original activation values, and 0 means that it loses the values.

Next, we specify an explanation model g_{CAM} to interpret $f^c(x)$. Since the explanation map of CAM L_{CAM}^c is linear w.r.t. the activation maps A_1, \dots, A_{N_l} by definition, it is reasonable to assume that the explanation model g_{CAM} is also linear w.r.t. the binary variables of the activation maps a'_1, \dots, a'_{N_l} as follows:

$$g_{CAM}(a') = \alpha_0 + \sum_{k=1}^{N_l} \alpha_k a'_k. \quad (5)$$

Under this assumption, the problem of determining α in Eq. 1 can be reformulated into the problem of determining g_{CAM} that follows Eq. 5.

Eq. 5 matches Eq. 2 exactly. Furthermore, each individual feature of Eq. 5 (i.e., each activation map) is expected to represent distinct high-level semantic information. Therefore, in this work, we determine g_{CAM} using additive feature




Figure 1. Proposed framework for determining the coefficients of CAM. First, we build a linear explanation model. Next, we determine the importance scores of the activation maps by optimizing the explanation model, using additive feature attribution methods. Last, we use the scores as the coefficients of CAM.

attribution methods ([10] in Sec. 3.3, [16] in Sec. 3.4, [2] and [12] in Supplementary Material). Once we obtain α on the basis of g_{CAM} , we can use the values to generate L_{CAM}^c . Figure 1 shows our proposed framework which is described in this section.

3.3. SHAP values of activation maps

SHAP [10] is a model-agnostic method and the local accuracy, missingness, and consistency [10] are still desirable in Eq. 5. Accordingly, we adopt the SHAP values of the activation maps as a unified solution for our framework. Let F be a latter part of the original model f , from layer $l+1$ to layer $L-1$, where L represents the total number of layers in f . Namely, we have $F(A) = f^{[L-1]}(x)$. Additionally, we define a mapping function h_A that converts a' into the embedding space of A ; it satisfies $A = h_A(A')$, where A' is a vector of ones. Specifically, $a'_k = 1$ is mapped to A_k and $a'_k = 0$ to $\mathbf{0}$, which has the same dimension as A_k . Note that this is reasonable because A_k exerts no influence on L_{CAM}^c when it has values of 0 for all activation neurons in Eq. (1).

Now, the SHAP values of the activation maps w.r.t. the class c are formulated by:

$$\alpha_k^{\text{shap}} = \sum_{a' \subset A'} \frac{(N_l - |a'|)! (|a'| - 1)!}{N_l!} [F^c(h_A(a')) - F^c(h_A(a' \setminus k))] \quad (6)$$

where α_k^{shap} is the SHAP value of A_k and F^c denotes the target output of F . The above equation implies that α_k^{shap} can be obtained by averaging marginal prediction differences between *presence* and *absence* of A_k across Π_{total} that denotes a set of all possible feature orderings of $\{1, \dots, N_l\}$.

To reduce the computational burden, we suggest to use a subset $\Pi \subset \Pi_{\text{total}}$ instead of Π_{total} to estimate $\alpha^{\text{shap}} = (\alpha_1^{\text{shap}}, \dots, \alpha_{N_l}^{\text{shap}})$, as described in Algorithm 1. We refer

²It represents the logit layer which precedes the softmax layer.

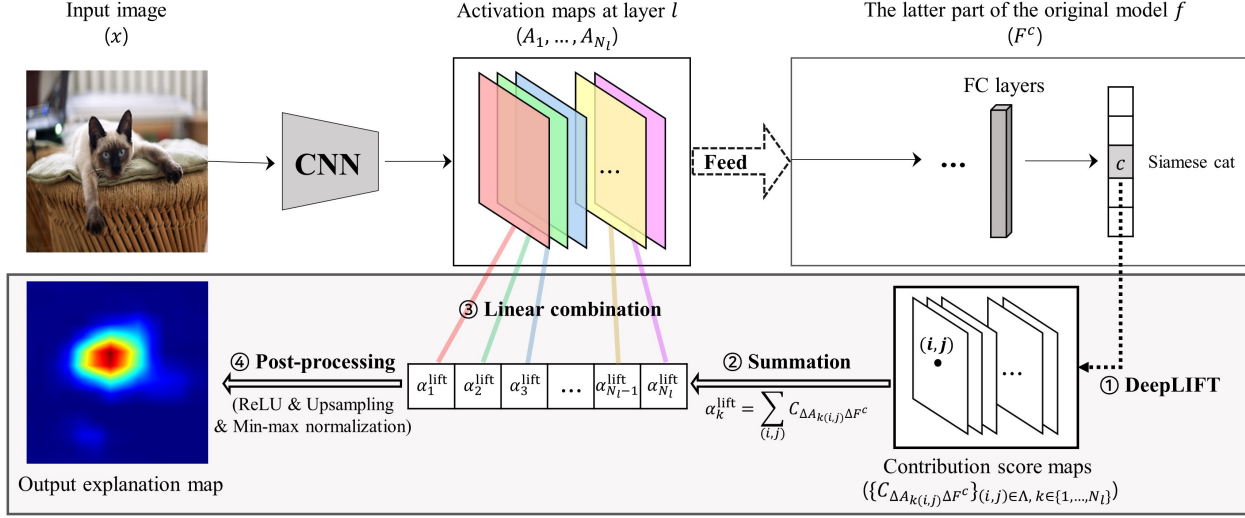


Figure 2. Overview of LIFT-CAM. First, we apply DeepLIFT from the target output up to the activation maps and acquire the contribution score maps, in which each pixel represents $C_{\Delta A_k(i,j)\Delta F^c}$. Next, we quantify the importance of each activation map by summing all the contribution scores of itself. Then, we perform a linear combination of $(\alpha_1^{\text{lift}}, \dots, \alpha_{N_l}^{\text{lift}})$ and A_1, \dots, A_{N_l} . Finally, we rectify the resulting map, upsample the map to the original image dimension, and normalize the map using the min-max normalization function.

to the algorithm of using $|\Pi|$ orderings as SHAP-CAM $_{|\Pi|}$ throughout the paper. The higher $|\Pi|$, the α from SHAP-CAM $_{|\Pi|}$ converges to α^{shap} by the law of large numbers. We validate the benefits of these SHAP attributions in terms of the faithfulness of L_{CAM}^c in Sec. 4.1. Analyses for acknowledged approximation methods for SHAP values, DeepSHAP [10] and KernelSHAP [10], are provided in Supplementary Material.

Algorithm 1 SHAP-CAM $_{|\Pi|}$

Input : F, c, h_A , and a subset $\Pi \subsetneq \Pi_{\text{total}}$

Output: $\alpha = (\alpha_1, \dots, \alpha_{N_l})$

Initialize: $\alpha \leftarrow \mathbf{0}$

for each ordering π in Π **do**

$a' \leftarrow \mathbf{0}$

for $i = 1, \dots, N_l$ **do**

$a'_{\pi(i)} \leftarrow 1$

$\alpha_{\pi(i)} \leftarrow \alpha_{\pi(i)} + F^c(h_A(a')) - F^c(h_A(a' \setminus \pi(i)))$

end

end

$\alpha \leftarrow \alpha / |\Pi|$

3.4. Efficient approximation: LIFT-CAM

Through the experiment in Sec. 4.1, we demonstrate that a faithful L_{CAM}^c can be achieved by α^{shap} . However, calculating the exact α^{shap} is almost impossible. Therefore, we need to consider an approximation approach. In this study, we propose a novel method, LIFT-CAM, that efficiently ap-

proximates α^{shap} using DeepLIFT³ [16].

First, we calculate the contribution score for every activation neuron at layer l using DeepLIFT through a single backward pass. Considering the summation-to-delta property of DeepLIFT, we define the contribution score of a specific activation map as the summation of the contribution scores of all neurons in that activation map, as follows:

$$\alpha_k^{\text{lift}} = C_{\Delta A_k \Delta F^c} = \sum_{(i,j) \in \Lambda} C_{\Delta A_k(i,j) \Delta F^c} \quad (7)$$

where $\Lambda = \{1, \dots, H\} \times \{1, \dots, W\}$ is a discrete activation dimension and $A_k(i,j)$ is an activation value at the (i,j) location of A_k . Note that Δ denotes the difference-from-reference and the reference values (i.e., the values corresponding to the *absent* feature) of all activation neurons are set to 0, aligning with SHAP. By this definition, $\alpha^{\text{lift}} = (\alpha_1^{\text{lift}}, \dots, \alpha_{N_l}^{\text{lift}})$ becomes a reliable solution for Eq. (5) while satisfying the local accuracy⁴ of SHAP as below:

$$\sum_{k=1}^{N_l} \alpha_k^{\text{lift}} = F^c(A) - F^c(\mathbf{0}). \quad (8)$$

Consequently, LIFT-CAM can estimate α^{shap} with a single backward pass while alleviating the gradient saturation problem [16]. Figure 2 shows an overview of our proposed

³DeepLIFT-Rescale is used for approximation because the method can be easily implemented by overriding gradient operators. This convenience enables LIFT-CAM to be easily applied to a large variety of tasks.

⁴ $\sum_{k=1}^{N_l} \alpha_k^{\text{shap}} = F^c(A) - F^c(\mathbf{0})$.

LIFT-CAM. Additionally, the following rationale motivates us to employ DeepLIFT for this problem.

DeepLIFT linearizes non-linear components to estimate the SHAP values [10]. Therefore, DeepLIFT attributions tend to deviate from the true SHAP values when passed through many overlapping non-linear layers during back-propagation (see Supplementary Material for details). However, for CAM, only the non-linearities in F matter. Since CAM usually uses the outputs of the last convolutional layer as its activation maps, almost all F of state-of-the-art architectures contain few non-linearities (e.g., the VGG family), or are even fully linear (e.g., the ResNet family). Thus, the SHAP values of the activation maps can be approximated with high precision by LIFT-CAM. Particularly, we can acquire the exact SHAP values using LIFT-CAM (i.e., $\alpha^{\text{lift}} = \alpha^{\text{shap}}$) for the architectures of linear F . The proof of this statement is provided in Supplementary Material.

3.5. Rethinking Ablation-CAM

The recently proposed Ablation-CAM [4] can be reinterpreted by our framework. Ablation-CAM defines the coefficients as below:

$$\alpha_k^{\text{ablation}} = \frac{F^c(h_A(A')) - F^c(h_A(A' \setminus k))}{F^c(h_A(A'))}. \quad (9)$$

Since Ablation-CAM uses this specific marginal difference as the coefficient, it can be deemed as another approximation method for α^{shap} . However, Ablation-CAM is computationally expensive requiring N_l forward simulations. In addition, the method does not satisfy the local accuracy of SHAP (i.e., $\sum_{k=1}^{N_l} \alpha_k^{\text{ablation}} \neq F^c(A) - F^c(\mathbf{0})$). This mismatch leads to less precise approximations compared to LIFT-CAM, resulting in less reliable explanations.

4. Experiments and Results

We now describe our experiments and show the results. In Sec. 4.1, we first validate the superiority of α^{shap} by evaluating the faithfulness of L_{CAM}^c generated by SHAP-CAM. Then, we demonstrate how closely LIFT-CAM can estimate α^{shap} in Sec. 4.2. These two experiments provide justification to opt for LIFT-CAM as a responsible method of determining α of CAM. We then evaluate the performance of LIFT-CAM on the object recognition task in the context of image classification, comparing it with the other CAMs: Grad-CAM, Grad-CAM++, XGrad-CAM, Score-CAM, and Ablation-CAM in Sec. 4.3. Finally, we apply LIFT-CAM to the visual question answering (VQA) task in Sec. 4.4 to check the scalability of the method.

For all experiments except VQA, we employ the public classification datasets: ImageNet [13] (ILSVRC 2012 validation set), PASCAL VOC [5] (2007 test set), and MS

COCO [9] (2014 validation set). In addition, the VGG16 network trained on each dataset is analyzed for the experiments (see Supplementary Material for the experiments of the ResNet50). We refer to the pretrained models from the torchvision⁵ package for ImageNet and the TorchRay package⁶ for VOC and COCO. For VQA, we use the fundamental architecture⁷ proposed by [1] and the VQA v2.0 dataset [8].

4.1. Validation of SHAP values

Faithfulness evaluation metrics. Intuitively, an explanation image w.r.t. the target class c can be generated using an original image x and a related visual explanation map L_{CAM}^c as below:

$$e^c = s(u(L_{\text{CAM}}^c)) \circ x \quad (10)$$

where $u(\cdot)$ indicates the upsampling operation into the original image dimension and $s(\cdot)$ denotes the min-max normalization function. The operator \circ refers to the Hadamard product. Hence, e^c preserves the information of x only in the region which L_{CAM}^c considers important.

In general, L_{CAM}^c is expected to recognize the regions which contribute the most to the model’s decision. Thus, we can evaluate the faithfulness of L_{CAM}^c on the object recognition task via the two metrics proposed by [3]: Increase in Confidence (IC) and Average Drop (AD), which are defined as:

$$\text{IC} = \frac{1}{N} \sum_{i=1}^N \mathbf{1}_{[Y_i^c < O_i^c]} \times 100, \quad (11)$$

$$\text{AD} = \frac{1}{N} \sum_{i=1}^N \frac{\max(0, Y_i^c - O_i^c)}{Y_i^c} \times 100, \quad (12)$$

where Y_i^c and O_i^c are the model’s softmax outputs of an i -th input image x_i and the associated explanation image e_i^c , respectively. N denotes the number of images and $\mathbf{1}_{[\cdot]}$ is an indicator function. Higher is better for the IC and lower is better for the AD.

However, both IC and AD evaluate the performance of the explanations via the *preservation* perspective; the region which is considered to be influential is maintained. We can also evaluate the performance through the opposite perspective (i.e., *deletion*); if we mute the region which is responsible for the target output, the softmax probability is expected to drop significantly. From this viewpoint, we suggest the Average Drop in Deletion (ADD) which can be defined as below:

$$\text{ADD} = \frac{1}{N} \sum_{i=1}^N \frac{(Y_i^c - D_i^c)}{Y_i^c} \times 100 \quad (13)$$

⁵<https://github.com/pytorch/vision/blob/master/torchvision>

⁶<https://github.com/facebookresearch/TorchRay>

⁷https://github.com/tbmoon/basic_vqa

	Increase in Confidence (%)			Average Drop (%)			Average Drop in Deletion (%)		
	<i>ImageNet</i>	<i>VOC</i>	<i>COCO</i>	<i>ImageNet</i>	<i>VOC</i>	<i>COCO</i>	<i>ImageNet</i>	<i>VOC</i>	<i>COCO</i>
*SHAP-CAM ₁	25.9	37.4	35.2	28.16	22.93	23.98	32.64	17.35	24.07
*SHAP-CAM ₁₀	26.2	42.7	40.1	27.54	17.53	19.07	32.99	19.95	27.50
*SHAP-CAM ₁₀₀	26.4	43.6	41.4	27.48	16.71	18.27	33.03	20.64	27.65

Table 1. Faithfulness evaluation on the object recognition task for SHAP-CAM₁, SHAP-CAM₁₀, and SHAP-CAM₁₀₀. The symbol * denotes averaging for 10 runs. We analyze 1,000 randomly selected images for each dataset. Higher is better for the IC and ADD. Lower is better for the AD.

where D^c is the softmax output of the inverted explanation image $e_{\text{inv}}^c = (\mathbf{1} - s(u(L_{\text{CAM}}^c))) \circ x$. Higher is better for this metric.

Faithfulness evaluation. Table 1 presents the comparative results of the IC, AD, and ADD between SHAP-CAM₁, SHAP-CAM₁₀, and SHAP-CAM₁₀₀. Each case is averaged for 10 simulations. We discover two important implications from Table 1. First, as $|\Pi|$ increases, the IC and ADD increase and the AD decreases, showing performance improvement. This result indicates that the closer the importances of the activation maps are to α^{shap} , the more effectively the distinguishable region of the target object is found. Second, even compared to the other CAMs (see Table 3), SHAP-CAM₁₀₀ shows the best performances for all cases. It reveals the adequacy of α^{shap} as the coefficients of CAM. However, this approach of averaging the marginal contributions of multiple orderings is impractical due to the significant computational burden. Therefore, we propose a cleverer approximation method: LIFT-CAM.

4.2. Approximation performance of LIFT-CAM

In this section, we quantitatively assess how precisely LIFT-CAM estimates α^{shap} . Since the exact α^{shap} is unattainable, we regard α from SHAP-CAM_{10k} as α^{shap} for comparison (see Supplementary Material for the justification of this assumption). Table 2 shows the cosine similarities between α from state-of-the-art CAMs, including LIFT-CAM, and α from SHAP-CAM_{10k}.

As shown in Table 2, LIFT-CAM presents the highest similarities for all datasets (greater than 0.9), which indicates high relevance between α^{lift} and α^{shap} . Even if Ablation-CAM also exhibits high similarities, the method falls behind LIFT-CAM due to dissatisfaction of the local accuracy. The other CAMs cannot approximate α^{shap} , providing consistently low similarities.

4.3. Performance evaluation of LIFT-CAM

The experimental results from Secs. 4.1 and 4.2 motivate us to generate visual explanations using LIFT-CAM. We verify the effectiveness of our LIFT-CAM by comparing the performances of the method with those of previous state-of-the-art saliency methods in terms of the quality of

	<i>ImageNet</i>	<i>VOC</i>	<i>COCO</i>
Grad-CAM	0.489	0.404	0.441
Grad-CAM++	0.385	0.329	0.412
XGrad-CAM	0.406	0.327	0.350
Score-CAM	0.195	0.181	0.157
Ablation-CAM	0.972	0.888	0.908
LIFT-CAM	0.980	0.918	0.924

Table 2. Cosine similarities between the coefficients from various CAMs and those from SHAP-CAM_{10k}. We analyze 500 randomly selected images for each dataset.

visualization, faithfulness, and localization.

4.3.1 Qualitative evaluation

Figure 3 provides qualitative comparisons between various saliency methods via visualization. Each row represents the visual explanation maps for each dataset. When compared to the other methods, our proposed method, LIFT-CAM, yields visually interpretable explanation maps for all cases. It clearly pinpoints the essential parts of the specific objects which are responsible for the classification results. This can be observed in the notebook case (row 1), for which the other visualizations cannot decipher the lower part of the notebook. Furthermore, LIFT-CAM alleviates pixel noise without highlighting trivial evidence. In the sheep case (row 2), the artifacts of the image are eliminated by LIFT-CAM and the exact location of the sheep is captured by the method. Last, the method successfully locates multiple objects in the apple case (row 3) by providing the clear object-focused map.

4.3.2 Faithfulness evaluation

IC, AD, and ADD. Table 3 shows the results of the IC, AD, and ADD for various CAMs. The three metrics can represent the object recognition performances of the saliency methods in a complementary way. LIFT-CAM provides the best results in most cases, with an exception of the ADD in COCO, where Ablation-CAM outperforms LIFT-




Figure 3. Visual explanation maps of state-of-the-art saliency methods and our proposed LIFT-CAM. Note that we apply a smoothing technique to Gradient [17] to acquire comparable visual explanation maps to those of CAMs, referring to [6].

	Increase in Confidence (%)			Average Drop (%)			Average Drop in Deletion (%)		
	<i>ImageNet</i>	<i>VOC</i>	<i>COCO</i>	<i>ImageNet</i>	<i>VOC</i>	<i>COCO</i>	<i>ImageNet</i>	<i>VOC</i>	<i>COCO</i>
Grad-CAM	24.0	32.7	31.9	31.89	30.73	30.74	30.60	17.43	25.66
Grad-CAM++	23.1	33.8	33.5	30.53	17.20	20.87	27.98	15.85	24.16
XGrad-CAM	25.0	30.5	31.3	31.36	30.04	29.92	30.48	17.09	24.95
Score-CAM	22.8	29.4	23.9	29.91	17.49	23.66	27.52	14.12	17.35
Ablation-CAM	24.1	34.4	35.0	29.41	25.49	23.99	32.52	19.42	26.75
LIFT-CAM	25.2	38.7	39.3	29.15	17.15	18.65	32.95	20.09	26.34

Table 3. Comparative evaluation of faithfulness on the object recognition task between various CAMs. We analyze 1,000 randomly selected images for each dataset (the same image samples as Table 1). Higher is better for the IC and ADD. Lower is better for the AD.

CAM. However, LIFT-CAM also provides a comparable result, showing that the difference is negligible. In addition, it should be noted that LIFT-CAM is much faster than Ablation-CAM since it requires only a single backward pass to calculate the coefficients. Thus, LIFT-CAM can determine which object is responsible for the model’s prediction, accurately and efficiently.

Area under probability curve. The above three metrics tend to be advantageous for methods which provide explanation maps of large magnitude. To exclude the influence of the magnitude, we can binarize the explanation map with two opposite perspectives: *insertion* and *deletion* [11]. First, we threshold $s(u(L_{CAM}^c))$ with $\delta \in [0 : 0.025 : 1]$ (i.e., 1 for top $100 \times \delta$ percentile pixels and 0 for the others) and acquire the corresponding target softmax outputs O^c for the insertion and D^c for the deletion. Using the softmax outputs, we can draw a probability curve as a function of δ . Finally, we can calculate the area under the probability curve (AUC).

As shown in Table 4, LIFT-CAM provides the most reliable results, presenting the highest insertion AUC and the lowest deletion AUC. Through this experiment, we demon-

	Insertion	Deletion
Grad-CAM	0.4427	0.0891
Grad-CAM++	0.4350	0.0969
XGrad-CAM	0.4457	0.0883
Score-CAM	0.4345	0.1002
Ablation-CAM	0.4685	0.0873
LIFT-CAM	0.4712	0.0866

Table 4. AUC results in terms of the insertion and deletion. The values are averaged for 1,000 randomly selected images from ImageNet. Higher is better for the insertion and lower is better for the deletion.

strate that LIFT-CAM succeeds in sorting the pixels according to the contributions to the target result.

4.3.3 Localization evaluation

It is reasonable to expect that a dependable explanation map overlaps with a target object. Therefore, we can also assess the reliability of the map via localization ability in ad-

	Grad-CAM	Grad-CAM++	XGrad-CAM	Score-CAM	Ablation-CAM	LIFT-CAM
Proportion (%)	47.76	49.14	47.91	51.14	51.87	52.43

Table 5. Proportions of energy located in bounding boxes for various CAMs. The values are averaged for 1,000 randomly selected images from ImageNet. Higher is better.

dition to the softmax probability. [18] newly proposed an improved version of a pointing game [19], named as an energy-based pointing game. This gauges how much energy of the explanation map interacts with the bounding box of the target object. For this, an evaluation metric can be formulated as follows:

$$\text{Proportion} = \frac{\sum_{(i,j) \in \text{bbox}} s(u(L_{\text{CAM}}^c))_{(i,j)}}{\sum_{(i,j) \in \Lambda'} s(u(L_{\text{CAM}}^c))_{(i,j)}} \quad (14)$$

where $\Lambda' = \{1, \dots, H'\} \times \{1, \dots, W'\}$ is an original image dimension and $s(u(L_{\text{CAM}}^c))_{(i,j)}$ denotes a min-max normalized importance at pixel location (i, j) . Higher is better for this metric.

The proportions of various methods are reported in Table 5. LIFT-CAM shows the highest proportion compared to the other methods. This implies that LIFT-CAM produces a compact explanation map which focuses on the essential parts of the images without trivial noise.

4.4. Application to VQA

We also apply LIFT-CAM to VQA to demonstrate the applicability of the method. We consider the standard VQA model [1] which consists of a CNN and a recurrent neural network in parallel. They function to embed images and questions, respectively. The two embedded vectors are fused and entered into a classifier to produce an answer.

Figure 4 illustrates the explanation maps of Grad-CAM and LIFT-CAM for VQA. LIFT-CAM highlights the regions in the given images that are more relevant to the question and answer pairs than those identified with Grad-CAM. Additionally, since this is a classification problem, the IC, AD, and ADD can be evaluated with fixed question embeddings. Table 6 shows the comparison results between Grad-CAM and LIFT-CAM in terms of those metrics. As demonstrated in the table, LIFT-CAM outperforms Grad-CAM for all of the metrics. This indicates that LIFT-CAM is better at figuring out the essential parts of images, which can serve as evidences for the answers to the questions.

5. Conclusion

In this work, we propose a novel analytic framework which determines the coefficients of CAM by optimizing a linear explanation model, using additive feature attribution methods. As desirable solutions for the explanation

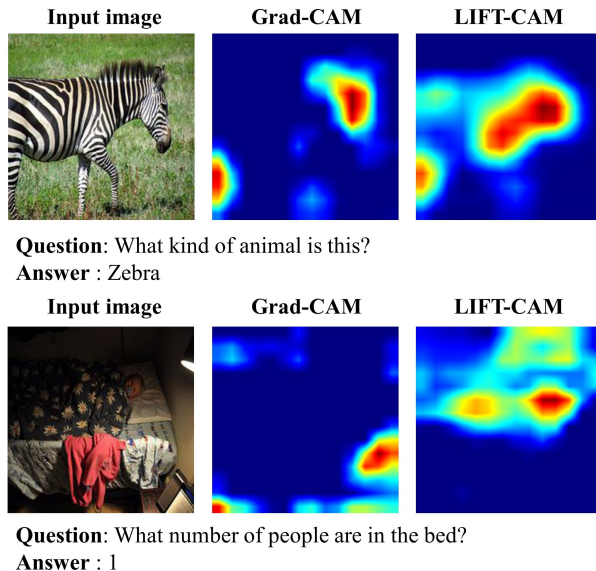


Figure 4. Visual explanation maps of Grad-CAM and LIFT-CAM for VQA.

	Grad-CAM	LIFT-CAM
Increase in Confidence (%)	41.95	43.39
Average Drop (%)	16.71	14.14
Average Drop in Deletion (%)	9.09	9.58

Table 6. Faithfulness evaluation for the VQA task. The values are averaged for the complete validation set. Higher is better for the IC and ADD. Lower is better for the AD.

model, we introduce several approaches including LIFT-CAM, throughout this paper and Supplementary Material. In addition, we show that Ablation-CAM can also be unified into this framework.

Our proposed LIFT-CAM approximates the SHAP values of the activation maps, which is a unified solution for the explanation model, with a single backward pass. The method provides qualitatively enhanced visual explanations compared with the other CAMs. Furthermore, it achieves state-of-the-art results on various quantitative evaluation metrics.

References

- [1] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015.
- [2] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS one*, 10(7):e0130140, 2015.
- [3] Aditya Chattopadhyay, Anirban Sarkar, Prantik Howlader, and Vineeth N Balasubramanian. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 839–847. IEEE, 2018.
- [4] Saurabh Desai and Harish G Ramaswamy. Ablation-cam: Visual explanations for deep convolutional network via gradient-free localization. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 972–980. IEEE, 2020.
- [5] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010.
- [6] Ruth Fong, Mandela Patrick, and Andrea Vedaldi. Understanding deep networks via extremal perturbations and smooth masks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2950–2958, 2019.
- [7] Ruigang Fu, Qingyong Hu, Xiaohu Dong, Yulan Guo, Yinghui Gao, and Biao Li. Axiom-based grad-cam: Towards accurate visualization and explanation of cnns. In *31th British Machine Vision Conference, BMVC 2020*.
- [8] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [9] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [10] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in neural information processing systems*, pages 4765–4774, 2017.
- [11] Vitali Petsiuk, Abir Das, and Kate Saenko. Rise: Randomized input sampling for explanation of black-box models. In *29th British Machine Vision Conference, BMVC 2018*.
- [12] Marco Tulio Ribeiro, Sameer Singh, and Carlos 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*, pages 1135–1144, 2016.
- [13] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [14] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [15] Lloyd S Shapley. A value for n-person games. Technical report, Rand Corp Santa Monica CA, 1952.
- [16] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. In *International Conference on Machine Learning*, pages 3145–3153, 2017.
- [17] K. Simonyan, A. Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *CoRR*, abs/1312.6034, 2014.
- [18] Haofan Wang, Zifan Wang, Mengnan Du, Fan Yang, Zijian Zhang, Sirui Ding, Piotr Mardziel, and Xia Hu. Score-cam: Score-weighted visual explanations for convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 24–25, 2020.
- [19] Jianming Zhang, Sarah Adel Bargal, Zhe Lin, Jonathan Brandt, Xiaohui Shen, and Stan Sclaroff. Top-down neural attention by excitation backprop. *International Journal of Computer Vision*, 126(10):1084–1102, 2018.
- [20] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016.