

# Enhancing Interpretability for Vision Models via Shapley Value Optimization

Kanglong Fan<sup>1</sup>, Yunqiao Yang<sup>1</sup>, Chen Ma<sup>1\*</sup>

<sup>1</sup>Department of Computer Science, City University of Hong Kong  
kanglofan2-c@my.cityu.edu.hk, yqyang.cs@my.cityu.edu.hk, chenma@cityu.edu.hk

## Abstract

Deep neural networks have demonstrated remarkable performance across various domains, yet their decision-making processes remain opaque. Although many explanation methods are dedicated to bringing the obscurity of DNNs to light, they exhibit significant limitations: post-hoc explanation methods often struggle to faithfully reflect model behaviors, while self-explaining neural networks sacrifice performance and compatibility due to their specialized architectural designs. To address these challenges, we propose a novel self-explaining framework that integrates Shapley value estimation as an auxiliary task during training, which achieves two key advancements: 1) a fair allocation of the model prediction scores to image patches, ensuring explanations inherently align with the model’s decision logic, and 2) enhanced interpretability with minor structural modifications, preserving model performance and compatibility. Extensive experiments on multiple benchmarks demonstrate that our method achieves state-of-the-art interpretability.

## Introduction

Deep neural networks (DNNs) have achieved remarkable success across numerous applications (Redmon et al. 2016; Vinyals et al. 2015; Antol et al. 2015). Despite their impressive capabilities, a significant challenge persists: the inherent lack of interpretability in their decision-making processes. This limitation raises critical concerns about the reliability and safety of DNNs, particularly in high-stakes applications where model interpretability is crucial for ensuring reliability and accountability (Molnar 2020; Borys et al. 2023).

Current approaches for explaining DNN predictions can be broadly categorized into two paradigms: post-hoc explaining and self-explaining methods. Post-hoc explanation methods, including gradient-based (Sundararajan, Taly, and Yan 2017; Selvaraju et al. 2017; Yang, Wang, and Bilgic 2023; Li et al. 2023a), perturbation-based techniques (Petsiuk, Das, and Saenko 2018; Fong, Patrick, and Vedaldi 2019; Jethani et al. 2022; Covert, Kim, and Lee 2023), counterfactual-generation-based (Bass et al. 2022; Xie et al. 2024) and attention-based (Abnar and Zuidema 2020; Chefer, Gur, and Wolf 2021a; Qiang et al. 2022;

Wu et al. 2024a), are typically applied independently of model training. While widely adopted, these methods often produce unfaithful explanations that inadequately represent model behaviors (Adebayo et al. 2018; Yang and Kim 2019; Kindermans et al. 2019; Hesse, Schaub-Meyer, and Roth 2024). In contrast, self-explaining neural networks (SENNs) (Chen et al. 2019; Brendel and Bethge 2019; Wang, Wang, and Inouye 2021; Hesse, Schaub-Meyer, and Roth 2021; Böhle, Fritz, and Schiele 2022; Chen et al. 2023; Nauta et al. 2023; De Santi et al. 2024; Arya et al. 2024) integrate interpretability directly into their models through specialized architecture designs. By construction, SENNs generate intrinsic explanations aligned with the model’s decision logic, offering greater faithfulness compared to post-hoc methods. However, SENNs face three major limitations: 1) they often require training from scratch, limiting their compatibility with pre-trained models (Arya et al. 2024); 2) their specialized architectural designs (Chen et al. 2019; Hesse, Schaub-Meyer, and Roth 2021; Chen et al. 2023) often lead to degraded performance compared to standard DNNs; 3) the inclusion of interpretability modules (Brendel and Bethge 2019; Wang, Wang, and Inouye 2021; Chen et al. 2023) introduces memory and computational overhead, hindering scalability.

Recent research (Lundberg and Lee 2017; Jethani et al. 2022; Covert, Kim, and Lee 2023) establishes the Shapley value (Shapley 1953) as a principled approach for model interpretation, as it quantifies the marginal contribution of individual input components (e.g., image patches) to predictions. While many approaches have made progress in incorporating Shapley value for model interpretation, they still face significant limitations in computational efficiency and attribution accuracy. Conventional approaches (Castro, Gómez, and Tejada 2009; Strumbelj and Kononenko 2010; Lundberg and Lee 2017; Covert and Lee 2021; Mitchell et al. 2022) require extensive model inferences to approximate Shapley values, making them computationally costly and impractical for many applications. Moreover, the discrepancy between masked data used for Shapley value estimation during testing and the unmasked data during training may introduce additional attribution errors. Alternative methods such as FastSHAP (Jethani et al. 2022) and ViT-Shapley (Covert, Kim, and Lee 2023) introduce an auxiliary surrogate model to process masked images, followed by

\*Corresponding author.

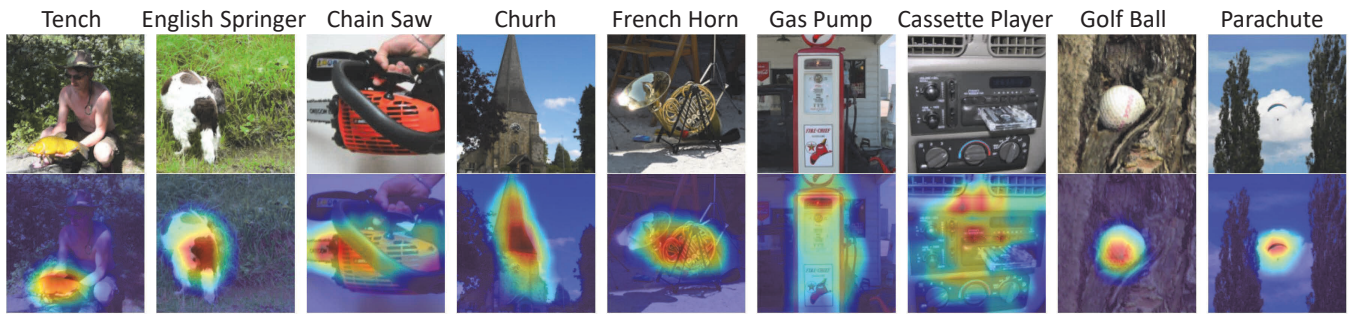


Figure 1: Visualization of our method’s explanations on ImageNet. The top row shows the input images and the bottom displays the corresponding explanations.

training an explainer to fastly estimate Shapley values. However, the explainer ultimately interprets the surrogate model rather than the model to be explained.

To address these challenges, we propose a multitask learning framework that intrinsically integrates Shapley value estimation directly into the model’s optimization process. By simultaneously optimizing the model for the primary task (e.g., image classification) and Shapley value estimation with an appropriate trade-off parameter, our framework (see Figure 2) achieves self-interpretable predictions without compromising the performance of the primary task. Unlike post-hoc explanation methods that risk misrepresenting a model’s decision logic, our approach learns explanations as part of the model’s reasoning process, ensuring alignment between explanations and decision logic (see Figure 1). Compared to existing SENNs, our framework requires minor architectural modifications and does not rely on external interpretation modules, preserving computational efficiency and model compatibility. Furthermore, by naturally incorporating masked images during training, our approach circumvents the discrepancy problem associated with post-hoc Shapley estimation without the need for a surrogate model.

In summary, our key contributions are as follows:

- we propose a novel optimization framework that jointly learns Shapley value estimation and the primary task, enhancing explanation faithfulness while preserving the primary task performance;
- we validate our method on image classification and image-text matching tasks, and demonstrate that it outperforms state-of-the-art post-hoc explaining and self-explaining methods.

## Related Work

Contemporary explainable artificial intelligence encompasses post-hoc explaining and self-explaining methods (Nauta 2023). Post-hoc methods including gradient-based approaches (Sundararajan, Taly, and Yan 2017; Selvaraju et al. 2017; Yang, Wang, and Bilgic 2023; Li et al. 2023a), perturbation-based techniques (Ribeiro, Singh, and Guestrin 2016; Lundberg and Lee 2017; Petsiuk, Das, and Saenko 2018; Jethani et al. 2022; Covert, Kim, and Lee 2023), counterfactual-generation-based (Bass et al. 2022; Xie et al. 2024) and attention-based (Abnar and Zuidema

2020; Chefer, Gur, and Wolf 2021a; Qiang et al. 2022; Wu et al. 2024a), often face challenges with reliability and faithfulness (Adebayo et al. 2018; Yang and Kim 2019; Hesse, Schaub-Meyer, and Roth 2024). Among perturbation-based methods, conventional Shapley value approaches (Castro, Gómez, and Tejada 2009; Strumbelj and Kononenko 2010; Lundberg and Lee 2017; Covert and Lee 2021; Mitchell et al. 2022) are constrained by high computational complexity, while existing acceleration techniques (Jethani et al. 2022; Covert, Kim, and Lee 2023) often fail to produce accurate Shapley value estimates.

Models with inherent interpretability (Chen et al. 2019; Brendel and Bethge 2019; Hesse, Schaub-Meyer, and Roth 2021; Nauta, Van Bree, and Seifert 2021; Böhle, Fritz, and Schiele 2022; Nauta et al. 2023; De Santi et al. 2024) offer a more promising direction by incorporating transparency mechanisms into their architecture. ProtoPNet (Chen et al. 2019), PIPNet (Nauta et al. 2023) and PIMPNet (De Santi et al. 2024) leverages prototype comparison, BagNets (Brendel and Bethge 2019) employs patch-based feature extraction,  $\mathcal{X}$ -DNNs (Hesse, Schaub-Meyer, and Roth 2021) enables efficient Integrated Gradients computation, and B-cos networks (Böhle, Fritz, and Schiele 2022) implement structural alignment for feature identification. Recent advances (Wang, Wang, and Inouye 2021; Chen et al. 2023) seek to incorporate Shapley value into self-explaining architectures. While enhancing the interpretability, the specialized architecture designs could increase memory and computational overhead, and compromise the model’s compatibility and performance. B-cosified (Arya et al. 2024) extends compatibility with pre-trained backbone models such as ResNet (He et al. 2016), and Vision Transformers (ViT) (Dosovitskiy et al. 2020). Although with minor architectural changes, the recent work (Alkhatib et al. 2025) primarily focuses on tabular data applications and small-scale classification tasks. Our method integrates Shapley value attribution directly into the learning process as an auxiliary task, tailored for fine-tuning modern, large-scale pretrained models like ViT and CLIP. We demonstrate its effectiveness and superior explanation quality across multiple tasks, including image classification, object localization (segmentation), and image-text matching on large-scale benchmarks.




Figure 2: Schematic representation of our proposed dual-branch self-explaining framework. The upper branch processes the original image to get the explanations and classification logits, while the lower branch processes the masked image to supervise the explanations via Shapley Loss.

## Background

### Shapley Value

Consider a feature space denoted by  $\mathcal{X}$ , where each observation  $\mathbf{x}$  comprises  $N$  distinct components expressed as  $\mathbf{x} = (\mathbf{m}_1, \dots, \mathbf{m}_N)$ . In a classification framework, the target class  $y$  is drawn from the set  $\mathcal{Y} = \{1, \dots, K\}$ . The predictive model is represented by  $\mathbf{f}(\mathbf{x}; \theta)$ , where  $\theta$  represents the model parameters and  $f_y(\mathbf{x}; \theta)$  corresponds to the logit value for the class  $y$ . A binary mask  $\mathbf{s} \in \{0, 1\}^N$  indicates the selection of components from  $\mathbf{x}$ , where the subset  $\mathbf{x}_{\mathbf{s}} = \{\mathbf{m}_i | s_i = 1, i = 1, \dots, N\}$ . The vectors  $\mathbf{1}$  and  $\mathbf{0}$  in  $\mathbb{R}^N$  represent all-ones and all-zeros vectors, respectively, while  $\mathbf{e}_i \in \mathbb{R}^N$  represents the  $i$ -th canonical basis vector (a one-hot vector with 1 at the  $i$ -th position and 0 elsewhere).  $\mathbf{1}^\top \mathbf{s}$  yields the number of components in a subset. Given an image  $\mathbf{x}$ , the value function  $v: 2^N \mapsto \mathbb{R}$  quantifies the predictive value of the subset  $\mathbf{x}_{\mathbf{s}}$  for class  $y$ :

$$v_y(\mathbf{s}) = f_y(\mathbf{x}_{\mathbf{s}}; \theta). \quad (1)$$

Within cooperative game theory, the Shapley value provides a principled method to fairly allocate collective gains among coalition participants (Shapley 1953). Given the value function  $v$  that maps subsets of participants to their aggregated contributions, the Shapley value  $\phi_i$  quantifies the marginal contributions of component  $i$ , computed as:

$$\phi_i = \sum_{\mathbf{s} | s_i = 0} \frac{(\mathbf{1}^\top \mathbf{s})!(N - 1 - \mathbf{1}^\top \mathbf{s})!}{N!} (v(\mathbf{s} + \mathbf{e}_i) - v(\mathbf{s})). \quad (2)$$

Here, the summation extends over all masks  $\mathbf{s}$  excluding component  $i$ . This formulation evaluates the incremental impact of adding component  $i$  to every feasible coalition, adhering to a theoretically rigorous methodology for assessing individual contributions (Shapley 1953).

### KernelSHAP

KernelSHAP (Lundberg and Lee 2017) reformulates the Shapley value estimation as a weighted least squares optimization problem, leveraging strategic coalition sampling to

efficiently approximate the solution. Formally, the Shapley values  $\Phi \in \mathbb{R}^N$  are obtained by solving:

$$\begin{aligned} \min_{\Phi} \mathbb{E}_{p(\mathbf{s})} \left[ (v(\mathbf{s}) - v(\mathbf{0}) - \mathbf{s}^\top \Phi)^2 \right] \\ \text{s.t.} \quad \mathbf{1}^\top \Phi = v(\mathbf{1}) - v(\mathbf{0}). \end{aligned} \quad (3)$$

The coalition weighting kernel  $p(\mathbf{s})$ , derived from (Charnes et al. 1988), prioritizes coalitions with very small or large subsets:

$$p(\mathbf{s}) \propto \frac{(N - 1)(\mathbf{1}^\top \mathbf{s} - 1)!(N - \mathbf{1}^\top \mathbf{s} - 1)!}{N!}, \quad (4)$$

where  $0 < \mathbf{1}^\top \mathbf{s} < N$ . Theoretical analyses (Lundberg and Lee 2017; Covert and Lee 2021) demonstrate that this estimator is statistically consistent and unbiased under mild assumptions. However, the computational cost remains prohibitive, as separate optimization problems must be solved for every image  $\mathbf{x}$ .

## Method

In this section, we propose a dual-branch framework that jointly optimizes the primary task (e.g., image classification) and Shapley value estimation to enhance the interpretability of predictive models. First, the input image  $\mathbf{x}$  is partitioned into  $N$  patches, each treated as an independent component. The first branch processes the original image and employs a patch-based soft voting for the primary task, where the contribution of each patch is weighted by its predicted Shapley value (i.e., explanations). The second branch processes the masked image to supervise the explanations. The framework overview is shown in Figure 2.

### Patch-Level Decomposition and Representation

Given an input image  $\mathbf{x} \in \mathbb{R}^{C \times H \times W}$ , we decompose it into  $N = H_p \times W_p$  patches with patch size  $P$ , each treated as a component. These patches are processed through a vision model  $\Phi$  with parameters  $\theta$ , comprising an encoder  $\text{Enc}(\cdot)$

and a prediction head  $\text{Head}(\cdot)$ . The model directly predicts the Shapley value of each patch as follows:

$$\mathbf{E} = \text{Enc}(\mathbf{x}) \in \mathbb{R}^{N \times L}, \quad (5)$$

$$\Phi(\mathbf{x}; \theta) = \text{Head}(\mathbf{E}) \in \mathbb{R}^{N \times K}, \quad (6)$$

where  $L$  denotes the feature dimension per patch, and  $\Phi(\mathbf{x}; \theta)$  represents class-specific shapley values for  $K$  classes across all  $N$  patches.

### Patch-Based Soft Voting

In our framework, image-level predictions are then obtained by aggregating predicted Shapley values  $\Phi(\mathbf{x}; \theta)$  for each patch through a soft voting mechanism:

$$\mathbf{f}(\mathbf{x}; \theta) = \sum_{n=1}^N \Phi(\mathbf{x}; \theta)_{n,:} \in \mathbb{R}^K. \quad (7)$$

This aggregation mechanism, rooted in cooperative game theory, ensures fair attribution of each patch’s marginal contribution to the final outcome of the primary task. Taking image classification as an example, the objective is then optimized using a standard cross-entropy loss:

$$\ell_{\text{cls}} = \mathbb{E}_{p(\mathbf{x})} \text{CrossEntropy}(\mathbf{f}(\mathbf{x}; \theta), y^*), \quad (8)$$

where  $y^*$  denotes the ground truth label of  $\mathbf{x}$ .

### Shapley Value Estimation via Self-Supervised Learning

We present our self-supervised framework for Shapley value estimation using the formulation in Problem (3).

**Mask Sampling** The process begins by strategically sampling the binary mask  $\mathbf{s}$  from the Shapley kernel defined in Eq. 4. Each mask represents a coalition of patches, where  $s_i = 1$  indicates the inclusion of patch  $i$  in the coalition.

**Masked Image Processing** Building on the patch decomposition outlined previously, we process each input image  $\mathbf{x}$  by applying the sampled mask  $\mathbf{s}$  to discard the excluded patches. We employ specialized masking strategies for different architectures: standard zero-masking for CNN, and physical patch removal for ViT similar to the approach proposed in Masked Autoencoders (He et al. 2022). The resulting masked image  $\mathbf{x}_s$  is then encoded through:

$$\mathbf{E}_s = \text{Enc}(\mathbf{x}_s) \in \mathbb{R}^{M \times L}, \quad (9)$$

$$\Phi(\mathbf{x}_s; \theta) = \text{Head}(\mathbf{E}_s) \in \mathbb{R}^{M \times K}, \quad (10)$$

where  $M$  is the number of included patches.

**Shapley Loss Formulation** Following FastSHAP (Jethani et al. 2022), we estimate Shapley values through a parameterized function  $\Phi(\mathbf{x}; \theta)$ , which simultaneously serves as the prediction mechanism via the aggregation in Eq. 7. In addition to the classification loss, we optimize  $\theta$  through the Shapley loss derived from Problem (3):

$$\ell_{\text{shap}} = \mathbb{E}_{p(\mathbf{x}), p(y), p(\mathbf{s})} \left[ \left( v_y(\mathbf{s}) - v_y(\mathbf{0}) - \mathbf{s}^\top \Phi(\mathbf{x}; \theta)_{:,y} \right)^2 \right], \quad (11)$$

where  $p(y)$  denotes uniform class sampling, and  $\Phi(\mathbf{x}; \theta)_{:,y}$  denotes class-specific Shapley value for class  $y$ .

We implement patch masking through either zero-masking or physical removal. This methodology establishes the baseline value  $v(\mathbf{0})$  as equivalent to the model’s empty-input prediction. By standardizing  $v(\mathbf{0}) = 0$  to streamline optimization, the constraint in Problem (3) simplifies to:

$$\mathbf{1}^\top \Phi(\mathbf{x}; \theta)_{:,y} = v_y(\mathbf{1}), \quad (12)$$

and the constraint is inherently satisfied by the aggregation mechanism defined in Eq. 7 and the value function formulation in Eq. 1. The Shapley loss is thus simplified as a self-supervised objective:

$$\ell_{\text{shap}} = \mathbb{E} \left[ \left( \sum_{m=1}^M \Phi(\mathbf{x}_s; \theta)_{m,y} - \mathbf{s}^\top \Phi(\mathbf{x}; \theta)_{:,y} \right)^2 \right], \quad (13)$$

where  $\sum_{m=1}^M \Phi(\mathbf{x}_s; \theta)_{m,y}$  is the predictive value for class  $y$  given masked input  $\mathbf{x}_s$ , which is equal to  $v_y(\mathbf{s})$ .

### Training Procedure

The total loss is a combination of the primary objective (e.g., classification loss) and Shapley loss as:

$$\ell = \ell_{\text{cls}} + \lambda \cdot \ell_{\text{shap}}, \quad (14)$$

where  $\lambda$  serves as a trade-off hyperparameter. The pseudocode of the training procedure is provided in the Appendix.

## Experiments

This section presents a comprehensive evaluation of our proposed self-explaining model through multiple experiments across diverse datasets and applications. We begin with model implementations, followed by extensive comparisons with state-of-the-art methods. Code is available at <https://github.com/kalofan/SENNs-via-SVO>.

### Method Implementations

We evaluate our framework using two standard architectures, ViT-B and ResNet-50, with an input image size of  $3 \times 224 \times 224$ . While preserving the core functionality of each model, we introduce specific adaptations to achieve patch decomposition and spatial feature representation critical for interpretable Shapley value estimation.

**Feature Extraction.** For ViT, we remove the class token [cls] and retain all the  $N$  patch tokens, generating patch feature representations with the size of  $N \times L$ . For CNN architectures, we preserve the inherent structures while reshaping outputs to match the ViT format. Specifically, a standard ResNet-50 encoder produces feature maps of dimensions  $L \times \sqrt{N} \times \sqrt{N}$ , which we reshape to  $N \times L$  patch representations:  $\mathbf{E} = \text{Reshape}(\text{Enc}_{\text{CNN}}(\mathbf{x}))$ .

**Masked Feature Extraction.** Given masks  $\mathbf{s}$ , we retain only  $M$  unmasked patches and remove masked patches along with [cls] for ViT, thereby extracting patch representations with the size of  $M \times L$ . For CNN models, direct removal of masked patches is infeasible. Instead, we first generate masked input  $\mathbf{x}_s$  by zeroing out masked patches.  $\mathbf{x}_s$

Model	Acc@1	#Params	#FLOPs
BagNet	63.5	18M	16.4B
HarsanyiNet	67.1	45M	6.1B
ViT-B	75.0	86M	17.6B
B-cos-V	74.4	86M	17.6B
B-cosified-V	75.3	86M	17.8B
Ours-V	<b>76.8</b>	86M	17.6B
ResNet-50	75.3	26M	4.1B
ProtoPNet	72.3	29M	4.4B
$\mathcal{X}$ -DNNs	72.9	26M	4.1B
B-cos-R	75.9	26M	4.3B
B-cosified-R	76.5	26M	4.3B
Ours-R	<b>76.3</b>	26M	6.6B

Table 1: Performance comparison of different models on ImageNet classification. The table is organized into three sections: custom CNNs (top), ViT-B-based models ("-V", middle), and ResNet-50-based models ("-R", bottom).

is then processed through the CNN encoder, followed by a Gather operation to select patch features exclusively from unmasked positions:

$$E_s = \text{Gather}(\text{Reshape}(E_{\text{ENC}_{\text{CNN}}}(\mathbf{x}_s)), s) \quad (15)$$

**Hyperparameters.** For ViT-B, we employ a patch size of  $P = 16$ , producing  $N = 196$  with a feature dimension  $L = 768$ . For ResNet-50, we modify the stride parameter from 2 to 1 in the initial residual block’s convolutional layer in stage 4, generating feature maps of dimensions  $L \times \sqrt{N} \times \sqrt{N} = 2048 \times 14 \times 14$ , which are reshaped to  $N \times L = 196 \times 2048$ . To accelerate convergence, all backbones are initialized with ImageNet pre-trained weights and fine-tuned by our method, underscoring our method’s compatibility with standard architectures. Additional training details and time costs are provided in the Appendix.

## Vision Models

**Prediction Performance** Table 1 presents a comparative analysis of our method against pre-trained ResNet-50, ViT-B, and state-of-the-art SENNs in terms of classification performance on ImageNet (Russakovsky et al. 2015). Evaluated SENNs include ProtoPNet (Chen et al. 2019), BagNet (Brendel and Bethge 2019),  $\mathcal{X}$ -DNNs (Hesse, Schaub-Meyer, and Roth 2021), B-cos (Böhle, Fritz, and Schiele 2022), HarsanyiNet (Chen et al. 2023) and B-cosified (Arya et al. 2024). Metrics include top-1 accuracy (Acc@1), number of model parameters, and computational complexity (FLOPs). BagNet and HarsanyiNet, shown in the upper section of the table, employ custom CNN models. The models in the center section are based on ViT-B architectures, whereas those in the lower section primarily utilize ResNet-50 variants.

Our ViT-B implementation achieves 76.8% top-1 accuracy, surpassing the vanilla ViT-B (75.0%) by 1.8% while retaining identical parameters and FLOPs. This demonstrates that our interpretability framework enhances rather than compromises classification performance. Similarly, when

Backbone	XAI Method	AOPC $\uparrow$	LOdds $\downarrow$	SaCo $\uparrow$
ViT-B	Grad-CAM	0.547	-4.179	0.1187
ViT-B	Rollout	0.671	-5.186	0.2887
ViT-B	GAME	0.707	-5.391	0.4353
ViT-B	CheferLRP	0.715	-5.587	0.4411
ViT-B	ATTCAT	0.647	-4.387	0.3454
ViT-B	ViTShapley	0.676	-5.056	0.4095
ViT-B	TokenTM	0.755	-5.668	0.4827
ResNet-50	IG	0.643	-4.689	0.2929
ResNet-50	Grad-CAM	0.604	-4.275	0.2182
ResNet-50	FastSHAP	0.683	-4.976	0.3197
ResNet-50	IDGI	0.706	-5.225	0.4057
ResNet-50	CAE	0.601	-4.107	0.2224
ProtoPNet	ProtoPNet	0.577	-4.075	0.2053
BagNet	BagNet	0.794	-5.864	0.5124
$\mathcal{X}$ -DNNs	$\mathcal{X}$ -DNNs	0.636	-4.678	0.2703
B-cos-V	B-cos	0.714	-5.474	0.4334
B-cos-R	B-cos	0.717	-5.486	0.4175
HarsanyiNet	HarsanyiNet	0.798	-5.903	0.5147
B-cosified-V	B-cosified	0.712	-5.473	0.4323
B-cosified-R	B-cosified	0.718	-5.489	0.4207
Ours-V	Ours	0.803	-5.987	<b>0.5196</b>
Ours-R	Ours	<b>0.806</b>	<b>-6.036</b>	0.5193

Table 2: Comparison of explanation performance on the ImageNet dataset. The upper and middle sections report post-hoc methods applied to vanilla ViT-B and ResNet-50 backbones, respectively, while the lower section shows results for various SENNs. The "-V" and "-R" suffixes indicate ViT and ResNet backbones, respectively.

applied to the ResNet-50, our method achieves 76.3% accuracy, outperforming the vanilla ResNet-50 (75.3%). This improvement is particularly significant given that existing interpretable models such as ProtoPNet and  $\mathcal{X}$ -DNNs typically exhibit substantial performance degradation. These results validate that integrating Shapley value estimation during training not only maintains but enhances classification performance, while introducing a minor to no increase in computational complexity.

**Explanation Faithfulness** To evaluate the explanation faithfulness, we employ multiple complementary metrics on ImageNet: AOPC (Area Over the Perturbation Curve) (Chen, Zheng, and Ji 2020), LOdds (Log-odds) (Qiang et al. 2022), and SaCo (Saliency-guided Faithfulness Coefficient) (Wu et al. 2024b). AOPC measures the average drop in model confidence when perturbing the most salient pixels, and a higher AOPC is better. LOdds evaluates if the most important pixels are sufficient for the prediction by measuring the log-odds after perturbation, and a lower LOdds is better. SaCo evaluates faithfulness by correlating pixel saliency scores with their actual impact on the model’s prediction, and a higher SaCo indicates a more faithful explanation. In addition, the number of pixel subsets for SaCo computation is 10.

We compare our method with several SOTA explaining methods (XAI), including post-hoc methods and SENNs. For post-hoc explanation on the vanilla ViT-B backbone, we

Backbone	XAI Method	Pix. Acc.↑	mAP↑	mIoU↑
ViT-B	Grad-CAM	66.83	77.89	45.14
ViT-B	Rollout	58.18	73.63	39.25
ViT-B	GAME	78.17	85.15	60.18
ViT-B	CheferLRP	79.89	85.63	60.97
ViT-B	ATTCAT	41.18	47.89	28.05
ViT-B	ViTShapley	77.47	80.92	58.18
ViT-B	TokenTM	81.43	85.17	64.77
ResNet-50	IG	64.34	68.66	40.78
ResNet-50	Grad-CAM	68.92	79.51	45.34
ResNet-50	FastSHAP	73.92	79.53	56.67
ResNet-50	IDGI	76.42	79.19	58.82
ResNet-50	CAE	61.87	73.73	40.63
ProtoPNet	ProtoPNet	59.76	64.48	40.82
BagNet	BagNet	82.76	86.72	66.77
$\mathcal{X}$ -DNNs	$\mathcal{X}$ -DNNs	64.21	72.94	49.01
B-cos-V	B-cos	75.56	80.27	60.83
B-cos-R	B-cos	77.92	82.16	61.32
HarsanyiNet	HarsanyiNet	83.16	86.97	66.03
B-cosified-V	B-cosified	75.57	80.86	60.43
B-cosified-R	B-cosified	77.91	81.97	61.31
Ours-V	Ours	83.79	86.98	66.38
Ours-R	Ours	<b>85.78</b>	<b>87.67</b>	<b>66.96</b>

Table 3: Quantitative comparison of segmentation performance on the ImageNet-Segmentation dataset. Explanations from various methods are converted into binary segmentation masks and evaluated using Pix. Acc., mAP and mIoU. The upper and center sections present results for applying post-hoc explanation methods to vanilla ViT-B and ResNet-50 backbones, respectively, while the lower section shows results for various SENNs.

evaluate Grad-CAM (Selvaraju et al. 2017), Rollout (Abnar and Zuidema 2020), GAME (Chefer, Gur, and Wolf 2021a), CheferLRP (Chefer, Gur, and Wolf 2021b), ATTCAT (Qiang et al. 2022), ViTShapley (Covert, Kim, and Lee 2023), and TokenTM (Wu et al. 2024a). For vanilla ResNet-50 backbone, we evaluate IG (Sundararajan, Taly, and Yan 2017), Grad-CAM (Selvaraju et al. 2017), FastSHAP (Jethani et al. 2022), IDGI (Yang, Wang, and Bilgic 2023) and CAE (Xie et al. 2024). We also compare our self-explaining framework with representative SENNs, including ProtoPNet, BagNet,  $\mathcal{X}$ -DNN, B-cos Network, HarsanyiNet, and B-cosified Network.

As shown in Table 2, our method achieves superior performance on ImageNet, demonstrating that our explanations faithfully represent the model’s decision-making processes. Our framework’s intrinsic explanations significantly outperform post-hoc methods on ImageNet. Notably, our approach achieves consistent interpretability performance across different backbones. Though our ResNet-50 variant shows marginally better overall performance, the ViT-B implementation achieves competitive results across all metrics. Among existing SENNs, BagNet and HarsanyiNet demonstrate competitive interpretability, but suffer severe degradation of classification performance (Table 1). In contrast, our method maintains high performance in both classifica-

Dataset	Method	T2I		I2T	
		R@1	R@5	R@1	R@5
MS COCO	CLIP	42.2	66.9	58.7	80.7
	B-cosified	41.9	66.7	59.1	81.1
	Ours	42.1	67.3	58.9	80.9
Flickr30k	CLIP	71.2	91.0	85.7	96.7
	B-cosified	70.8	90.7	86.0	97.1
	Ours	70.8	91.1	85.5	97.1

Table 4: Comparison of zero-shot retrieval performance on MS COCO and Flickr30k datasets. We report Recall@1 (R@1) and Recall@5 (R@5) for both text-to-image (T2I) and image-to-text (I2T) retrieval tasks.

Method	Ins. ↑ on T2I		Ins. ↑ on I2T	
	R@1	R@5	R@1	R@5
Grad-CAM	0.1027	0.2216	0.1152	0.2327
Rollout	0.1294	0.2932	0.1753	0.3503
GAME	0.1537	0.3083	0.2097	0.3735
MaskCLIP	0.1423	0.2953	0.1891	0.3514
CLIPsurgery	0.1419	0.2941	0.1771	0.3384
M2IB	0.1469	0.3004	0.2058	0.3691
ECLIP	0.1576	0.3203	0.2056	0.3761
B-cosified	0.1558	0.3089	0.2016	0.3622
Ours	<b>0.1679</b>	<b>0.3360</b>	<b>0.2183</b>	<b>0.3902</b>

Table 5: Evaluation of explanation on MS COCO dataset: Ins. on image retrieval and text retrieval tasks.

tion and explanation, requiring minor architectural modifications.

**Localization Ability** Following previous works (Wu et al. 2024a; Zhao et al. 2024; Chefer, Gur, and Wolf 2021b), we perform segmentation on the ImageNet-Segmentation (Gao et al. 2022) dataset, treating explanation heatmaps as preliminary semantic signals and using a predefined threshold to produce binary segmentation maps. Based on the ground-truth maps, we evaluate the performance on three metrics: pixel-wise accuracy (Pix. Acc.), mean average precision (mAP), and mean intersection over union (mIoU).

The quantitative results, presented in Table 3, unequivocally demonstrate the superiority of our framework in generating spatially precise explanations that are directly applicable to semantic segmentation. Our method consistently outperforms a wide array of both post-hoc and self-explaining methods across two distinct architectures. The superior performance in this segmentation task validates that our method learns explanations that are not only faithful to the model’s classification logic but also possess high spatial fidelity. The learned Shapley values effectively highlight the complete extent of the object, making them a valuable and reliable signal for dense prediction tasks.

### Contrastive Language-Image Pre-Trained Models

To further demonstrate the generalization capabilities and multi-modal applicability of our proposed self-explaining

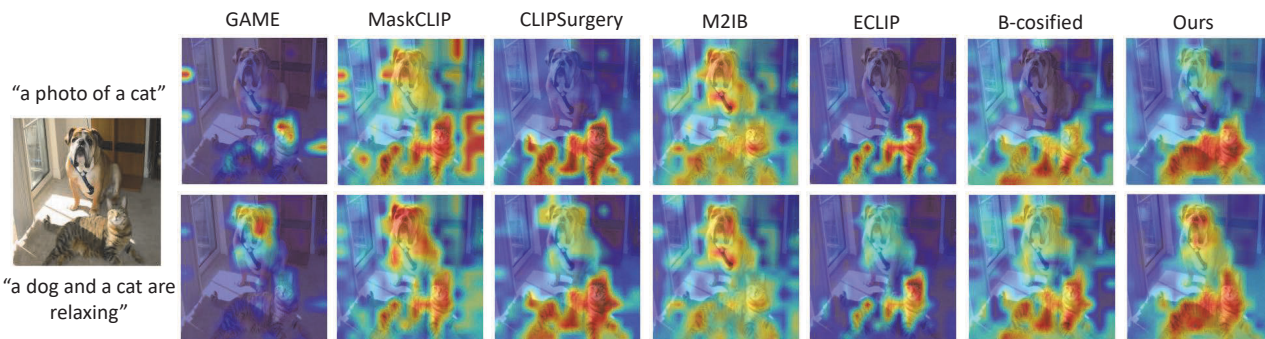


Figure 3: Visualization of different explanation methods for image-text matching. The left column shows input texts matched with the image, and subsequent columns display corresponding explanations generated by various methods.

Method	Del.↓ on T2I		Del.↓ on I2T	
	R@1	R@5	R@1	R@5
Grad-CAM	0.1717	0.3502	0.2161	0.4008
Rollout	0.1948	0.3946	0.2268	0.4238
GAME	0.1706	0.3552	0.1982	0.3800
MaskCLIP	0.1321	0.2841	0.1516	0.2949
CLIPSurgery	0.1794	0.3652	0.2381	0.4292
M2IB	0.1797	0.3671	0.2057	0.3905
ECLIP	0.1246	0.2670	0.1550	0.2933
B-cosified	0.1257	0.2633	0.1632	0.2975
Ours	<b>0.1175</b>	<b>0.2498</b>	<b>0.1401</b>	<b>0.2821</b>

Table 6: Evaluation of explanation on MS COCO dataset: Del. on image retrieval and text retrieval tasks.

framework, we extend its evaluation to the image-text matching task. Following ECLIP (Zhao et al. 2024), we adopt EVA-CLIP (ViT-B/16) (Sun et al. 2023) as the baseline model. Modifications are applied exclusively to the vision encoder to incorporate Shapley value estimation, while the text encoder remains unchanged. Mirroring the implementation in the classification framework in Figure 2, the dual-branch framework processes both unmasked and masked images through separate pathways. Key adaptations include: 1) replace the classification head with a cosine similarity measurement between patch and text embeddings, and 2) use the global sum-pooling to aggregate patch-wise similarities into an image-text similarity, which is capable of the standard CLIP contrastive learning.

Our training objective comprises the standard CLIP contrastive loss and the Shapley loss, enabling the model to generate explanations for image-text alignment through the identification of salient patches corresponding to textual descriptions, while maintaining cross-modal alignment capabilities. For a fair comparison, both our method and B-cosified (the competing self-explaining approach) are fine-tuned on the prevalent large-scale dataset CC3M (Sharma et al. 2018) dataset from EVA-CLIP pre-trained weights. The training details are in the Appendix.

**Prediction Performance** We evaluate zero-shot retrieval performance on MS COCO (Lin et al. 2014) and Flickr30k (Plummer et al. 2015) using Recall@1 (R@1) and Recall@5 (R@5). As presented in Table 4, our method achieves zero-shot retrieval performance that is highly competitive with the vanilla CLIP baseline across both MS COCO and Flickr30k datasets. This confirms that the integration of the Shapley-based optimization preserves primary task capability.

**Explanation Faithfulness** Following ECLIP (Zhao et al. 2024), we evaluate explanation faithfulness using rigorous perturbation tests on MS COCO: Insertion (Ins.) and Deletion (Del.) metrics. Our method is compared against state-of-the-art CLIP explanation approaches, including Grad-CAM, Rollout, GAME, MaskCLIP (Zhou, Loy, and Dai 2022), CLIPSurgery (Li et al. 2023b), M2IB (Wang, Rudner, and Wilson 2023), ECLIP (Zhao et al. 2024), and B-cosified. As Tables 5 and 6 demonstrate, our approach achieves optimal scores on both metrics, confirming that identified pixels are indeed critical to the model’s decision-making process. Visually, Figure 3 illustrates the explanation faithfulness of our method, showing precise alignment between highlighted patches and textual descriptions.

Further experiments, including an ablation study on the trade-off hyperparameter  $\lambda$ , and an analysis of Shapley value estimation errors are presented in the Appendix.

## Conclusion

In this paper, we introduce a novel self-explaining framework that integrates Shapley value estimation into the training process. Extensive experiments demonstrate that our method not only enhances primary task performance but also provides faithful explanations. These results substantiate that integrating Shapley value estimation through multitask learning offers an optimal balance between model performance and interpretability.

## Acknowledgments

This work was supported by the Early Career Scheme (No. CityU 21219323) and the General Research Fund (No. CityU 11220324) of the University Grants Committee

(UGC), the NSFC Young Scientists Fund (No. 9240127), and the Donation for Research Projects (No. 9229164 and No. 9220187).

## References

- Abnar, S.; and Zuidema, W. 2020. Quantifying attention flow in transformers. *arXiv preprint arXiv:2005.00928*.
- Adebayo, J.; Gilmer, J.; Muelly, M.; Goodfellow, I.; Hardt, M.; and Kim, B. 2018. Sanity checks for saliency maps. In *Advances in Neural Information Processing Systems*.
- Alkhatib, A.; Bresson, R.; Boström, H.; and Vazirgiannis, M. 2025. Prediction via Shapley Value Regression. *arXiv preprint arXiv:2505.04775*.
- Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Zitnick, C. L.; and Parikh, D. 2015. VQA: Visual question answering. In *IEEE International Conference on Computer Vision*, 2425–2433.
- Arya, S.; Rao, S.; Böhle, M.; and Schiele, B. 2024. B-cosification: Transforming deep neural networks to be inherently interpretable. In *Advances in Neural Information Processing Systems*, 62756–62786.
- Bass, C.; Da Silva, M.; Sudre, C.; Williams, L. Z.; Sousa, H. S.; Tudosiu, P.-D.; Alfaro-Almagro, F.; Fitzgibbon, S. P.; Glasser, M. F.; Smith, S. M.; et al. 2022. ICAM-reg: Interpretable classification and regression with feature attribution for mapping neurological phenotypes in individual scans. *IEEE Transactions on Medical Imaging*, 42(4): 959–970.
- Böhle, M.; Fritz, M.; and Schiele, B. 2022. B-cos networks: Alignment is all we need for interpretability. In *IEEE Conference on Computer Vision and Pattern Recognition*, 10329–10338.
- Borys, K.; Schmitt, Y. A.; Nauta, M.; Seifert, C.; Krämer, N.; Friedrich, C. M.; and Nensa, F. 2023. Explainable AI in Medical Imaging: An overview for clinical practitioners—Saliency-based XAI approaches. *European Journal of Radiology*, 162(110787).
- Brendel, W.; and Bethge, M. 2019. Approximating CNNs with bag-of-local-features models works surprisingly well on imagenet. *arXiv preprint arXiv:1904.00760*.
- Castro, J.; Gómez, D.; and Tejada, J. 2009. Polynomial calculation of the Shapley value based on sampling. *Computers & Operations Research*, 36(5): 1726–1730.
- Charnes, A.; Golany, B.; Keane, M.; and Rousseau, J. 1988. Extremal principle solutions of games in characteristic function form: core, Chebychev and Shapley value generalizations. In *Econometrics of Planning and Efficiency*, 123–133. Springer.
- Chefer, H.; Gur, S.; and Wolf, L. 2021a. Generic attention-model explainability for interpreting bi-modal and encoder-decoder transformers. In *International Conference on Computer Vision*, 397–406.
- Chefer, H.; Gur, S.; and Wolf, L. 2021b. Transformer interpretability beyond attention visualization. In *IEEE Conference on Computer Vision and Pattern Recognition*, 782–791.
- Chen, C.; Li, O.; Tao, D.; Barnett, A.; Rudin, C.; and Su, J. K. 2019. This looks like that: Deep learning for interpretable image recognition. In *Advances in Neural Information Processing Systems*.
- Chen, H.; Zheng, G.; and Ji, Y. 2020. Generating hierarchical explanations on text classification via feature interaction detection. *arXiv preprint arXiv:2004.02015*.
- Chen, L.; Lou, S.; Zhang, K.; Huang, J.; and Zhang, Q. 2023. Harsanyinet: Computing accurate shapley values in a single forward propagation. *arXiv preprint arXiv:2304.01811*.
- Covert, I.; Kim, C.; and Lee, S.-I. 2023. Learning to estimate shapley values with Vision Transformers. In *International Conference on Learning Representations*.
- Covert, I.; and Lee, S.-I. 2021. Improving kernelSHAP: Practical shapley value estimation using linear regression. In *International Conference on Artificial Intelligence and Statistics*, 3457–3465.
- De Santi, L. A.; Schlötterer, J.; Nauta, M.; Positano, V.; and Seifert, C. 2024. Patch-based Intuitive Multimodal Prototypes Network (PIMPNet) for Alzheimer’s Disease classification. *arXiv preprint arXiv:2407.14277*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*.
- Fong, R.; Patrick, M.; and Vedaldi, A. 2019. Understanding deep networks via extremal perturbations and smooth masks. In *IEEE International Conference on Computer Vision*, 2950–2958.
- Gao, S.; Li, Z.-Y.; Yang, M.-H.; Cheng, M.-M.; Han, J.; and Torr, P. 2022. Large-scale unsupervised semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(6): 7457–7476.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *IEEE Conference on Computer Vision and Pattern Recognition*, 16000–16009.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Hesse, R.; Schaub-Meyer, S.; and Roth, S. 2021. Fast axiomatic attribution for neural networks. *Advances in Neural Information Processing Systems*, 19513–19524.
- Hesse, R.; Schaub-Meyer, S.; and Roth, S. 2024. Benchmarking the Attribution Quality of Vision Models. *arXiv preprint arXiv:2407.11910*.
- Jethani, N.; Sudarshan, M.; Covert, I. C.; Lee, S.-I.; and Ranganath, R. 2022. Fastshap: Real-time shapley value estimation. In *International Conference on Learning Representations*.
- Kindermans, P.-J.; Hooker, S.; Adebayo, J.; Alber, M.; Schütt, K. T.; Dähne, S.; Erhan, D.; and Kim, B. 2019. The (un) reliability of saliency methods. *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, 11700: 267–280.

- Li, X.; Pan, D.; Li, C.; Qiang, Y.; and Zhu, D. 2023a. Negative flux aggregation to estimate feature attributions. *arXiv preprint arXiv:2301.06989*.
- Li, Y.; Wang, H.; Duan, Y.; and Li, X. 2023b. Clip surgery for better explainability with enhancement in open-vocabulary tasks. *arXiv e-prints*, arXiv-2304.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft COCO: Common objects in context. In *European Conference on Computer Vision*, 740–755.
- Lundberg, S. M.; and Lee, S.-I. 2017. A Unified Approach to Interpreting Model Predictions. In *Advances in Neural Information Processing Systems*.
- Mitchell, R.; Cooper, J.; Frank, E.; and Holmes, G. 2022. Sampling permutations for shapley value estimation. *Journal of Machine Learning Research*, 23(43): 1–46.
- Molnar, C. 2020. *Interpretable machine learning*. Lulu.com.
- Nauta, M. 2023. *Explainable AI and Interpretable Computer Vision: From Oversight to Insight*. Phd thesis, University of Twente.
- Nauta, M.; Schlötterer, J.; Van Keulen, M.; and Seifert, C. 2023. Pip-net: Patch-based intuitive prototypes for interpretable image classification. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2744–2753.
- Nauta, M.; Van Bree, R.; and Seifert, C. 2021. Neural prototype trees for interpretable fine-grained image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 14933–14943.
- Petsiuk, V.; Das, A.; and Saenko, K. 2018. Rise: Randomized input sampling for explanation of black-box models. *arXiv preprint arXiv:1806.07421*.
- Plummer, B. A.; Wang, L.; Cervantes, C. M.; Caicedo, J. C.; Hockenmaier, J.; and Lazebnik, S. 2015. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *IEEE International Conference on Computer Vision*, 2641–2649.
- Qiang, Y.; Pan, D.; Li, C.; Li, X.; Jang, R.; and Zhu, D. 2022. Attcat: Explaining transformers via attentive class activation tokens. In *Advances in neural information processing systems*, 5052–5064.
- Redmon, J.; Divvala, S.; Girshick, R.; and Farhadi, A. 2016. You only look once: Unified, real-time object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, 779–788.
- Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. “Why should I trust you?” Explaining the predictions of any classifier. In *ACM International Conference on Knowledge Discovery and Data Mining*, 1135–1144.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Bergs, A. C.; and Li, F.-F. 2015. ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, 115: 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 *IEEE International Conference on Computer Vision*, 618–626.
- Shapley, L. S. 1953. *A Value for n-Person Games*, 307–318. Princeton: Princeton University Press.
- Sharma, P.; Ding, N.; Goodman, S.; and Soricut, R. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Annual Meeting of the Association for Computational Linguistics*, 2556–2565.
- Strumbelj, E.; and Kononenko, I. 2010. An efficient explanation of individual classifications using game theory. *The Journal of Machine Learning Research*, 11: 1–18.
- Sun, Q.; Fang, Y.; Wu, L.; Wang, X.; and Cao, Y. 2023. Evalclip: Improved training techniques for clip at scale. *arXiv preprint arXiv:2303.15389*.
- Sundararajan, M.; Taly, A.; and Yan, Q. 2017. Axiomatic attribution for deep networks. In *International Conference on Machine Learning*, 3319–3328.
- Vinyals, O.; Toshev, A.; Bengio, S.; and Erhan, D. 2015. Show and tell: A neural image caption generator. In *IEEE Conference on Computer Vision and Pattern Recognition*, 3156–3164.
- Wang, R.; Wang, X.; and Inouye, D. I. 2021. Shapley explanation networks. *arXiv preprint arXiv:2104.02297*.
- Wang, Y.; Rudner, T. G.; and Wilson, A. G. 2023. Visual explanations of image-text representations via multi-modal information bottleneck attribution. In *Advances in Neural Information Processing Systems*, 16009–16027.
- Wu, J.; Duan, B.; Kang, W.; Tang, H.; and Yan, Y. 2024a. Token transformation matters: Towards faithful post-hoc explanation for vision transformer. In *IEEE Conference on Computer Vision and Pattern Recognition*, 10926–10935.
- Wu, J.; Kang, W.; Tang, H.; Hong, Y.; and Yan, Y. 2024b. On the faithfulness of vision transformer explanations. In *IEEE Conference on Computer Vision and Pattern Recognition*, 10936–10945.
- Xie, R.; Chen, J.; Jiang, L.; Xiao, R.; Pan, Y.; and Cai, Y. 2024. Accurate explanation model for image classifiers using class association embedding. In *IEEE International Conference on Data Engineering*, 2271–2284.
- Yang, M.; and Kim, B. 2019. Benchmarking attribution methods with relative feature importance. *arXiv preprint arXiv:1907.09701*.
- Yang, R.; Wang, B.; and Bilgic, M. 2023. IDGI: A framework to eliminate explanation noise from integrated gradients. In *IEEE Conference on Computer Vision and Pattern Recognition*, 23725–23734.
- Zhao, C.; Wang, K.; Zeng, X.; Zhao, R.; and Chan, B. A. 2024. Gradient-based Visual Explanation for CLIP. In *International Conference on Machine Learning*.
- Zhou, C.; Loy, C. C.; and Dai, B. 2022. Extract free dense labels from clip. In *European Conference on Computer Vision*, 696–712.