

RESEARCH ARTICLE

T-TAME: Trainable Attention Mechanism for Explaining Convolutional Networks and Vision Transformers

MARIANO V. NTROUGKAS¹, (Graduate Student Member, IEEE),

NIKOLAOS GKALELIS¹, (Member, IEEE),

AND VASILEIOS MEZARIS¹, (Senior Member, IEEE)

Centre for Research and Technology Hellas (CERTH), Information Technologies Institute (ITI), 570 01 Thessaloniki, Greece

Corresponding authors: Mariano V. Ntroukas (ntroukas@iti.gr) and Vasileios Mezaris (bmezaris@iti.gr)

This work was supported by the European Union (EU) Horizon 2020 Program under Grant H2020-101021866 CRITERIA.

ABSTRACT The development and adoption of Vision Transformers and other deep-learning architectures for image classification tasks has been rapid. However, the “black box” nature of neural networks is a barrier to adoption in applications where explainability is essential. While some techniques for generating explanations have been proposed, primarily for Convolutional Neural Networks, adapting such techniques to the new paradigm of Vision Transformers is non-trivial. This paper presents T-TAME, *Transformer-compatible Trainable Attention Mechanism for Explanations* (<https://github.com/IDT-ITI/T-TAME>), a general methodology for explaining deep neural networks used in image classification tasks. The proposed architecture and training technique can be easily applied to any convolutional or Vision Transformer-like neural network, using a streamlined training approach. After training, explanation maps can be computed in a single forward pass; these explanation maps are comparable to or outperform the outputs of computationally expensive perturbation-based explainability techniques, achieving SOTA performance. We apply T-TAME to three popular deep learning classifier architectures, VGG-16, ResNet-50, and ViT-B-16, trained on the ImageNet dataset, and we demonstrate improvements over existing state-of-the-art explainability methods. A detailed analysis of the results and an ablation study provide insights into how the T-TAME design choices affect the quality of the generated explanation maps.

INDEX TERMS CNN, vision transformer, deep learning, explainable AI, model interpretability, attention.

I. INTRODUCTION

Vision Transformers (ViTs) [1] have been found to match or outperform Convolutional Neural Networks (CNNs) in many important visual tasks such as natural image classification [2], classification of masses in breast ultrasound [3], skin cancer classification [4], and face recognition [5]. As a result of the complex multi-layer nonlinear structure and end-to-end learning strategy of models, such as CNNs and ViTs, they typically act as “black box” models that lack transparency [6]. This fact makes it difficult to convince users in critical fields, such as healthcare, law, and governance to

The associate editor coordinating the review of this manuscript and approving it for publication was Zhenhua Guo¹.

trust and employ such systems [7], thus limiting the adoption of Artificial Intelligence [6], [8]. Therefore, it is necessary to develop solutions that address the transparency challenge of deep neural networks.

Explainable artificial intelligence (XAI) is an active research area in the field of machine learning. XAI focuses on developing explainable techniques that help users of AI systems comprehend, trust, and more efficiently manage them [9], [10]. For the image classification task, several explanation approaches have been proposed to tackle the explainability problem for CNN and ViT models [10]. These methods typically produce an explanation map, also referred to as a saliency map, highlighting the salient input features. We must stress that explainability methods should




FIGURE 1. An explanation produced by T-TAME for the ViT-B-16 backbone classifier. The input image (left) belongs to the class “Siamese cat” and is correctly classified by ViT-B-16. The produced explanation (right) highlights the salient features of the image that explain the decision of this specific classifier (areas in red color in the explanation map), which do not necessarily coincide with the image region where the human-recognizable “Siamese cat” object appears. In this example, the explanation map reveals that it is primarily the cat’s head that this classifier relied on to render its decision.

not be confused with approaches targeting weakly supervised learning tasks such as weakly supervised object localization or segmentation [11], which also generate superficially similar heatmaps as an intermediate step. Contrary to the latter, the goal of explainability approaches is to explain the classifier’s decision rather than to locate the region of the target object (for an example, see Fig. 1).

The existing explanation approaches for image classifiers can be roughly categorized as follows. Gradient-based methods, such as Grad-CAM and Grad-CAM++, were pioneering approaches in explaining CNNs [12], [13] and were also among the first methods applied to ViTs [14]. Since these approaches utilize gradient information, they are subject to associated limitations such as gradient saturation and noise issues, resulting in explanations that may include high-frequency variations [15], [16]. Relevance-based methods, on the other hand, use a Taylor decomposition of a relevance function they define to propagate the relevance of pixel information through the examined network [14], [17], [18]. These methods do not directly rely on gradient information and are therefore less prone to the limitations associated with gradient-based approaches; however, their difficulty to be adapted to novel classifier architectures restricts their applicability [19]. Finally, perturbation- [20], [21] and response-based approaches [22], [23], [24], [25] observe the output’s sensitivity to a multitude of small input changes, and combine intermediate network representations to derive an explanation, respectively. The methods within these categories operate without using gradients and thus avoid relevant drawbacks; however, their process for generating an explanation is computationally very expensive.

Distinguished from the above works, L-CAM [26] is a trainable response-based method: it utilizes an appropriate objective function to guide the training of an attention mechanism in order to derive explanation maps of high quality in one forward pass. However, L-CAM uses the feature maps of only the last convolutional layer of the

frozen CNN model to be explained (hereafter also referred to as the “backbone network”), thus may not be able to adequately capture the information used within this backbone for making a classification decision. Additionally, L-CAM is not applicable to ViTs, because ViT feature maps are not three-dimensional, unlike CNN feature maps, and because of the different ways in which ViTs handle input perturbations (see [27] for a comparison w.r.t. robustness between ViTs and CNNs).

To this end, we propose T-TAME: Transformer-compatible Trainable Attention Mechanism for Explanations. T-TAME is inspired by the learning-based paradigm of L-CAM. Unlike L-CAM, T-TAME exploits intermediate feature maps extracted from multiple layers of the backbone network. These features are then used to train a multi-branch hierarchical attention architecture for generating class-specific explanation maps in a single forward pass. Additionally, T-TAME introduces components that manage the compatibility of the trainable attention mechanism with the backbone network, enabling its use with both CNN and ViT backbones. We demonstrate that T-TAME generates higher quality explanation maps over the current SOTA explainability methods, by performing a rich set of qualitative and quantitative comparisons. A preliminary version of this work, still applicable only to CNN backbones, was presented in [28].

In summary, the contributions of this paper are:

- We present the first, to the best of our knowledge, trainable post-hoc method for generating explanation maps for both CNN and Transformer-based image classification networks, which utilizes an attention mechanism to process feature maps from multiple layers.
- We provide a comprehensive evaluation study of the proposed T-TAME method for three heterogeneous backbones: the widely used CNN models VGG-16 [29] and ResNet-50 [30], as well as the breakthrough ViT model ViT-B-16 [1].
- Based on example explanations produced by T-TAME and ablation experiments, we gain insights into the ViT classifier. Specifically, we demonstrate ViT’s global view of input images, thanks to its multi-head attention layer, and we confirm its robustness to out-of-sample distributions of input images.

II. RELATED WORK

We start by briefly discussing the broader domain of XAI. The ability to provide an explanation for why a specific decision was made is now seen as a desirable feature of intelligent systems [31]. These explanations serve to help users understand the AI system’s underlying model, facilitating its effective use and maintenance. Additionally, they assist users in identifying and correcting errors in the AI system’s outputs, thus aiding in debugging. Furthermore, explanations can be used for educational purposes, helping users to explore and understand new concepts within a particular domain. Finally, explanations contribute to users’ trust and cogency

by offering actionable insights and convincing them that the system's decisions can be trusted.

What constitutes a “good” explanation for an AI system is still an open research question. Three important properties for explanations have been identified by social science research on how humans explain their decisions to each other [32]; here we briefly discuss how the current paradigm of explanation methods for vision classifiers aligns with these properties. First, explanations are counterfactual; they justify a decision in opposition to other choices, i.e., why a backbone network classified a specific image as a certain class instead of another possible class. An explanation method can be counterfactual by providing explanation maps for each class that is considered by the backbone network, thus, allowing the user to compare explanation maps for different possible classification decisions. Second, explanations are selected in a biased manner, so as to not overwhelm the user with information. To this end, in the vision classifier domain, the most common form of explanation is a heatmap (a.k.a. explanation map). Third, explanations are social, thus they need to align with the mental model of the user of an AI system. When a user views an image, typically they pay more attention to some parts of the image than to others. In a direct analogy, the user would expect an image classification model to focus more or less on specific regions of the input image for making its classification decision; these are the image regions that are highlighted by the explanation map.

There is a wide range of explainability methods, which are often also referred to as feature attribution methods. Based on the scope of their explanations, i.e., whether they are used to produce explanations for single predictions or for the overall model, these methods can be characterized as local or global [33]. Another important distinction regarding an explainability method arises from its relationship with the model it aims to explain, classifying it as either ante-hoc or post-hoc. The former approaches require architectural modifications that have to be applied prior to the training of the classifier. Several intrinsically explainable classifiers that fall in this category have been developed [34]. Contrarily, a method that can be directly applied to an already-trained classifier is a post-hoc method. Post-hoc explainability approaches can be applied to existing off-the-shelf classifiers, thus providing users with the freedom to choose a top-performing classifier without compromising on model explainability [35]. These approaches can be further categorized as model-specific or model-agnostic, depending on whether they are applicable to only specific models or any type of model. For a more comprehensive review of the different taxonomies of explanation methods and the different approaches therein, the interested reader is referred to [9], [10], [36], and [37].

Among the above-described classes of explainability methods, local post-hoc methods are most widely applicable to the task of explaining deep learning-based image classification models. In the following, we survey the state-of-the-art approaches in this category that are most closely related

to ours. These approaches can be roughly categorized into gradient-, relevance-, perturbation- and response-based. Gradient-based methods [12], [13] compute the gradient of a given input with backpropagation and modify it in various ways to produce an explanation map. Grad-CAM [12], one of the first in this category, uses global average pooling in the gradients of the backbone network's logits with respect to the feature maps to compute weights. The explanation maps are obtained as the weighted combination of feature maps, using the computed weights. Grad-CAM++ [13] similarly uses gradients to generate explanation maps. These methods suffer from the same issues as the gradients they use: neural network gradients can be noisy and suffer from saturation problems for typical activation functions such as ReLU, GELU, and Sigmoid [15].

Relevance-based methods [14], [17], [18] use a Taylor approximation of the gradients to propagate the relevance of pixel information through the examined network. The propagation function is a modified version of backpropagation, aimed at reducing noise and retaining layer-wise salient information. Relevance is propagated to the input image, producing an explanation map. An early work of this class, Deep Taylor Decomposition (DTD) [17], directly uses gradients, propagating them throughout the network and accumulating the contribution to the output prediction from each layer of the network. Layer-wise Relevance Propagation (LRP) [18] cemented the use of Taylor approximation to explain general network architectures. In contrast to methods like Grad-CAM, this method combines information from all of the layers in the network. An extension of the LRP method for Transformer-based architectures, including ViTs, is presented in [14]. However, applying these methods to novel architectures and new network layers is not a straightforward task, requiring the careful fulfillment of the relevance propagation rules through each network operation and dealing with practical issues that may arise, such as numerical instability; thus, their applicability is limited.

Perturbation-based methods [20], [21] attempt to repeatedly alter the input and produce explanations based on the observed changes in the confidence of the original prediction; thus, avoid gradient-related problems such as vanishing or noisy gradients. For instance, RISE [20] utilizes Monte Carlo sampling to generate random masks, which are then used to perturb the input image and generate a respective CNN classification score. Using the computed scores as weights, the explanation map is derived as the weighted combination of the generated random masks. Score-CAM [21], on the other hand, utilizes the feature maps from the final layer of the network as masks by upsampling them to the size of the input image, instead of generating random masks. Thus, RISE and Score-CAM, as most methods in this category, require many forward passes through the network (in the order of hundreds or thousands) to generate an explanation, considerably increasing the inference time and computational cost.

Response-based methods [22], [23], [24], [25], [26] use feature maps or activations of the backbone's layers in the inference stage to interpret the decision-making process of the backbone neural network. One of the first methods in this category, CAM [38], uses the output of the backbone's global average pooling layer as weights, and computes the weighted average of the feature maps at the final convolutional layer. CAM requires the presence of such a global average pooling layer in the target model's architecture, restricting its applicability. SISE [22], and later Ada-SISE [23], aggregate feature maps in a cascading manner to produce explanation maps of any CNN model. Similarly, Poly-CAM [24] uses feature maps from multiple layers, upscales them to the largest spatial dimension present in the set, and then combines them in a cascading manner. Iterated Integrated Attributions (IIA) [25] is a generalization of Integrated Gradients [39] that further employs gradients from internal feature maps. It is also applied to ViT models by using attention matrices as feature maps; the usage of gradients of the input and feature maps from the last two layers before the classification stage is considered. Similarly to perturbation-based methods, the above methods require either multiple forward passes in the case of SISE, Ada-SISE, and Poly-CAM, or multiple backward passes in the case of IIA, to produce an explanation.

Finally, the category of trainable response-based explanation methods is represented by L-CAM [26]. L-CAM mitigates the limitations of response-based methods by using a learned attention mechanism to compute class-specific explanations in one forward pass. However, it can only harness the salient information of feature maps from a single layer of a CNN backbone. The proposed T-TAME method is a trainable response-based method that addresses the limitations of L-CAM, by using feature maps from multiple layers and by being applicable to both CNN and Transformer-based architectures. In contrast to the majority of the approaches described above, which traverse the network multiple times to provide an explanation, the proposed approach is computationally inexpensive at the inference stage, requiring only a single forward pass.

We should also note that the methods of [40] and [41] take a somewhat similar approach to ours in that they produce explanation maps using an attention mechanism and multiple sets of feature maps. However, these methods are ante-hoc, jointly training the attention model with the CNN backbone that learns to perform the desired image classification task. In contrast, T-TAME does not modify the trained target (a.k.a. backbone) model, whose weights remain frozen. I.e., T-TAME is a post-hoc method, exclusively optimizing the attention mechanism in an unsupervised learning manner to generate visual explanations. Thus, no direct comparisons can be drawn with [40] and [41] as they provide explanations for a different, concurrently-trained classifier rather than an already optimized backbone. Finally, as T-TAME is based on an attention mechanism, special tribute must be paid to [42] for the first use of hierarchical attention, inspired by early primate vision, in the field of image processing.

TABLE 1. Main symbols used in Section III.

Symbols in bold denote tensors or sets. Scalars and operators are denoted in normal font.	
Symbols	Description
$f()$	Classifier neural network
\mathbf{I}	Input image
Cls	Scalar specifying the number of classes on which classifier f has been trained on
C, W, H	Dimensions of input image tensor: number of channels, width, height
N, D, P	Parameters of ViT-based backbone: number of patches, hidden size of tokens, width (& height) of a single patch
L_i	i th feature map
L^s	Set of feature maps extracted from s layers
A_i	Attention map of the i th feature branch
A^s	Set of attention maps from s feature branches
C_i, W_i, H_i	Dimensions of the i th feature map (same as the i th attention map): number of channels, width, height
E	Explanation maps (for all Cls classes)
E_n	Explanation map for the n th class
W_e, H_e	Dimensions of explanation map E_n : width, height
Ψ	Set of explanation maps for a subset of the Cls classes

III. METHODOLOGY

A. PROBLEM FORMULATION

Let f be a trained backbone network for which we want to generate explanation maps,

$$f: \text{Sp}(\mathbf{I}) \rightarrow [0, 1]^{\text{Cls}}, \quad (1)$$

where $\text{Sp}(\mathbf{I})$ is the space of three-dimensional input images,

$$\begin{aligned} \text{Sp}(\mathbf{I}) &= \{\mathbf{I} \mid \mathbf{I}: \mathbf{A} \rightarrow \mathbb{R}\}, \\ \mathbf{A} &= \{1, \dots, C\} \times \{1, \dots, W\} \times \{1, \dots, H\}, \end{aligned} \quad (2)$$

$C, W, H \in \mathbb{N}$ are the input image tensor dimensions, i.e., number of channels, width, and height, respectively [20], [22]; and Cls is the number of classes that f has been trained to classify. E.g., for RGB images in the ImageNet dataset, typically $H = W = 224$ is the image height/width, $C = 3$ is the number of channels, Cls = 1000, and the image tensor \mathbf{I} is the mapping from the 3D coordinates to pixel values, commonly in the range $[0, 1]$.

The input image \mathbf{I} is transformed to the output $[0, 1]^{\text{Cls}}$ through various discrete computation steps, called layers. A neural network consists of numerous layers, depending on its specific architecture; a layer's output is referred to as a "feature map". Suppose feature maps are extracted from s layers of the backbone network f ; this set of feature maps is represented as

$$L^s = \{L_i \mid i \in \{1, \dots, s\}\}. \quad (3)$$

A feature map L_i of a neural network can take different shapes depending on the type of the backbone network. For CNNs, a feature map is typically represented as

$$L_i: \{1, \dots, C_i\} \times \{1, \dots, W_i\} \times \{1, \dots, H_i\} \rightarrow \mathbb{R}, \quad (4)$$

where, $C_i, W_i, H_i \in \mathbb{N}$ are the respective channel, width, and height dimensions of the i th feature map in the feature map set. In ViTs [1], the feature map is represented as

$$L_i: \{1, \dots, N + 1\} \times \{1, \dots, D\} \rightarrow \mathbb{R}, \quad (5)$$

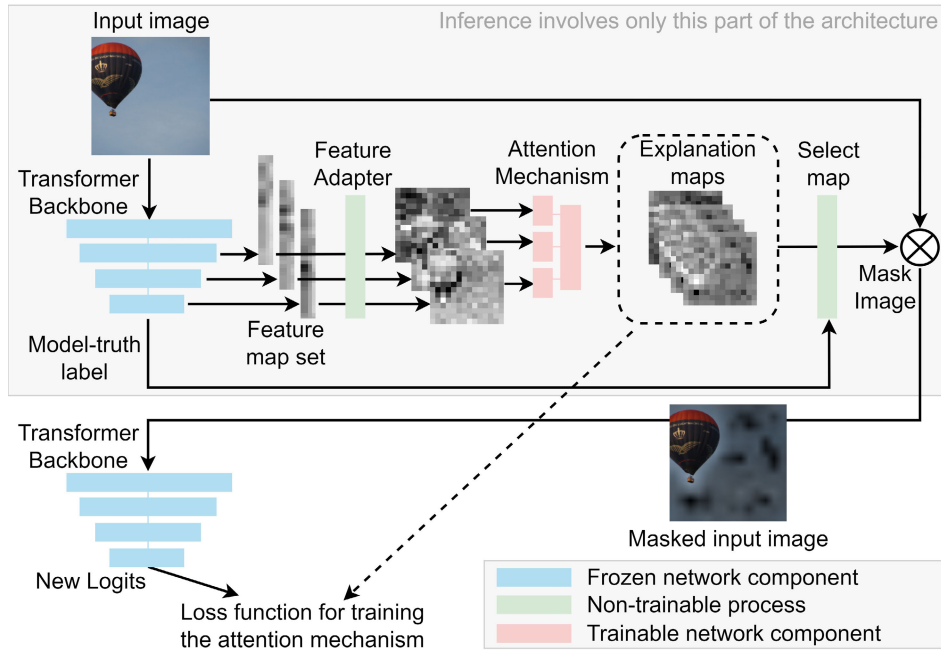


FIGURE 2. Overview of the T-TAME method, showing both the overall architecture used for training the explanation-generating attention mechanism and the inference-stage use of the trained attention mechanism. In this illustration, T-TAME is applied on a ViT backbone.

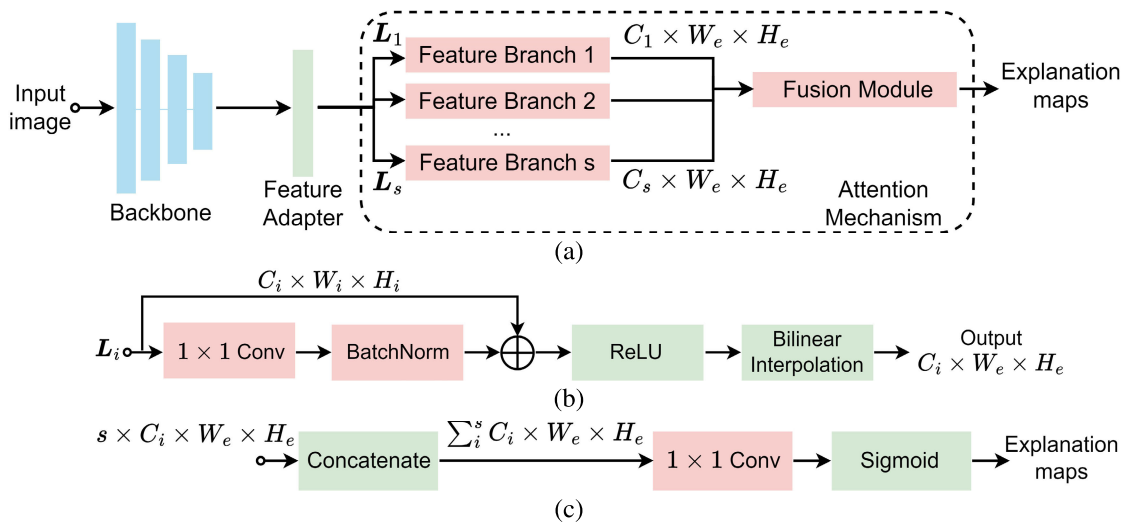


FIGURE 3. Structure of the core architecture of the proposed T-TAME method: (a) Overall structure (feature map adapter and attention mechanism), (b) detailed structure of a feature branch of the attention mechanism, (c) detailed structure of the fusion module of the attention mechanism. Color coding retains the same meaning as in Fig. 2.

where $N, D \in \mathbb{N}$ are the number of patches and the common hidden size through all its layers, respectively. The former (N) equals HW/P^2 , where $P \in \mathbb{N}$ is the width (& height) of a single square patch of the input image. P, N and D are architecture-dependent values. For instance, for the ViT-B-16 architecture and input image resolution $W = H = 224, P = 16, N = 14^2 = 196$ and $D = 768$. The extra token in the ViT feature map (i.e., the one that increases the map’s dimension from N to $N + 1$) is called the “class token” and is used by the classification layer. Thus, in CNNs, feature maps are 3D tensors, while in ViTs they are 2D tensors.

Assume an attention mechanism defined as

$$\text{AM: Sp}(L^s) \rightarrow \text{Sp}(E), \tag{6}$$

where

$$E: \{1, \dots, \text{Cls}\} \times \{1, \dots, W_e\} \times \{1, \dots, H_e\} \rightarrow [0, 1] \tag{7}$$

are the explanation maps produced by the attention mechanism, having spatial dimensions W_e, H_e . $\text{Sp}(L^s)$ and $\text{Sp}(E)$ denote the space of feature map sets and explanation maps, respectively. The explanations are class-discriminative, i.e., each slice of E along its first dimension, E_n ,

$n \in \{1, \dots, \text{Cls}\}$, is the explanation map corresponding to the n th class on which the classifier f has been trained on.

Given the above general formulation, we propose T-TAME: a trainable attention mechanism architecture, along with a compatible training method. The proposed attention mechanism is applicable to a wide range of classifier backbones, i.e., vastly different CNNs and ViTs. An overview of the T-TAME method is given in Fig. 2.

B. T-TAME OVERALL ARCHITECTURE

The T-TAME method, as illustrated in Fig. 3(a), is composed of the following components:

- A feature map adapter
- The feature branches of the Attention Mechanism
- The fusion module of the Attention Mechanism

These components are trained, as illustrated in Fig. 2, using a suitable loss function, together with a mask selection and an image masking procedure.

The feature map adapter reshapes the feature map set output by the backbone network so that it can be input to the attention mechanism, which consists of the feature branches and the fusion module. Each feature branch has a one-to-one mapping with each feature map in the feature map set and processes them separately. The fusion module combines the attention maps from each feature branch into the final class-discriminate explanation maps. Specific masks are then selected, in an unsupervised manner, and used to mask the image. The loss function takes as input a subset of the produced explanation maps, i.e., a number of slices along the channel dimension, and the logits generated by passing the masked image through the backbone network. In the next section, we specify each of these components.

C. T-TAME ARCHITECTURE COMPONENTS

1) ATTENTION MECHANISM

For a feature map set L^s , the attention mechanism consists of s feature branches and the fusion module. The feature branch structure consists of a 1×1 convolution layer with the same number of input and output channels, a batch normalization layer, a skip connection, and a ReLU activation, as illustrated in Fig. 3(b). Each feature branch

$$FB : \text{Sp}(L_i) \rightarrow \text{Sp}(A_i) \tag{8}$$

takes as input a single CNN-type feature map L_i (as defined in Eq. (4)) and outputs an attention map

$$A_i : \{1, \dots, C_i\} \times \{1, \dots, W_e\} \times \{1, \dots, H_e\} \rightarrow \mathbb{R}, \tag{9}$$

where $W_e = \max_i W_i$ and $H_e = \max_i H_i$. That is, the attention map A_i has the same channel dimension as L_i , and the same spatial dimensions as the explanation maps E (Eq. (7)). The dimensions W_e, H_e are equal to the spatial dimensions of the largest input feature map. This is achieved by applying bilinear interpolation¹ where necessary (Fig. 3(b)), i.e., on

¹Bilinear interpolation is chosen because it produces smoother explanation maps than simple nearest-neighbor interpolation, while still being rather computationally inexpensive.

the feature branches whose input feature map dimensions are smaller than W_e and H_e . The resulting attention maps $A^s = \{A_i \mid i \in \{1, \dots, s\}\}$ are forwarded into the fusion module

$$FS : \text{Sp}(A^s) \rightarrow \text{Sp}(E), \tag{10}$$

consisting of a concatenation operator, a 1×1 convolutional layer, and a sigmoid activation, as illustrated in Fig. 3(c). Specifically, the attention maps are initially concatenated into a single attention map (a 3D tensor with $\sum_{i=1}^s C_i$ channels, each channel of spatial dimensions W_e, H_e), and then processed (by the convolution and sigmoid layers) to generate the explanation map.

2) FEATURE MAP ADAPTER

In the context of a CNN backbone network, the feature maps inherently conform to the required input shape (C_i, W_i, H_i) (as seen in Fig. 3(b)), thus there is no need to adapt the feature maps to the attention mechanism. In this case, the feature map adapter is the identity function $a(L_i) = L_i$. When the backbone network is Transformer-based, as in the case of ViTs, the feature maps are defined as in Eq. (5). The feature map adapter first excludes the class token, as it lacks spatial information, and then reshapes the feature map into a 3D format that mirrors the structure of feature maps typically found in a CNN backbone, as defined in Eq. (4), where $C_i = D, W_i = H_i = \sqrt{N}$. This is essentially the inverse of the ViT architecture input processing step.² The reshaping step is visually presented in Fig 4.




FIGURE 4. Illustration of tensor reshaping in the feature map adapter: Converting a (4, 2) tensor to a (2, 2, 2) tensor.

3) LOSS FUNCTION, MASK SELECTION, AND MASKING METHOD

The loss function used for training the proposed attention mechanism is the weighted sum of two loss functions,

$$\text{Loss}(\Psi, \text{logits}, y) = \lambda_1 \text{CE}(\text{logits}, y) + \lambda_2 \text{TV}'(\Psi), \tag{11}$$

where $\text{CE}(), \text{TV}'()$ are the cross-entropy and modified total variation loss, respectively; λ_1, λ_2 are the corresponding

²In ViT, the feature map produced by the initial convolution layer, with dimensions (D, \sqrt{N}, \sqrt{N}) is initially reshaped into a 2D format with dimensions (D, N) . Then, the order of dimensions is permuted, i.e., the dimensions become (N, D) and the class token is introduced, resulting in a feature map with dimensions $(N + 1, D)$.

summation weights; and y is the predicted class of the backbone network (a.k.a. model truth): $y = \arg \max f(\mathbf{I})$. Ψ is defined as

$$\Psi: \{E_n \mid n \in \text{Cls}_\Psi \subset \{1, \dots, \text{Cls}\}\}, \quad (12)$$

i.e., Ψ is a set containing any number of explanation maps E_n . For each input image, in a batched training scenario with batch size B , we include in Ψ the explanation map corresponding to the predicted class y of the backbone network, E_y , and additional $B - 1$ explanation maps for randomly selected classes. The incorporation of explanation maps corresponding to other classes besides the predicted class in the loss function helps the attention mechanism to learn to generate class-discriminative explanation maps.

The cross-entropy loss uses the logits generated by the backbone network for the masked input image and the predicted class to compute a loss value. This term trains the attention mechanism to focus on salient, class-relevant parts of the input image. The masking procedure involves taking the element-wise product (also known as the Hadamard product), denoted as \odot , between the raw image and the mask of the predicted class using,

$$\text{CNN Masking}(E_y, \mathbf{I}) = |E_y \odot \mathbf{I}|, \quad (13)$$

$$\text{ViT Masking}(E_y, \mathbf{I}) = E_y \odot |\mathbf{I}|, \quad (14)$$

where $||$ denotes element-wise standardization (also known as Z-score normalization) using the dataset mean and standard deviation [43]. This operation shifts and scales each element of the input tensor based on the mean and standard deviation of the dataset. We should note that masking removes features from the input image and renders it out-of-distribution [44]. CNNs are more sensitive to such a transformation in comparison to Transformer-like architectures, as shown in [27]. To this end, in the case of CNN, the explanation map is first used as a mask to perturb the input image and then the standardization is applied (Eq. (13)). This is the typical order of applying a perturbation (e.g. masking, augmentation, multiplicative/additive noise) in an input image, with the aim of causing a minimal shift to the input data distribution [45]. On the other hand, in the case of ViT backbones, the image \mathbf{I} is first standardized, and then used in the Hadamard product (Eq. (14)). This different approach is shown to perform better (see Table 7 in the Experiments section), and is motivated by considering what happens when standardizing only the input image: the explanation map, when used as a mask, behaves as a local perturbation, i.e., certain regions of the input image remain intact while the global statistics of the image change. Since ViT-like models [1], [44], [46] focus on certain image sub-regions and also examine global information, this type of perturbation is beneficial [27], [47].

The modified total variation loss, inspired by total variation denoising [48], is the sum of the squares of the total variation norm of the explanation maps Ψ and the mean of element-wise exponentiation of the explanation maps. This term reduces noise and overactivation in the generated

explanation maps. The modified total variation loss is defined as,

$$\text{TV}'(\Psi) = E(\Psi) + \lambda_3 V(\Psi), \quad (15)$$

with $E()$ defined as

$$E(\Psi) = \frac{1}{S} \sum_{n,j,k} E_{n,j,k}^{\lambda_4}, \quad E_n \in \Psi, \quad (16)$$

and $V()$ defined as

$$V(\Psi) = \frac{1}{2S} \sum_{n,j,k} (|E_{n,j+1,k} - E_{n,j,k}|^2 + |E_{n,j,k+1} - E_{n,j,k}|^2), \quad E_n \in \Psi, \quad (17)$$

where $E_{n,j,k}$ denotes the value of the explanation map E_n in indices (j, k) and $S = B \cdot W_e \cdot H_e$ is the number of such values included in the summation of Eq. (16). $\text{TV}'(\Psi)$ forces the attention mechanism to output less noisy explanation maps that emphasize smaller and more focused regions in the input image instead of arbitrarily large areas. Without this term in the loss function, the trivial solution for minimizing the cross-entropy loss would be not masking the input image at all, with a homogeneous and appropriately scaled explanation map. The scalars λ_3 and λ_4 are additional hyperparameters of the loss function. By modifying the original total variation loss with the addition of these hyperparameters, we gain an additional degree of freedom to generate smoother and more focused explanation maps.

D. TRAINING AND INFERENCE

During the training of T-TAME, batches from the dataset that was used to originally train the backbone network are used to generate feature map sets and logits. The feature maps are then input to the attention mechanism to produce explanation maps. Using the predicted classes from the backbone's logits, specific explanation maps are selected and used to mask the input images. The batch of masked images is input to the backbone to produce new logits. The new logits and a subset of explanation maps corresponding to the predicted classes, as well as other random classes, are input to the loss function. Through backpropagation, the weights of the attention mechanism are optimized to produce more salient explanation maps.

During inference, only the upper half of the architecture illustrated in Fig. 3 is used: as typically done for classifying an input image, the image is input to the backbone classifier to generate a decision and, as an intermediate result of this process, a feature map set. Then, the produced feature map set is input to the trained attention mechanism for generating explanation maps for all classes of the backbone classifier.

We should clarify here that, at the inference stage, the sigmoid activation function of Fig. 3(c) is replaced by a min-max scaling step. This is done to produce a heatmap in the $[0, 1]$ range, for a fair comparison with all of the examined explainability methods that typically introduce such a scaling step, e.g., [12], [13], and [21]. Contrarily, the sigmoid function illustrated in Fig. 3(c) is used during training,

because the gradient of the min-max scaling operation is very noisy, impeding the training of the explanation mechanism.

IV. EXPERIMENTS

A. DATASETS AND BACKBONE NETWORKS

We choose three neural network models that are widely used for image classification, as the backbones for which we will generate explanations using T-TAME: VGG-16 [29], ResNet-50 [30] and ViT-B-16 [1]. This choice is further motivated by the diversity among these models: there are significant differences between the two chosen CNN architectures, and between them and the ViT architecture. All 3 backbones have been trained on the ImageNet dataset [49]; we obtain the trained models from the `torchvision.models` library.

For training and evaluating T-TAME on each of these backbones, we use the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) 2012 dataset [49] (i.e., the same dataset that the backbones have been trained on). This dataset contains 1000 classes, 1.3 million, and 50k images for training and evaluation, respectively. Out of the 50k evaluation images, we use a set of 2000 randomly selected images as the validation set and a different, disjoint set of 2000 randomly selected images for testing the explainability results (the same as in [26], [28] to allow for a fair comparison). The validation set is utilized for optimizing the T-TAME training hyperparameters, including the hyperparameters of the loss function: λ_1 , λ_2 , λ_3 , and λ_4 , as well as the number of training epochs and learning rate. Testing on 2000 images is chosen not only for consistency with [26] and [28] but additionally because executing the perturbation-based approaches that we use in the experimental comparisons is computationally expensive [20], [21] (up to almost four orders of magnitude more expensive than T-TAME and gradient-based methods).

B. EVALUATION MEASURES

For quantitative evaluation and comparisons, we employ two widely used evaluation measures, Increase in Confidence (IC) and Average Drop (AD) [13]. Additionally, we employ the promising Noisy Imputation method from the Remove and Debias (ROAD) evaluation framework recently introduced in [50]. For completeness, we briefly describe these two evaluation approaches in the following.

1) IC AND AD

These two measures are defined as follows:

$$\text{AD}(v) = \sum_{v=1}^{\Upsilon} \frac{\max\{0, \psi_v - \psi_v^{\phi_v}\}}{\Upsilon \psi_v} \cdot 100, \quad (18)$$

$$\text{IC}(v) = \sum_{v=1}^{\Upsilon} \frac{\text{int}(\psi_v^{\phi_v} > \psi_v)}{\Upsilon} \cdot 100, \quad (19)$$

where Υ represents the number of test images; $y_v = \arg \max f(\mathbf{I}_v)$ is the model-truth label for the v th test image \mathbf{I}_v ; and $\psi_v = \max f(\mathbf{I}_v)$ is the classifier's output score (confidence) for the model-truth class. $\psi_v^{\phi_v}$ is the classifier's output score for the model-truth class when input to the

classifier is a modified image, i.e. one that is masked according to the explanation map for the same class, \mathbf{E}_{y_v} (generated by the explainability method under evaluation). That is,

$$\psi_v^{\phi_v} = \mathbf{e}_{y_v} \cdot f(\mathbf{I}_v \odot \phi_v(\mathbf{E}_{y_v})), \quad (20)$$

$$\mathbf{e}_{y_v} = (0, \dots, 1 \text{ at position } y_v, \dots, 0), \quad (21)$$

where $\phi_v()$ represents a threshold function to select the top $v\%$ higher-valued pixels of the explanation map \mathbf{E}_{y_v} , and $\text{int}()$ returns 1 when the input condition is satisfied and 0 otherwise.

Intuitively, AD measures how much, on average, the produced explanation maps, when used to mask the input images, reduce the confidence of the model. The implicit assumption is that by masking the input image using the explanation, confusing and irrelevant background information is removed, and thus, the average drop in confidence should be minimized. In contrast, IC measures how often explanation maps, when applied in the same manner, increase the model's confidence. By eliminating confounding background information, the classification confidence likely will increase, hence IC should be maximized. A naive all-ones mask would result in a 0% AD, the optimal result, and 0% IC, the worst result. Therefore, for a more comprehensive evaluation, we use the combination of these measures. Furthermore, since the explanation maps produced by each method vary in their intensity of activation, we also apply a threshold $v\%$ to the explanation maps, as discussed above, to assess how effectively the pixels are ordered based on importance. Using a smaller threshold (e.g. $v = 15\%$) creates a more challenging evaluation setup since a smaller percentage of the image pixels is retained. This way, we can compare methods more fairly, since methods that produce highly activated explanation maps could initially generate good results without thresholding, but when a threshold is applied they may struggle, revealing a subpar ordering of pixel importance in the explanation map. This evaluation protocol has been adopted in most previous works, including [13], [20], [21], [25], [26], [51], and [14].

2) ROAD

The Remove and Debias evaluation framework [50] aims to improve the process of assessing the quality of explanation maps of different explainability techniques with pixel perturbations. The authors of [50] first prove, using an information theory analysis, that simpler methods of removing areas of an image using a binary mask leak information about the shape of the mask. The shape of the mask could reveal class information. Thus, the ROAD framework aims to remove salient information rather than simply removing salient pixels. An example of the effect of different imputation methods is shown in Fig. 5. In this example, we observe that by imputing the images in a straightforward way (Fig. 5(c)), i.e., replacing the removed pixels with the mean of the original image, the region of the modification of the original image is evident. This can leak information about the class

contained in the image. Resolving the discrepancy between the removal of pixels and the removal of the information contained in the removed pixels is the aim of the noisy imputation method of ROAD. In Fig. 5(d), where the noisy imputation method is employed, we observe that it is now much harder to detect which pixels were removed, reducing the leakage of class information contained in the binary mask.

Two evaluation measures are defined in ROAD, namely, MoRF (Most Relevant First) and LeRF (Least Relevant First). In the former/latter, a binary mask generated from the explanation map is used that highlights the $\nu\%$ most/least important regions in the image. This binary mask is then utilized to impute the input image, and the confidence in the target class is calculated. The ROAD score is then computed using,

$$\text{MoRF}(\nu) = \sum_{v=1}^{\Upsilon} \frac{\psi_v^{\hat{\theta}_v}}{\Upsilon} \cdot 100, \quad (22)$$

$$\text{LeRF}(\nu) = \sum_{v=1}^{\Upsilon} \frac{\psi_v^{\check{\theta}_v}}{\Upsilon} \cdot 100, \quad (23)$$

where $\psi_v^{\hat{\theta}_v} = \mathbf{e}_{y_v} \cdot f(\hat{\theta}_v(\mathbf{I}_v, \mathbf{E}_{y_v}))$, $\psi_v^{\check{\theta}_v} = \mathbf{e}_{y_v} \cdot f(\check{\theta}_v(\mathbf{I}_v, \mathbf{E}_{y_v}))$, \mathbf{e}_{y_v} is defined in Eq. (21), and $\hat{\theta}_v(\cdot)$, $\check{\theta}_v(\cdot)$ represent the ROAD imputation operation applied to $\nu\%$ of the most or least important pixels of the input image, respectively. In the case of MoRF, a sharp decline in model confidence should be observed, as the removal of important class information should rapidly deteriorate the model's performance. In the case of LeRF, removing irrelevant information should minimally affect the confidence of the model. We compute the ROAD measures only when comparing with other methods (i.e., in Section IV-D), as ROAD is significantly more computationally expensive than computing the AD and IC measures. This new evaluation protocol has been adopted in the very recent works [52], [53].

C. EXPERIMENTAL SETUP

Feature maps from three layers are extracted from each backbone to which T-TAME is applied (i.e., $s = 3$). The VGG-16 backbone model consists of five blocks of convolutions separated by 2×2 max-pooling operations, as shown in Fig. 6. We choose one layer from each of the last three blocks, namely the feature maps output by the max-pooling layers of each block. Alternatively, we also experiment with the use of feature maps output by the last convolution layer of each block. The results of this alternate choice of feature maps are discussed in Section IV-E1. ResNet-50 consists of five stages, as depicted in Fig. 7. For this backbone, we utilize the feature maps from its final three stages. Finally, for the ViT-B-16 backbone, which consists of eleven encoder blocks, we use the feature maps of the last three encoder blocks, as shown in Fig. 8.

T-TAME is trained using the loss function defined in Eq. (11) with the SGD (Stochastic Gradient Descent) algorithm. The OneCycleLR policy [54] was utilized to vary the learning rate during the training procedure. The largest batch size that can fit in the employed GPU's memory is used,

as recommended in [55]. The rest of the hyperparameters were identified using the validation dataset and the IC(15%) and AD(15%) measures. IC and AD were preferred over ROAD because they are simpler to interpret and much less computationally expensive; and, we opted for IC and AD at the $\nu = 15\%$ threshold because they are the most challenging ones to improve upon and provide more focused explanation maps. To this end, the optimal hyperparameters of the loss function (Eq. (11), Eq. (16), Eq. (15)) were empirically identified as: $\lambda_1 = 1.5$, $\lambda_2 = 2$, $\lambda_3 = 0.005$, $\lambda_4 = 0.3$. We observed surprising robustness across the different architectures using the above set of hyperparameters. Thus, the hyperparameter values do not vary between backbones. The maximum learning rate in the OneCycleLR policy was optimized using a grid search. Finally, the number of epochs was identified by varying it from one to eight and selecting the optimal one.

During training, the same image preprocessing employed in the original backbone network [1], [29], [30] is used, i.e., the smallest spatial dimension of each image is resized to 256 pixels, the image is then random-cropped to dimensions $W = H = 224$, and standardized using the channel-wise statistics calculated on the ImageNet dataset (mean = [0.485, 0.456, 0.406], std = [0.229, 0.224, 0.225]). During the testing phase, the image is again resized so that the smallest spatial dimension becomes 256 pixels, however, center-cropping is used instead of random-cropping, again as in [1], [29], and [30]. Subsequently, the appropriate masking procedure is selected, depending on the type of backbone network, as discussed in Section III-C3. This protocol is used unaltered for every considered explainability method, to ensure a fair comparison. Feature maps are extracted from the backbone networks using the `torchvision.models.feature_extraction` library.

D. QUANTITATIVE ANALYSIS

The following state-of-the-art methods are quantitatively compared with the proposed T-TAME, on all three considered backbones, using the evaluation measures described in Section IV-B: Grad-CAM [12], Grad-CAM++ [13], Score-CAM [21], Ablation-CAM [51], RISE [20] and Iterated Integrated Attributions (IIA) [25]. Additionally, we compare with L-CAM-Img [26] and Transformer Layer-wise Relevance Propagation (LRP) [14], only on CNN and Transformer backbones, respectively (because L-CAM-Img and Transformer LRP are only applicable to these specific backbones). These methods are selected because they are among the top-performing methods in the visual XAI domain and their source code is publicly available. For Grad-CAM [12], Grad-CAM++ [13], Score-CAM [21], Ablation-CAM [51], we use the implementations of the `pytorch_gradcam` library [56]. For RISE [20] and Iterated Integrated Attributions (IIA) [25] we use the original implementations available at <https://github.com/eclique/RISE> and <https://github.com/iaa-iccv23/iaa>, respectively. For L-CAM-Img [26], which is only




FIGURE 5. In this synthetic example, a typical imputation approach is compared to the noisy imputation method of the ROAD framework. In the naive case, information about the mask's shape is clearly leaked. Fig 5(d) shows how ROAD removes pixels in a more nuanced way to avoid revealing the shape of the binary mask.

TABLE 2. Comparison of T-TAME with other methods using the AD and IC measures (CNN backbones).

Backbone	Measure	Grad-CAM [12]	Grad-CAM++ [13]	Score-CAM [21]	Ablation-CAM [51]	RISE [20]	IJA [25]	L-CAM-Img [26]	T-TAME
VGG-16	AD↓(100%)	32.12%	30.75%	27.75%	34.87%	8.74%	25.42%	12.15%	9.33%
	IC↑(100%)	22.10%	22.05%	22.80%	19.25%	51.30%	22.90%	40.75%	<u>50.00%</u>
	AD↓(50%)	58.65%	54.11%	45.60%	49.23%	42.42%	64.43%	37.37%	36.50%
	IC↑(50%)	9.50%	11.15%	14.10%	11.45%	17.55%	5.60%	<u>20.25%</u>	22.45%
	AD↓(15%)	84.15%	82.72%	75.70%	76.96%	78.70%	87.68%	74.23%	73.29%
	IC↑(15%)	2.20%	3.15%	4.30%	36.50%	<u>4.45%</u>	1.45%	<u>4.45%</u>	5.60%
	FP↓	1	1	512	2048	4000	<u>12</u>	1	1
	AD↓(100%)	13.61%	13.63%	11.01%	13.58%	11.12%	21.58%	11.09%	7.81%
IC↑(100%)	38.10%	37.95%	<u>39.55%</u>	37.05%	<u>46.15%</u>	26.90%	43.75%	54.00%	
ResNet-50	AD↓(50%)	29.28%	30.37%	26.80%	30.33%	36.31%	41.31%	29.12%	27.88%
	IC↑(50%)	23.05%	23.45%	24.75%	22.30%	21.55%	14.95%	24.10%	<u>27.50%</u>
	AD↓(15%)	<u>78.61%</u>	79.58%	78.72%	79.62%	82.05%	87.84%	79.41%	78.58%
	IC↑(15%)	<u>3.40%</u>	3.40%	3.60%	3.50%	3.20%	1.45%	<u>3.90%</u>	4.90%
	FP↓	1	1	2048	8192	8000	<u>12</u>	1	1

TABLE 3. Comparison of T-TAME with other methods using the AD and IC measures (ViT backbone).

Backbone	Measure	Grad-CAM [12]	Grad-CAM++ [13]	Score-CAM [21]	Ablation-CAM [51]	RISE [20]	Transformer LRP [14]	IJA [25]	T-TAME
ViT-B-16	AD↓(100%)	37.19%	57.21%	38.09%	35.50%	13.52	67.07%	56.54%	8.19%
	IC↑(100%)	12.75%	5.55%	15.35%	8.90%	<u>37.00</u>	2.25%	8.05%	38.35%
	AD↓(50%)	40.74%	72.77%	44.20%	42.16%	31.94	63.19%	64.01%	23.64%
	IC↑(50%)	12.30%	4.85%	14.50%	10.55%	<u>33.10</u>	3.90%	7.65%	40.40%
	AD↓(15%)	<u>73.11%</u>	92.51%	77.50%	80.86%	79.56	88.68%	86.44%	72.89%
	IC↑(15%)	<u>5.40%</u>	0.80%	4.85%	2.95%	<u>6.85</u>	0.90%	4.00%	9.40%
	FP↓	1	1	768	768	8000	1	<u>44</u>	1

applicable to CNN backbones, we use the original implementation, available at <https://github.com/bmezaris/L-CAM>. Finally, for the Transformer Layer-wise Relevance Propagation (LRP) method [14], which is only applicable to ViT backbones, we use the original implementation, available at <https://github.com/hila-chefer/Transformer-Explainability>.

The results in terms of the AD(v) and IC(v) measures with $v = 15\%$, 50% , 100% for CNN and ViT models are shown in Tables 2 and 3. The respective results for the MoRF(v) and LeRF(v) measures are shown in Figs. 9 to 14, where v varies from 10% to 90%. In order to acquire a single value for each ROAD measure, model, and examined explainability

TABLE 4. Comparison of T-TAME with other methods using the ROAD measures (CNN backbones).

Backbone	Measure	Grad-CAM [12]	Grad-CAM++ [13]	Score-CAM [21]	Ablation-CAM [51]	RISE [20]	IIA [25]	L-CAM-Img [26]	T-TAME
VGG-16	MoRF↓	21.40%	23.14%	22.54%	20.51%	22.78%	32.07%	19.32%	17.74%
	LeRF↑	71.22%	72.26%	73.46%	73.47%	76.08%	64.21%	<u>70.75%</u>	<u>73.50%</u>
	FP↓	1	1	512	2048	4000	<u>12</u>	1	1
ResNet-50	MoRF↓	25.95%	27.15%	28.70%	27.20%	23.80%	32.46%	24.62%	25.83%
	LeRF↑	<u>80.35%</u>	79.22%	78.50%	79.16%	80.58%	73.95%	<u>76.69%</u>	75.35%
	FP↓	1	1	2048	8192	8000	<u>12</u>	1	1

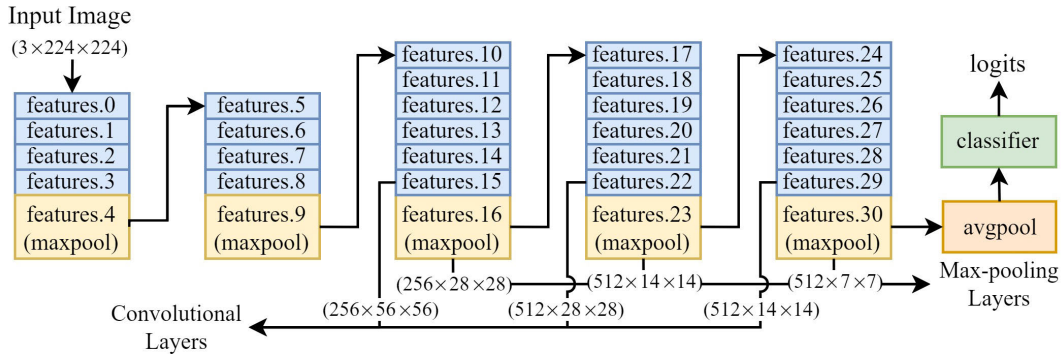


FIGURE 6. The layers from which feature maps are extracted when applying T-TAME to a VGG-16 backbone. We also indicate in this diagram the dimensions of the extracted feature maps. We experiment with two separate sets of layers in the ablation study (Table 6), where we denote by “Max-pooling Layers” the last three max-pooling layers, and by “Convolutional Layers” the three layers before the last three max-pooling layers. We use the same layer naming as the torchvision.models.feature_extraction library.

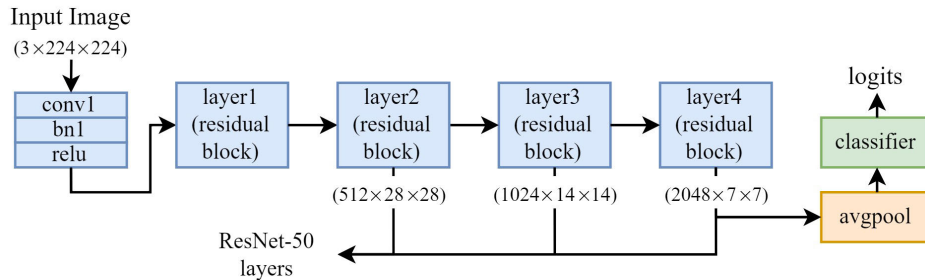


FIGURE 7. The layers from which feature maps are extracted when applying T-TAME to a ResNet-50 backbone. We also indicate in this diagram the dimensions of the extracted feature maps. The outputs of the final three residual blocks are used. We use the same layer naming as the torchvision.models.feature_extraction library.




FIGURE 8. The layers from which feature maps are extracted when applying T-TAME to a ViT-B-16 backbone. We also indicate in this diagram the dimensions of the extracted feature maps. The outputs of the final three encoder blocks are used. We use the same layer naming as the torchvision.models.feature_extraction library.

method, we also compute the average confidence score across all percentages v described above; these results are presented in Tables 4 and 5.

In all tables, for each comparison (i.e., each row), the best and second-best results are shown in bold and underline. From the obtained results, we observe the following:

TABLE 5. Comparison of T-TAME with other methods using the ROAD measures (ViT backbone).

Backbone	Measure	Grad-CAM [12]	Grad-CAM++ [13]	Score-CAM [21]	Ablation-CAM [51]	RISE [20]	Transformer LRP [14]	IIA [25]	T-TAME
ViT-B-16	MoRF↓	29.29%	51.74%	35.00%	36.68%	38.77%	29.28%	42.37%	26.48%
	LeRF↑	78.49%	69.64%	67.95%	78.51%	82.81%	81.52%	73.97%	81.60%
	FP↓	1	1	768	768	8000	1	<u>44</u>	1

TABLE 6. Ablation study: different architectural choices of the attention mechanism of T-TAME.

Model	Feature Extraction	Architecture Variant	AD↓ (100%)	IC↑ (100%)	AD↓ (50%)	IC↑ (50%)	AD↓ (15%)	IC↑ (15%)
VGG-16	Max-pooling layers	T-TAME (Proposed)	9.33	50	36.5	22.45	<u>73.29</u>	<u>5.6</u>
		No skip connection	10.09	45.25	36.44	20.65	74.85	5.15
		No skip + No batch norm	5.92	57.9	<u>34.49</u>	24.2	74.58	5.15
		Sigmoid in feature branch	<u>7.22</u>	<u>55.65</u>	38.4	21.6	79	4.85
		FMs from two layers	10.72	45.45	34.48	<u>23.05</u>	71.94	5.75
		FMs from ne layer	12.1	42.1	35.81	20.8	74.19	4.85
VGG-16	Convolutional layers	T-TAME (Proposed)	9.07	51.1	40.72	20.9	<u>77.05</u>	<u>4.8</u>
		No skip connection	6.22	<u>58.85</u>	41.47	20.9	79.12	3.8
		No skip + No batch norm	<u>6.62</u>	56.6	40.48	<u>20.75</u>	77.84	4.95
		Sigmoid in feature branch	6.8	60	42.17	19.75	80.73	4.1
		FMs from two layers	10.99	45.85	<u>40.89</u>	19.55	76.66	<u>4.8</u>
		FMs from one layer	13.09	39.65	42.3	17.7	78.02	3.8
ResNet-50	Stage Outputs	T-TAME (Proposed)	<u>7.81</u>	<u>54</u>	<u>27.88</u>	27.5	<u>78.58</u>	4.9
		No skip connection	5.7	62.65	46.58	18.25	89.32	2.3
		No skip + No batch norm	9.29	50.25	29.43	<u>25.95</u>	79.81	3.95
		Sigmoid in feature branch	9.11	53.3	45.68	18.1	86.95	3.15
		FMs from two layers	9.48	47.05	27.83	25	77.95	<u>4.25</u>
		FMs from one layer	11.32	43.45	29.85	24.25	79.59	3.55
ViT-B-16	Block Outputs	T-TAME (Proposed)	8.19	<u>38.35</u>	23.64	<u>40.40</u>	72.89	9.40
		No skip connection	10.16	38.65	24.42	41.65	73	<u>9.35</u>
		No skip + No batch norm	<u>8.91</u>	36.75	24.79	39.75	72.98	8.70
		Sigmoid in feature branch	12.75	35.50	26.79	38.85	72.99	8.85
		FMs from two layers	9.23	37.70	<u>24.08</u>	40	73.69	9.15
		FMs from one layer	9.88	37.60	24.96	40.20	<u>72.97</u>	9.40




FIGURE 9. Comparison of methods using the MoRF measure of ROAD on the VGG-16 backbone.




FIGURE 10. Comparison of methods using the LeRF measure of ROAD on the VGG-16 backbone.

(i) For the CNN backbones, T-TAME generally provides the best performance. Specifically, in the case of the VGG-16 backbone, for the AD and IC measures,

T-TAME provides the best results for the more challenging $v = 50%$ and $v = 15%$ thresholds and is only outperformed in the less-challenging $v = 100%$ setup

TABLE 7. Ablation study: comparison of mismatched backbone-masking procedure combination.

Backbone	Measure	ViT Masking (Eq. (14))	CNN Masking (Eq. (13))
VGG-16	AD↓(100%)	12.44	9.33
	IC↑(100%)	42.05	50.00
	AD↓(50%)	41.38	36.50
	IC↑(50%)	18.05	22.45
ResNet-50	AD↓(15%)	76.81	73.29
	IC↑(15%)	4.90	5.60
	AD↓(100%)	16.97	7.81
	IC↑(100%)	37.20	54.00
ViT-B-16	AD↓(50%)	73.56	27.88
	IC↑(50%)	6.85	27.50
	AD↓(15%)	96.10	78.58
	IC↑(15%)	0.90	4.90
ResNet-50	AD↓(100%)	8.19	9.76
	IC↑(100%)	38.35	39.45
	AD↓(50%)	23.64	25.49
	IC↑(50%)	40.40	41.6
ViT-B-16	AD↓(15%)	72.89	74.85
	IC↑(15%)	9.40	8.20




FIGURE 11. Comparison of methods using the MoRF measure of ROAD on the ResNet-50 backbone.

by the perturbation-based method RISE, which requires 4000 forward passes to generate a single explanation (thus, being 4000 times more computationally expensive than T-TAME at inference time). In the case of the ResNet-50 backbone, for the AD and IC measures, T-TAME is overall the top-performing method, while being second-best in one instance. In that instance, it is outperformed by the perturbation-based method ScoreCAM, which however requires 2048 forward passes (instead of one, for T-TAME) to generate a single explanation.

From the averaged ROAD measures of Table 4, in the case of the VGG-16 backbone, we observe that T-TAME

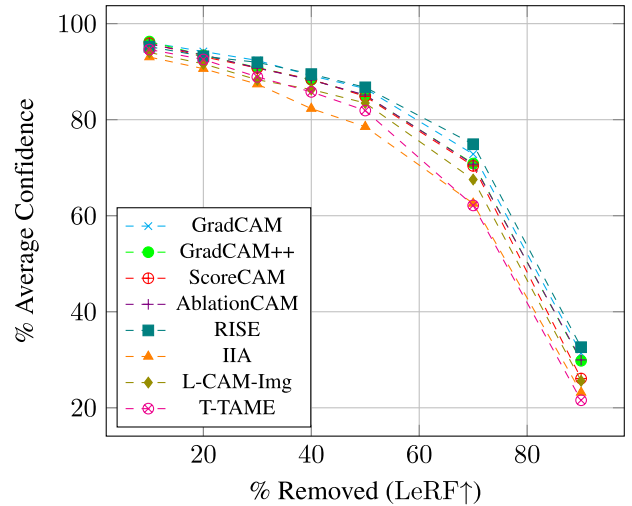


FIGURE 12. Comparison of methods using the LeRF measure on the ResNet-50 backbone.




FIGURE 13. Comparison of methods using the MoRF measure of ROAD on the ViT-B-16 backbone.

achieves the best results w.r.t. the MoRF measure. According to the LeRF measure, it is outperformed only by RISE, as in the case of the AD and IC metrics. In the case of the ResNet-50 backbone, from Table 4 and Fig. 11 we observe that w.r.t. MoRF, the explanation maps of RISE produce the lowest (i.e., best) average confidence but are overtaken by T-TAME in the higher removal percentages (50% or more, in Fig. 11). These results suggest that T-TAME correctly identifies the important regions, but the exact pixel-wise importance ordering is noisy. Additionally, w.r.t. LeRF for this backbone, RISE has the highest average confidence. To explain the LeRF results of T-TAME in the case of ResNet-50, we should recall that LeRF is computed by removing the less important features of the input image, and takes into account only the ordering of pixels according to the explanation map. As can be seen in




FIGURE 14. Comparison of methods using the LeRF measure of ROAD on the ViT-B-16 backbone.

Fig. 15, because of the low-resolution feature maps that are due to the specifics of the ResNet-50 architecture, the produced explanations are overly smooth. While they highlight the important regions of the input image, the ordering of less important pixels is noisy. Since for the computation of the ROAD measures the ordering of pixel importance is the only consideration for their computation, this quality is detrimental. Still, the ability of T-TAME to generate explanation maps in a single forward pass is a significant advantage for practical applications.

- (ii) For the ViT-B-16 backbone, T-TAME is the top-performing method across the board. It performs best for all thresholds of the AD and IC measures. It is the second-best method only in the case of the LeRF measure, being outperformed by RISE, which, in this case, requires 8000 forward passes to generate a single explanation. Moreover, T-TAME outperforms the Transformer-specific LRP-based method. In the case of MoRF, as observed in Fig. 13, T-TAME exhibits the overall best performance for all percentages except for the initial $v = 10\%$ removal percentage. Particularly for $v = 30\%$ to $v = 70\%$, the difference between T-TAME and the second-best method, Transformer LRP, is large. This suggests that, except for the very fine-grained ordering examined in the case of $v = 10\%$, T-TAME correctly identifies the most important pixels for the ViT-B-16 backbone. In the case of LeRF, RISE is initially the top-performing method, being outperformed by T-TAME in the higher removal percentages. This again suggests a more globally correct ordering of importance, with less finely-grained orderings in the lower percentages. Considering that T-TAME requires only one forward pass to compute an explanation, it is significant that it can compete with and in most cases outperform perturbation-based approaches.

E. ABLATION STUDIES

In this section, we perform several ablation studies to assess the effects of different architectural choices of the T-TAME attention mechanism and to observe the effect of the different masking procedures when a CNN (Eq. (13)) or ViT backbone (Eq. (14)) is used. We measure the performance utilizing only the AD(v) and IC(v) measures, to allow a more straightforward interpretation of the results, and additionally, because the ROAD measures are much more computationally expensive.

1) DIFFERENT ARCHITECTURAL CHOICES OF THE ATTENTION MECHANISM

Results of this set of ablation experiments are reported in Table 6, where we indicate with bold/underline the best/second best results according to each measure for each model and layer selection. For the VGG-16 model, inspired by similar works in the literature suggesting that the last layers of the network provide more salient features [40], we report two sets of experiments, one that uses feature maps extracted from the three last max-pooling layers and one where feature maps are extracted from the layers directly before the last three max-pooling layers (Fig. 6). There is a difference in the spatial dimensions of the explanation maps generated using the former and the latter layers for feature extraction, i.e., 28×28 versus 56×56 , since the dimension of the explanation maps obtained by T-TAME is dictated by that of the employed feature maps (as explained in Section III-C1). For the ResNet-50 model, we extract feature maps from the outputs of the final three stages, resulting in an explanation map of 28×28 pixels. In the case of the ViT-B-16 model, feature maps are extracted from the outputs of the final three encoder blocks, resulting in an explanation map of 14×14 pixels. For each backbone and set of considered feature maps, we examine the following variants of the proposed architecture:

No skip connection: It has been shown that the inclusion of a skip connection promotes a smoother loss landscape [57] and preserves gradients that might otherwise be lost or diluted by passing through multiple layers, thus improving the training of very deep neural networks. Even for shallower neural networks, such as the proposed attention mechanism, we can benefit from using a skip connection. We see that by omitting the skip connection shown in Fig. 3(a), we get worse results in ResNet-50 for the more challenging $v = 50\%$ and $v = 15\%$ measures. Similarly, for the VGG-16 backbone, we report worse performance for the harder $v = 50\%$ and $v = 15\%$ measures. In the case of ViT-B-16, the proposed architecture that includes this skip connection prevails in the more challenging $v = 15\%$ metric.

No skip + No batch norm: Batch normalization is used in neural networks for speeding up training and combating internal covariate shift [58]. Compared to the proposed architecture of Fig. 3(a), we see that, in the case of VGG-16, this variant generally performs better in the $v = 100\%$ measures, but this does not hold for the other measures.

Sigmoid in feature branch: In this variant, we replace the ReLU function of Fig. 3(a) with the sigmoid function, which squeezes the input from $(-\infty, \infty)$ to the output $(0, 1)$. It is well known that the sigmoid function in deeper neural networks causes the vanishing gradient problem, making it more difficult to train the early layers of the neural network. We see again that the proposed architecture of Fig. 3(a) prevails for the more challenging $\nu = 15\%$ measures.

Two layers and One layer: In this case, the proposed attention mechanism architecture is employed with feature maps from fewer than three layers. The results when using just one layer, i.e., omitting the two earlier layers of the backbone (Fig. 6), are very similar to the L-CAM-Img method (as shown in Table 2), which also uses just one feature map. In the case of CNN backbones, all measures are improved when utilizing a second feature map instead of just one, i.e., excluding only the third (earliest) layer in Figs. 6, 7, 8. When shifting from using feature maps from two to three layers, the results are somewhat mixed; these mixed results could be attributed to the extra noise of feature maps taken earlier in the backbone's pipeline. However, considering these results across all backbones supports the choice of utilizing three feature maps in T-TAME.

Overall remarks on the attention mechanism: We note that by omitting both the skip connection and the batch normalization in the feature branch architecture, we obtain generally better results in the case of the VGG-16 model, but this is not the case for the same architecture applied to the ResNet-50 model. In addition, all the examined architecture variations struggle under the more challenging $\nu = 15\%$ measures, being in most cases outperformed by the proposed T-TAME architecture; the latter is shown to generalize the best across different backbone models.

2) MISMATCHED BACKBONE-MASKING PROCEDURE COMBINATION

As discussed in Section III-C3, CNN backbones are generally sensitive to out-of-distribution samples. Thus, in T-TAME, we introduced different procedures for masking the input with the explanation maps when working with Convolutional (Eq. (13)) or Transformer-like (Eq. (14)) backbones. In this ablation experiment, we assess the effect of switching these procedures, i.e., conversely applying our CNN-specific masking procedure on the ViT backbone and our ViT-specific masking procedure on the CNN backbones. We can see in Table 7 that for the ViT-B-16 backbone, using our ViT-specific masking procedure is beneficial, especially when looking at the challenging $\nu = 15\%$ measures. For ResNet-50, the performance differences caused by switching the masking procedure are much greater, demonstrating the sensitivity of the skip connections and of the overall ResNet architecture to out-of-distribution inputs. Similarly, in the case of the VGG-16 backbone: the degradation of performance when masking inputs using the ViT-specific procedure is clear, although less pronounced than what it was for ResNet-50. This can be attributed to the fact that the VGG-16 architecture has no skip connections: in [59] it

has been shown that limiting the number of skip connections improves robustness. Summarily, this ablation experiment demonstrates the importance of handling the perturbation of inputs in the case of CNN and ViT backbones differently, in agreement with what we proposed in Section III-C3. Additionally, an interesting observation is that ViT is less sensitive to the choice of masking procedure than the two examined CNNs; this is consistent with the findings of [27] on the robustness of the ViT architecture to out-of-distribution samples.

F. QUALITATIVE ANALYSIS

An extensive qualitative analysis is performed using images from the evaluation partition of the ILSVRC 2012 dataset. Specifically, we present visualization examples across different backbones for the T-TAME method (Fig. 15); and, focusing on the ViT backbone, for T-TAME and all other compared methods of Table 3 (Fig. 16). Additionally, we conduct model randomization sanity checks (following the protocol of [15]) on the T-TAME method (Fig. 17). Finally, in Subsection IV-F4 we provide examples where the T-TAME-generated explanations can help us to gain specific insights about the backbone model and the dataset (Figs. 18 and 19).

1) QUALITATIVE COMPARISON OF T-TAME EXPLANATIONS ACROSS DIFFERENT BACKBONES

The qualitative differences between explanations produced using T-TAME for the VGG-16, ResNet-50, and ViT-B-16 backbones are examined in Fig. 15. We observe that explanations produced for the VGG-16 and ResNet-50 models are generally more focused on specific regions compared to the ViT-B-16 backbone, and explanations produced for the three different backbone types primarily attend to different areas of the image. This can be explained by the fact that T-TAME is essentially trained by perturbing the original input image. ViTs are more robust to occlusions and perturbations [44]. By leveraging disjoint and spatially separate regions, ViTs retain high accuracy even when using masked inputs (see also Section III-C3). This result suggests that VGG-16, ResNet-50, and ViT-B-16 classify images in fundamentally different ways, focusing on different features of an input image to make their predictions. The more global way in which ViT-B-16 (and Transformers, in general) interprets input images could be one of the reasons that such Transformer-based architectures perform better on the ILSVRC 2012 dataset.

2) EXPLANATION MAPS FOR THE VISION TRANSFORMER

In Fig. 16, explanation maps for the ViT-B-16 backbone produced using different explanation methods are depicted. We observe that the proposed T-TAME (last row) generates the most activated explanation maps, followed by Ablation-CAM (row six) and Score-CAM (row four). Most other methods activate only on a small, and usually a different, part of the object in the image. For instance, observing the explanation maps in the second column of Fig. 16 concerning

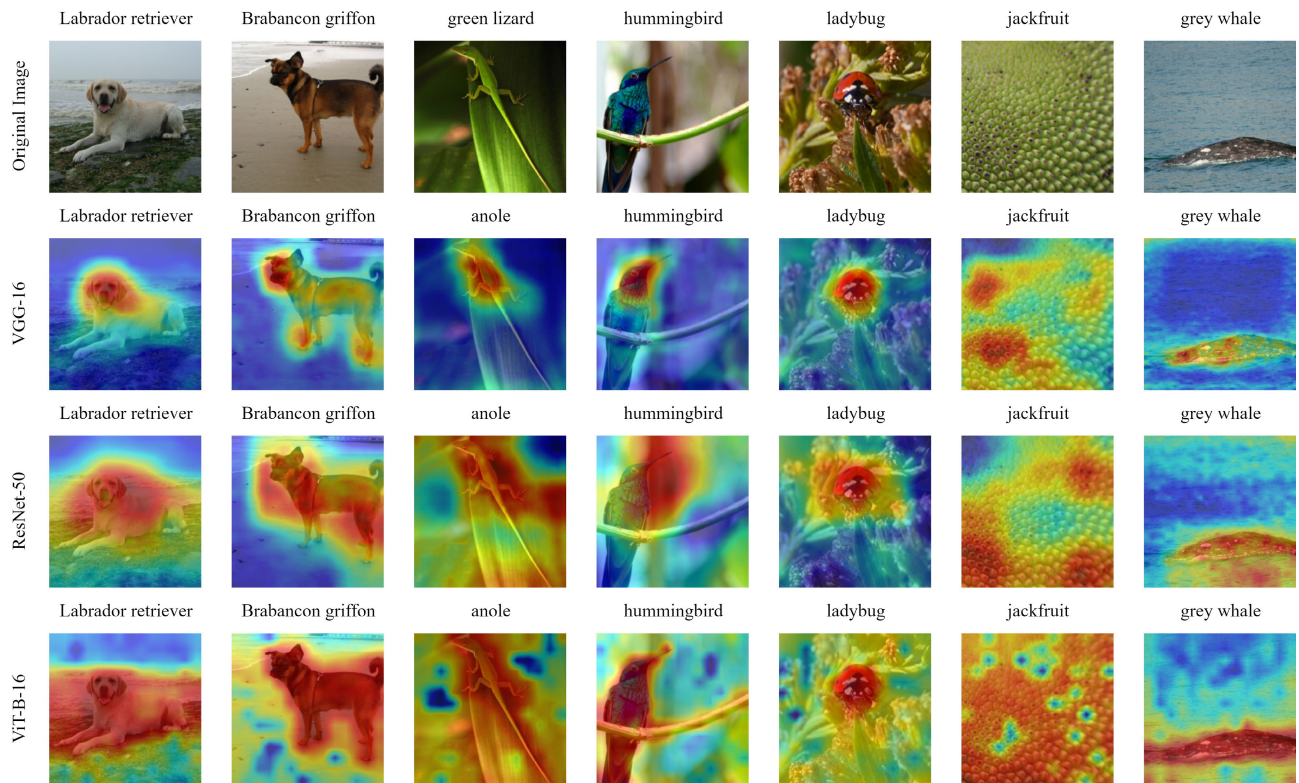


FIGURE 15. T-TAME applied to VGG-16, ResNet-50 and ViT-B-16 backbones. We report the ground truth classes for each input image (top) and the predicted classes for each backbone (above the corresponding explanation map). A general observation is that the explanation maps produced using the ViT-B-16 backbone attribute significance to larger image regions in comparison to the CNN backbones, highlighting the global view of the input thanks to the Transformer's Multi-head Attention layer.

the Brabancon griffon, we see that all methods besides T-TAME focus on the body, on the neck and back part of the head, or the mouth and nose.

Contrarily to the other methods, the explanation maps of T-TAME tend to highlight the overall region of the object corresponding to the model-truth label, and at the same time provide the required granularity in the activation values so that the parts of the object that explain mostly the decision of the classifier are activated at a higher degree, as shown by the very good results with the AD, IC and ROAD measures (reported in Tables 3 and 5). This shows the effectiveness of T-TAME in revealing the long-term relations between patches captured by the ViT-B-16 multi-head attention layer and its ability to identify the salient image regions. Additionally, this demonstrates the importance of evaluating the various explainability methods using the AD and IC measures at multiple ν thresholds (Table 6), and particularly the significance of the $\nu = 15\%$ over the $\nu = 100\%$ and $\nu = 50\%$ threshold measures in judging the quality of the generated explanations.

3) SANITY CHECKS OF T-TAME

Sanity checks for explanation maps [15] aim to ensure that explainability methods produce explanations that are dependent on the specific mechanism by which the backbone network processes its inputs to reach a classification decision.

By randomizing the backbone network, or the dataset image-label pairs, we expect to see drastic changes in the produced explanation maps. If these changes are not observed, the method of explanation map generation does not explain the specific backbone's decision-making mechanism. It may instead simply detect image edges, or simulate other basic image filtering methods to generate superficially-convincing explanation maps. The methods used for the comparison studies (Tables 2 and 3) have been observed in [14], [15], [20], [21], [25], and [51] to pass the sanity checks, so we will focus on T-TAME. We conduct two types of sanity checks on T-TAME.

In the first case, depicted in Fig. 17(a), we gradually randomize the layers of the ViT-B-16 backbone network from the output layer to the input layer. We examine the effects that layer randomization has on the explanations produced for a specific image. We witness significant and abrupt differentiation between the produced explanations and the original explanation. Specifically, after randomizing the logit-producing layer, and the fifth encoder layer, we notice a major shift in the highlighted salient regions. After having randomized the entire backbone, the produced explanation bears very little resemblance to the initial explanation. This is the expected and desired result since a randomized backbone produces random results, thus no reasonable explanations for its decisions can be produced.




FIGURE 16. Qualitative comparison between T-TAME and the other explainability methods of Table 3 for the ViT-B-16 backbone. We observe that T-TAME produces more activated explanation maps, demonstrating the global context used by the ViT-B-16 architecture.

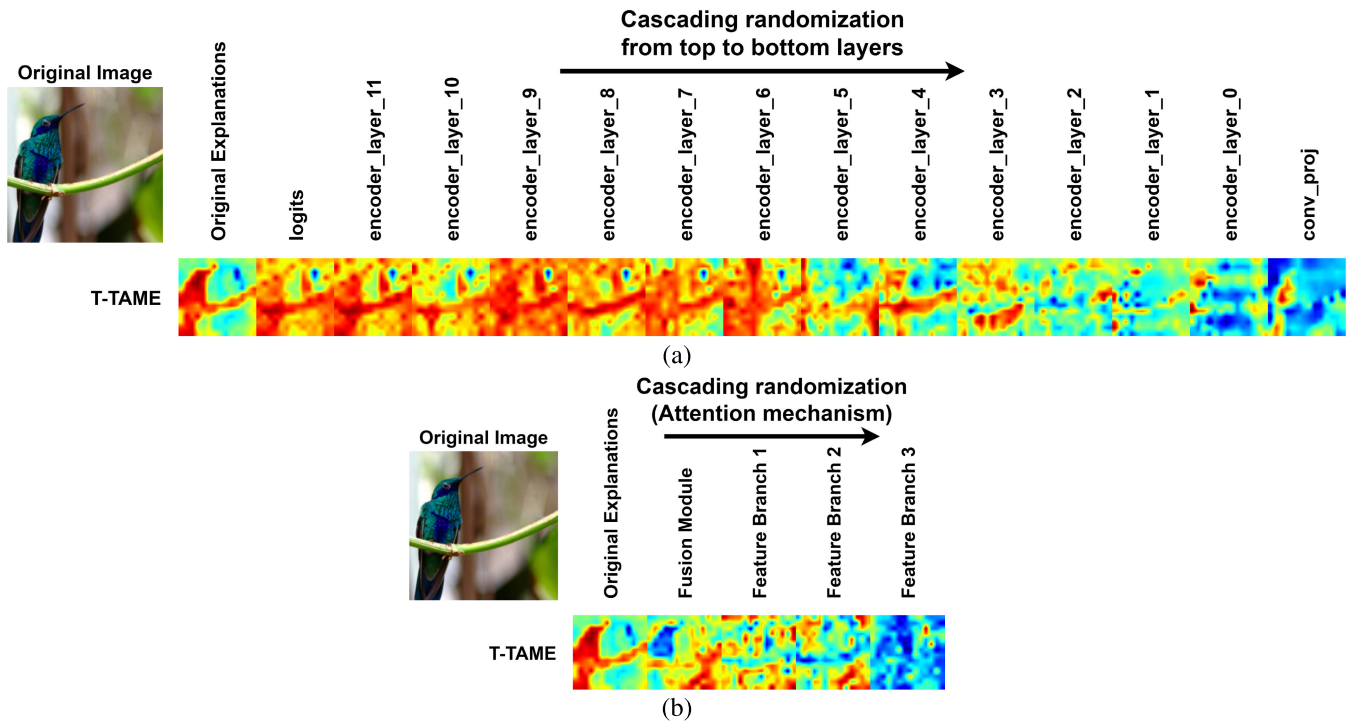


FIGURE 17. Qualitative sanity check of the proposed T-TAME method. In (a) we randomize the weights of the backbone network (ViT-B-16) in a cascading manner. In (b) we gradually randomize the attention mechanism of T-TAME. We can observe a drastic drop in the quality of the produced explanation when randomizing the backbone, starting with the logit-producing layer and finishing with the initial patch-processing convolutional layer. When randomizing the attention mechanism, the result is also a dramatic change in the produced explanation map.

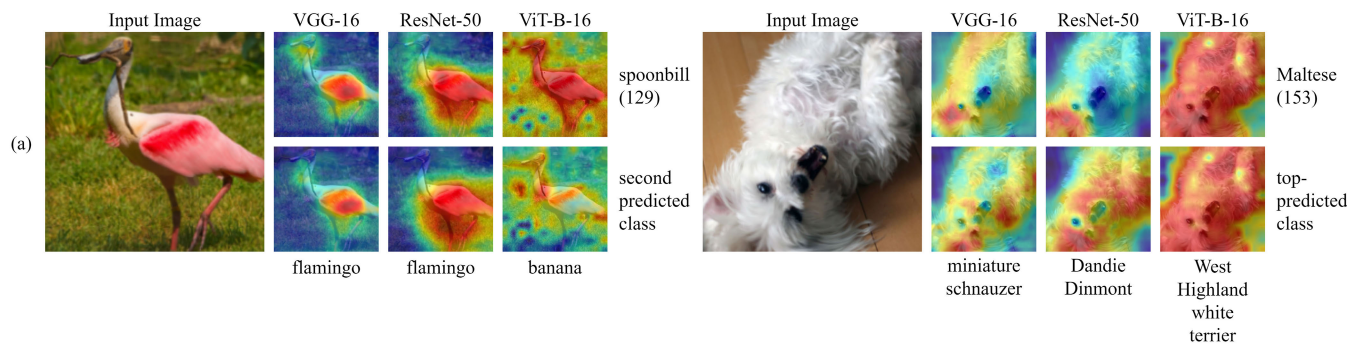


FIGURE 18. Counterfactual explanations for two input images. In each case, we display six class-specific explanations for VGG-16, ResNet-50, and ViT-B-16. The first row of explanations for each image corresponds to the image’s ground-truth class, whereas the second row to the other classes: for the image on the left that is correctly classified by all three backbones, these are the second-best predictions of each backbone, while for the image on the right that is misclassified by all three backbones, these are the erroneously-predicted class of each backbone.

In the second case, depicted in Fig. 17(b), we randomize the trained attention mechanism of the T-TAME method in a cascading manner (that is, this sanity check is specific to T-TAME). After randomizing the fusion module, we observe a considerable change in the produced explanation. The produced explanation map further resembles a random heatmap, as feature branches are consecutively randomized. This is again the desirable result of this sanity check, as it demonstrates that the training step of the T-TAME attention mechanism results in weights that are necessary for producing meaningful explanation maps.

4) EXAMPLE INSIGHTS ON IMAGENET CLASSIFIERS

In Fig. 18, we provide class-specific explanation maps referring to the ground truth class but also to an erroneous class, for the three examined backbones, to examine how T-TAME can assist in model interpretability. Interpretability refers to the rationale employed by a model to generate its decisions. It is different from explainability because the focus is on the model instead of a specific classification decision. The first image (left side example) is correctly classified by all of the examined backbones. The explanations for the second-highest predicted class, by each backbone,

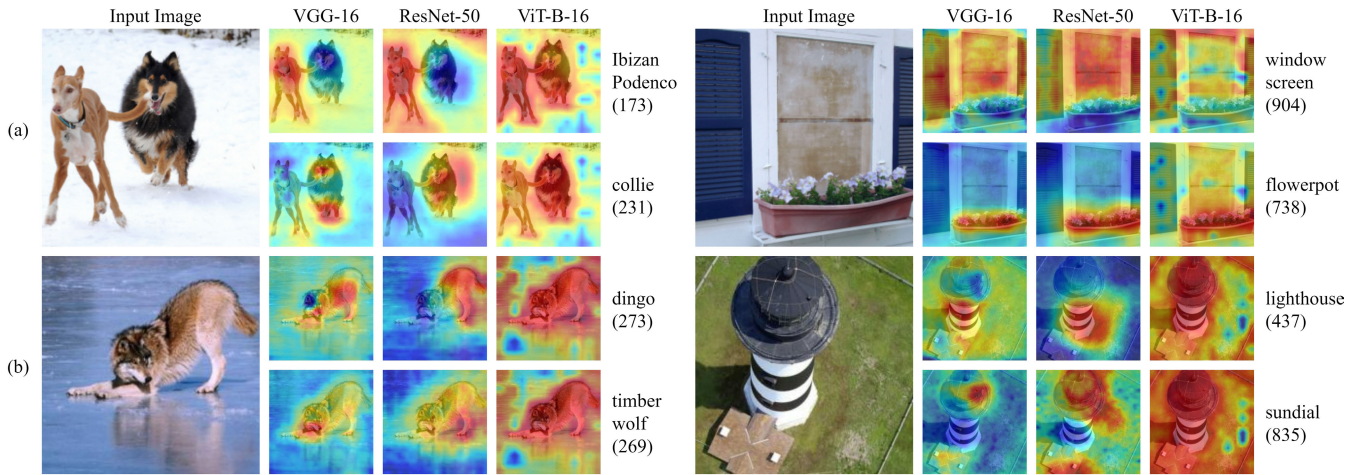


FIGURE 19. Explanations for four input images. In each case, we display six class-specific explanations, i.e., of the true (ground truth) (top) and erroneous (bottom) class prediction of the input image, for VGG-16, ResNet-50, and ViT-B-16. In (a), example images with multiple classes, along with generated explanations for each respective class are depicted. In (b), two cases of misclassification are provided: dataset misclassification (left-side example) and model misclassification (right-side example).

are also depicted. The second image (on the right) is incorrectly classified by all of the examined backbones. The explanations for the class predicted by each model, along with the explanations for the ground truth class, are shown. By comparing the explanation maps for adversarial classes, we can probe at the underlying decision strategy and possibly gain new insights for the classifier. For example, for the first image, which depicts a “spoonbill”, in the case of the CNN backbones, the second-highest predicted class is the class “flamingo”. These two animals share many visual characteristics, such as body shape and color. In the case of the ViT backbone, the second-highest predicted class is “banana”, a seemingly unrelated class to the input image. Both CNNs seem to generate their decision from generic visual characteristics such as color, shape, and background. The Transformer-based architecture seems to employ a different strategy: the image has been classified as a spoonbill with high confidence, and no other class is considered possible, so the decision “banana” has near zero confidence. The second image (right side example) is incorrectly classified by all of the examined backbones; the top-predicted class is different for each backbone. The ground truth class’s explanation map is also depicted. The initially predicted classes are all visually similar dog breeds to the ground truth class, but even for the ViT backbone the confidence in its prediction is not high: the model recognizes that classification is unclear in this instance, instead of always outputting a single prediction of high confidence.

The examples of Fig. 19 (a) demonstrate the potential of the explanation maps to be used for explaining multiple different classes contained in a single image, i.e., the “Ibizan Podenco” and “collie” image, and the “window screen” and “flowerpot” image. All models can clearly distinguish between the various classes contained in the images. Interestingly, the ViT backbone highlights both dogs in the first example, varying only in the intensity of the

explanation map, instead of considering the second dog a negative presence in the image, as do the CNN backbones. This corroborates with our findings that ViT models interpret the input image more globally and relationally (it may be more likely for multiple animals to exist, rather than a single animal, in an image of the ImageNet dataset).

Finally, in Fig. 19 (b) we provide two cases of images that have been misclassified, i.e., the predicted class is not in agreement with the ground truth label of the dataset, and we use the explanations to understand what went wrong. The first image of Fig. 19 (b) belongs according to its ground truth label to the “dingo” class (273) but is misclassified as “timber wolf” by all three backbones. Visual inspection reveals that the image evidently belongs to the “timber wolf” class, hence this is a case of dataset mislabeling; the backbone classifiers correctly focused on meaningful parts of the image to make their decisions. The second image depicts a lighthouse. VGG-16 misclassified this image as a “sundial”. Again, using the explanations generated by T-TAME, we can understand which features led the model to produce a wrong decision. For instance, in this case, we see that for both CNN models, the “sundial” explanations focus on the lighthouse roof, which might resemble a sundial, explaining the erroneous classification decision of VGG-16. ViT-B-16 correctly classifies this image. The ViT-B-16 explanation does focus more on the roof as well, but it is much less concentrated on a specific region, and in this lighthouse example also focuses on the perimeter fence of the building, again showing that this classifier utilizes information from multiple parts of the image.

V. CONCLUSION

We proposed T-TAME, a novel method for generating visual explanations for deep-learning-based image classifiers. This is accomplished by training a hierarchical attention mechanism to make use of feature maps that are extracted from

multiple layers of the backbone classifier. These feature maps are appropriately transformed according to the type of the backbone network, making T-TAME compatible with both CNN and Transformer-based classifier architectures. Experimental results verified that T-TAME clearly outperforms gradient-based and non-trainable relevance-based explainability methods, and outperforms or is on par with perturbation-based methods while, in contrast to them, it requires only a single forward pass to generate explanations. Possible future directions include the application of T-TAME to medical image classification problems; and, the investigation of how we could mitigate the effects of masking with low-resolution feature maps in backbones such as ResNet-50, where the output of the backbone's last stages is inevitably of low spatial resolution.

REFERENCES

- [1] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16×16 words: Transformers for image recognition at scale," 2020, *arXiv:2010.11929*.
- [2] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jegou, "Training data-efficient image transformers & distillation through attention," in *Proc. 38th Int. Conf. Mach. Learn.*, vol. 139, 2021, pp. 10347–10357.
- [3] B. Gheflati and H. Rivaz, "Vision transformers for classification of breast ultrasound images," in *Proc. 44th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. (EMBC)*, Jul. 2022, pp. 480–483.
- [4] C. Xin, Z. Liu, K. Zhao, L. Miao, Y. Ma, X. Zhu, Q. Zhou, S. Wang, L. Li, F. Yang, S. Xu, and H. Chen, "An improved transformer network for skin cancer classification," *Comput. Biol. Med.*, vol. 149, Oct. 2022, Art. no. 105939.
- [5] Y. Zhong and W. Deng, "Face transformer for recognition," 2021, *arXiv:2103.14803*.
- [6] R. Hamon, H. Junklewitz, I. Sanchez, G. Malgieri, and P. De Hert, "Bridging the gap between AI and explainability in the GDPR: Towards trustworthiness-by-design in automated decision-making," *IEEE Comput. Intell. Mag.*, vol. 17, no. 1, pp. 72–85, Feb. 2022.
- [7] Ł. Górski and S. Ramakrishna, "Explainable artificial intelligence, lawyer's perspective," in *Proc. 18th Int. Conf. Artif. Intell. Law*, Jun. 2021, pp. 60–68.
- [8] J. Amann, A. Blasimme, E. Vayena, D. Frey, and V. I. Madai, "Explainability for artificial intelligence in healthcare: A multidisciplinary perspective," *BMC Med. Informat. Decis. Making*, vol. 20, no. 1, pp. 1–9, Dec. 2020.
- [9] A. Barredo Arrieta, N. Díaz-Rodríguez, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, S. Garcia, S. Gil-Lopez, D. Molina, R. Benjamins, R. Chatila, and F. Herrera, "Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," *Inf. Fusion*, vol. 58, pp. 82–115, Jun. 2020.
- [10] W. Samek, G. Montavon, S. Lapuschkin, C. J. Anders, and K.-R. Müller, "Explaining deep neural networks and beyond: A review of methods and applications," *Proc. IEEE*, vol. 109, no. 3, pp. 247–278, Mar. 2021.
- [11] P.-T. Jiang, C.-B. Zhang, Q. Hou, M.-M. Cheng, and Y. Wei, "LayerCAM: Exploring hierarchical class activation maps for localization," *IEEE Trans. Image Process.*, vol. 30, pp. 5875–5888, 2021.
- [12] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 618–626.
- [13] A. Chattopadhyay, A. Sarkar, P. Howlader, and V. N. Balasubramanian, "Grad-CAM++: Generalized gradient-based visual explanations for deep convolutional networks," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Mar. 2018, pp. 839–847.
- [14] H. Chefer, S. Gur, and L. Wolf, "Transformer interpretability beyond attention visualization," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 782–791.
- [15] J. Adebayo, J. Gilmer, M. Muelly, I. Goodfellow, M. Hardt, and B. Kim, "Sanity checks for saliency maps," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 9525–9536.
- [16] P.-J. Kindermans, S. Hooker, J. Adebayo, M. Alber, K. T. Schütt, S. Dähne, D. Erhan, and B. Kim, *The (Un)reliability of Saliency Methods*. Cham, Switzerland: Springer, 2019, pp. 267–280.
- [17] G. Montavon, S. Lapuschkin, A. Binder, W. Samek, and K.-R. Müller, "Explaining nonlinear classification decisions with deep Taylor decomposition," *Pattern Recognit.*, vol. 65, pp. 211–222, May 2017.
- [18] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek, "On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation," *PLoS ONE*, vol. 10, no. 7, Jul. 2015, Art. no. e0130140.
- [19] A. Holzinger, R. Goebel, R. Fong, T. Moon, K. R. Müller, and W. Samek, "XXAI-Beyond explainable artificial intelligence," in *Proc. ICML*, Jul. 2020, pp. 3–10.
- [20] V. Petsiuk, A. Das, and K. Saenko, "RISE: Randomized input sampling for explanation of black-box models," in *Proc. BMVC*, Sep. 2018, pp. 1–19.
- [21] H. Wang, Z. Wang, M. Du, F. Yang, Z. Zhang, S. Ding, P. Mardziel, and X. Hu, "Score-CAM: Score-weighted visual explanations for convolutional neural networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2020, pp. 111–119.
- [22] S. Sattarzadeh, M. Sudhakar, A. Lem, S. Mehryar, K. N. Plataniotis, J. Jang, H. Kim, Y. Jeong, S. Lee, and K. Bae, "Explaining convolutional neural networks through attribution-based input sampling and block-wise feature aggregation," in *Proc. AAAI*, vol. 35, 2021, pp. 11639–11647.
- [23] M. Sudhakar, S. Sattarzadeh, K. N. Plataniotis, J. Jang, Y. Jeong, and H. Kim, "Ada-SISE: Adaptive semantic input sampling for efficient explanation of convolutional neural networks," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Jun. 2021, pp. 1715–1719.
- [24] A. Englebret, O. Cornu, and C. de Vleeschouwer, "Backward recursive class activation map refinement for high resolution saliency map," in *Proc. 26th Int. Conf. Pattern Recognit. (ICPR)*, Aug. 2022, pp. 2444–2450.
- [25] O. Barkan, Y. Elisha, Y. Asher, A. Eshel, and N. Koenigstein, "Visual explanations via iterated integrated attributions," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2023, pp. 2073–2084.
- [26] I. Gkartzonika, N. Gkalelis, and V. Mezaris, "Learning visual explanations for DCNN-based image classifiers using an attention mechanism," in *Proc. Workshop Vis. Biased Scarce Data*, Oct. 2022, pp. 1–19.
- [27] Y. Bai, J. Mei, A. Yuille, and C. Xie, "Are transformers more robust than CNNs?" 2021, *arXiv:2111.05464*.
- [28] M. Ntroukas, N. Gkalelis, and V. Mezaris, "TAME: Attention mechanism based feature fusion for generating explanation maps of convolutional neural networks," in *Proc. IEEE Int. Symp. Multimedia (ISM)*, Dec. 2022, pp. 58–65.
- [29] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. ICLR*, May 2015, pp. 1–15.
- [30] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [31] R. Confalonieri, L. Coba, B. Wagner, and T. R. Besold, "A historical perspective of explainable artificial intelligence," *WIREs Data Mining Knowl. Discovery*, vol. 11, no. 1, p. e1391, Jan. 2021.
- [32] T. Miller, "Explanation in artificial intelligence: Insights from the social sciences," *Artif. Intell.*, vol. 267, pp. 1–38, Feb. 2019.
- [33] W. Wu, Y. Su, X. Chen, S. Zhao, I. King, M. R. Lyu, and Y.-W. Tai, "Towards global explanations of convolutional neural networks with concept attribution," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 8649–8658.
- [34] M. Böhle, M. Fritz, and B. Schiele, "Holistically explainable vision transformers," 2023, *arXiv:2301.08669*.
- [35] B. Crook, M. Schlüter, and T. Speith, "Revisiting the performance-explainability trade-off in explainable artificial intelligence (XAI)," in *Proc. IEEE 31st Int. Requirements Eng. Conf. Workshops (REW)*, Sep. 2023, pp. 316–324.
- [36] V. Chamola, V. Hassija, A. R. Sulthana, D. Ghosh, D. Dhingra, and B. Sikdar, "A review of trustworthy and explainable artificial intelligence (XAI)," *IEEE Access*, vol. 11, pp. 78994–79015, 2023.
- [37] V. Hassija, V. Chamola, A. Mahapatra, A. Singal, D. Goel, K. Huang, S. Scardapane, I. Spinelli, M. Mahmud, and A. Hussain, "Interpreting black-box models: A review on explainable artificial intelligence," *Cognit. Comput.*, vol. 16, no. 1, pp. 45–74, Jan. 2024.

- [38] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2921–2929.
- [39] M. Sundararajan, A. Taly, and Q. Yan, "Axiomatic attribution for deep networks," 2017, *arXiv:1703.01365*.
- [40] S. Jetley, N. A. Lord, N. Lee, and P. H. S. Torr, "Learn to pay attention," 2018, *arXiv:1804.02391*.
- [41] H. Fukui, T. Hirakawa, T. Yamashita, and H. Fujiyoshi, "Attention branch network: Learning of attention mechanism for visual explanation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 10697–10706.
- [42] L. Itti, C. Koch, and E. Niebur, "A model of saliency-based visual attention for rapid scene analysis," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 20, no. 11, pp. 1254–1259, Nov. 1998.
- [43] X. Pei, Y. H. Zhao, L. Chen, Q. Guo, Z. Duan, Y. Pan, and H. Hou, "Robustness of machine learning to color, size change, normalization, and image enhancement on micrograph datasets with large sample differences," *Mater. Des.*, vol. 232, Aug. 2023, Art. no. 112086.
- [44] S. Jain, H. Salman, E. Wong, P. Zhang, V. Vineet, S. Vemprala, and A. Madry, "Missingness bias in model debugging," 2022, *arXiv:2204.08945*.
- [45] P. Chen, S. Liu, H. Zhao, X. Wang, and J. Jia, "GridMask data augmentation," 2020, *arXiv:2001.04086*.
- [46] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. NIPS*, 2017, pp. 6000–6010.
- [47] C. Zhang, M. Zhang, S. Zhang, D. Jin, Q. Zhou, Z. Cai, H. Zhao, X. Liu, and Z. Liu, "Delving deep into the generalization of vision transformers under distribution shifts," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 7267–7276.
- [48] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D, Nonlinear Phenomena*, vol. 60, nos. 1–4, pp. 259–268, Nov. 1992.
- [49] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet large scale visual recognition challenge," 2014, *arXiv:1409.0575*.
- [50] Y. Rong, T. Leemann, V. Borisov, G. Kasneci, and E. Kasneci, "A consistent and efficient evaluation strategy for attribution methods," 2022, *arXiv:2202.00449*.
- [51] S. Desai and H. G. Ramaswamy, "Ablation-CAM: Visual explanations for deep convolutional network via gradient-free localization," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Mar. 2020, pp. 972–980.
- [52] M. Nauta, J. Trienes, S. Pathak, E. Nguyen, M. Peters, Y. Schmitt, J. Schlötterer, M. van Keulen, and C. Seifert, "From anecdotal evidence to quantitative evaluation methods: A systematic review on evaluating explainable AI," *ACM Comput. Surveys*, vol. 55, no. 13s, pp. 1–42, Dec. 2023.
- [53] A. Hedström, P. Bommer, K. K. Wickström, W. Samek, S. Lapuschkin, and M. M.-C. Höhne, "The meta-evaluation problem in explainable AI: Identifying reliable estimators with MetaQuantus," 2023, *arXiv:2302.07265*.
- [54] L. N. Smith and N. Topin, "Super-convergence: Very fast training of neural networks using large learning rates," *Proc. Artif. Intell. Mach. Learn. Multi-Domain Operations Appl.*, vol. 11006, pp. 369–386, Jun. 2019.
- [55] L. N. Smith, "A disciplined approach to neural network hyper-parameters: Part 1—Learning rate, batch size, momentum, and weight decay," 2018, *arXiv:1803.09820*.
- [56] (2021). *Pytorch Library for CAM Methods*. [Online]. Available: <https://github.com/jacobgil/pytorch-grad-cam>
- [57] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein, "Visualizing the loss landscape of neural nets," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 31, 2018, pp. 1–11.
- [58] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 448–456.
- [59] B. F. Zhang and G. Q. Zhou, "Control the number of skip-connects to improve robustness of the NAS algorithm," *IET Comput. Vis.*, vol. 15, no. 5, pp. 356–365, Aug. 2021.



MARIANO V. NTROUKAS (Graduate Student Member, IEEE) received the M.Eng. degree in electrical and computer engineering from the Aristotle University of Thessaloniki, Greece, in 2023. He is currently a Research Assistant with the Center for Research and Technology Hellas, Information Technologies Institute. His research interests include explainable deep learning and deep learning for biomedical applications.



NIKOLAOS GKALELIS (Member, IEEE) received the B.Sc. degree in electronic engineering from the Technological Institute of Thessaloniki, Greece, in 2000, the M.Sc. degree (Hons.) in control systems engineering from The University of Sheffield, U.K., in 2002, the M.Sc. degree in computational engineering from the University of Erlangen–Nuremberg, Germany, in 2004, and the Ph.D. degree in electrical engineering from Imperial College London, U.K., in 2014.

He worked in research and development with the Aristotle University of Thessaloniki, Siemens, and ECB, and as a Research Associate and a Postdoctoral Researcher with the Centre for Research and Technology Hellas, Information Technologies Institute. He has coauthored six journal articles, one book chapter, 33 conference papers, and one patent. His research interests include deep learning, video analysis, computer vision, and digital signal processing. He serves regularly as a reviewer for many international journals and conferences.



VASILEIOS MEZARIS (Senior Member, IEEE) received the B.Sc. and Ph.D. degrees in electrical and computer engineering from the Aristotle University of Thessaloniki, Greece, in 2001 and 2005, respectively. He is currently a Research Director of the Centre for Research and Technology Hellas, Information Technologies Institute. He has coauthored more than 40 journal articles, 20 book chapters, 200 conference papers, and three patents. His research interests include multimedia understanding and artificial intelligence; in particular, image and video analysis and annotation, machine learning and deep learning for multimedia understanding and big data analytics, explainable AI and green AI, multimedia indexing and retrieval, and applications of multimedia understanding and artificial intelligence. He served as an Associate Editor for IEEE SIGNAL PROCESSING LETTERS and IEEE TRANSACTIONS ON MULTIMEDIA. He serves as a Senior Area Editor for IEEE SIGNAL PROCESSING LETTERS.

...