

CA-Stream: Attention-based pooling for interpretable image recognition

Felipe Torres¹, Hanwei Zhang², Ronan Sicro¹, Stéphane Ayache¹, Yannis Avrithis³

¹Centrale Marseille, Aix Marseille Univ, CNRS, LIS, France

²Institute of Intelligent Software, China

³Institute of Advanced Research on Artificial Intelligence (IARAI), Austria

{felipe.torres, ronan.sicro, stephane.ayache}@lis-lab.fr,
zhanghanwei0912@gmail.com, yannis@avrithis.net

Abstract

Explanations obtained from transformer-based architectures in the form of raw attention, can be seen as a class-agnostic saliency map. Additionally, attention-based pooling serves as a form of masking in feature space. Motivated by this observation, we design an attention-based pooling mechanism intended to replace Global Average Pooling (GAP) at inference. This mechanism, called Cross-Attention Stream (CA-Stream), comprises a stream of cross attention blocks interacting with features at different network depths. CA-Stream enhances interpretability in models, while preserving recognition performance.

1. Introduction

Convolutional neural networks (CNN) have attained tremendous success in computer vision [22, 30], but interpreting their predictions remains challenging. Most explanations are based on saliency maps, using methods derived from *class activation mapping* (CAM). *Vision transformers* [13] are now strong competitors of convolutional networks, characterized by global interactions between patch embeddings in the form of *self attention*. Based on the classification (CLS) token, an explanation map in the form of *raw attention* can be constructed. However, these maps are class-agnostic, often of low quality [6], and dedicated interpretability methods are required to explain models [9].

In CNNs, features are pooled into a global representation by *global average pooling* (GAP). In transformers, a global representation is obtained by cross-attention between patch embeddings and the CLS token. In this work, we make a connection between CAM-based saliency maps and raw attention from the CLS token, observing that attention-based pooling is a form of *masking in the feature space*. Motivated by this observation, we design a pooling mechanism that generates a global representation to be used at infer-

ence, replacing GAP and improving interpretability.

Our approach, called *Cross-Attention Stream* (CA-Stream), consists of a branch in parallel with the backbone network, allowing interactions between feature maps and the CLS token through cross-attention at different stages of the network. The CLS token embedding is a learnable parameter and, at the output of the stream, provides a global image representation for classification.

More specifically, we make the following contributions¹:

1. We demonstrate that attention-based pooling in vision transformers is the same as soft masking by a class-agnostic CAM-based saliency map (section 3.2).
2. We design an attention-based pooling mechanism, inject it in convolutional networks to replace GAP and study its effect on post-hoc interpretability (subsection 3.3).
3. We show improved explanations for a trained model and provides a class-agnostic raw attention map (section 4).

2. Related work

Deep neural networks interpretability is investigated through *Post-hoc interpretability* or *Transparency* [20, 28, 53].

Post-hoc interpretability considers the model as a black-box and provides explanations based on input and output observations. These methods can be grouped into sets of possibly overlapping categories. *Gradient-based methods* [1–3, 42–45] use gradient information to visualize the contribution of different input regions in an image. *CAM-based methods* [8, 12, 16, 25, 40, 47] compute saliency maps as a linear combination of feature maps to highlight salient regions in the input image. *Occlusion or masking-based methods* [14, 15, 32, 36, 39] instead compute saliency maps based on the prediction changes induced

¹This work has received funding from the UnLIR ANR project (ANR-19-CE23-0009), the Excellence Initiative of Aix-Marseille Université - A*Midex, a French “Investissements d’Avenir programme” (AMX-21-IET-017). Part of this work was performed using HPC resources from GENCI-IDRIS (Grant 2020-AD011011853 and 2020-AD011013110)




Figure 1. *Cross-Attention Stream (CA-Stream)* applied to ResNet-based architectures. Given a network f , we replace global average pooling (GAP) by a learned, attention-based pooling mechanism implemented as a stream in parallel to f . The feature tensor $F_\ell \in \mathbb{R}^{p_\ell \times d_\ell}$ (key) obtained by stage Res- ℓ of f interacts with a CLS token (query) embedding $\mathbf{q}_\ell \in \mathbb{R}^{d_\ell}$ in block CA- ℓ , which contains cross attention (6) followed by a linear projection (10) to adapt to the dimension of $F_{\ell+1}$. Here, p_ℓ is the number of patches (spatial resolution) and d_ℓ the embedding dimension. The query is initialized by a learnable parameter $\mathbf{q}_0 \in \mathbb{R}^{d_0}$, while the output \mathbf{q}_5 of the last cross attention block is used as a global image representation into the classifier.

by masking the input image. Finally, *learning-based methods* [7, 10, 33, 39, 58] learn additional models or branches to produce explanations for a given input.

Transparency modifies the model or its training process to explain it. These approaches are grouped according to the nature of the explanation they provide. *Rule-based methods* [48, 49] approximate the model using a decision tree as a proxy. *Hidden semantic-methods* [4, 52, 54, 56] learn disentangled semantics following a hierarchical structure or object-level concepts. *Prototype-based methods* learn prototypes seen in training images to explain models from intermediate representations. *Attribution-based methods* [17, 24, 38, 57] propose modifications to the network or its training process, improving interpretable properties of post-hoc attribution methods. Finally, saliency-guided training [24, 26] design and train a model that aligns images with their saliency based masks during training enhancing recognition and interpretability properties.

Our approach aligns with attribution-based methods. Specifically, we introduce a learnable cross-attention stream, producing a representation that replaces GAP.

Attention-based architectures Attention is a mechanism introduced into convolutional neural networks to enhance their recognition capabilities [5, 35, 41]. Following the success of vision transformers (ViT) [13], fully attention-based architectures are now competitive with convolutional neural networks, while drawing inspiration from them to enhance their recognition capabilities [19, 23, 29, 50].

Unlike similar approaches combining ideas from convolutions in transformers [27, 31, 46], we propose to add an attention-based pooling mechanism in convolutional models, enhancing post-hoc interpretability properties without degrading classification accuracy.

3. Method

3.1. Preliminaries and background

Notation Let $f : \mathcal{X} \rightarrow \mathbb{R}^C$ be a classifier network that maps an input image $\mathbf{x} \in \mathcal{X}$ to a logit vector $\mathbf{y} =$

$f(\mathbf{x}) \in \mathbb{R}^C$, where \mathcal{X} is the image space and C is the number of classes. A class probability vector is obtained by $\mathbf{p} = \text{softmax}(\mathbf{y})$. The logit and probability of class c are respectively denoted by y^c and $p^c = \text{softmax}(\mathbf{y})^c$. Let $\mathbf{F}_\ell \in \mathbb{R}^{w_\ell \times h_\ell \times d_\ell}$ be the feature tensor at layer ℓ of the network, where $w_\ell \times h_\ell$ is the spatial resolution and d_ℓ the embedding dimension, or number of channels. The feature map of channel k is denoted by $F_\ell^k \in \mathbb{R}^{w_\ell \times h_\ell}$.

CAM-based saliency maps Given a class of interest c and a layer ℓ , we consider the saliency maps $S_\ell^c \in \mathbb{R}^{w_\ell \times h_\ell}$ given by the general formula

$$S_\ell^c := h \left(\sum_k \alpha_k^c F_\ell^k \right), \quad (1)$$

where α_k^c are weights defining a linear combination over channels and h is an activation function. Assuming *global average pooling* (GAP) of the last feature tensor \mathbf{F}_L followed by a linear classifier, CAM [55] is defined for the last layer L only, with h being the identity mapping and α_k^c the classifier weight connecting channel k with class c .

Self-attention Let $X_\ell \in \mathbb{R}^{t_\ell \times d_\ell}$ denote the sequence of token embeddings of a vision transformer [13] at layer ℓ , where $t_\ell := w_\ell h_\ell + 1$ is the number of tokens, including patch tokens and the CLS token, and d_ℓ is the embedding dimension. The *attention matrix* $A \in \mathbb{R}^{t_\ell \times t_\ell}$ expresses pairwise dot-product similarities between queries (Q) and keys (K), normalized by softmax over rows:

$$A = \text{softmax} \left(\frac{QK^\top}{\sqrt{d_\ell}} \right). \quad (2)$$

For each token, the *self-attention* operation is then defined as an average of all values (V) weighted by attention A :

$$\text{SA}(X_\ell) := AV \in \mathbb{R}^{t_\ell \times d_\ell}. \quad (3)$$

At the last layer L , the CLS token embedding is used as a global representation for classification as it gathers information from all patches by weighted averaging, replacing GAP. Thus, at the last layer, it is only cross-attention between CLS and the patch tokens that matters.

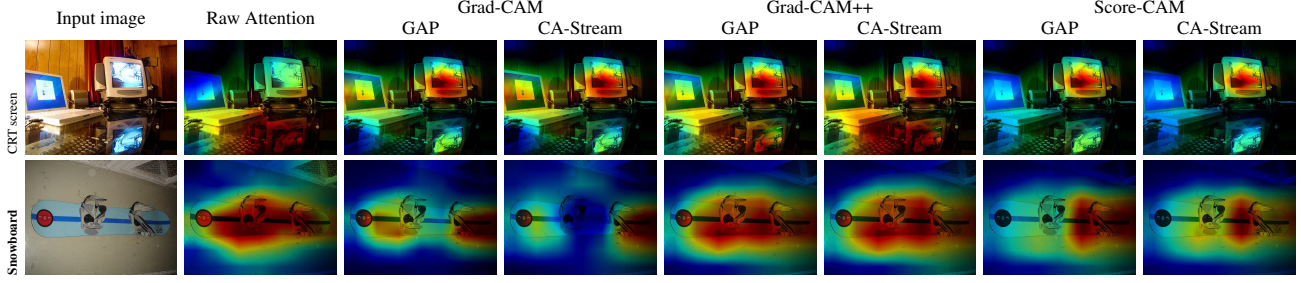


Figure 2. Comparison of saliency maps generated by different CAM-based methods, using GAP and our CA-Stream, on ImageNet images. The raw attention is the one used for pooling by CA-Stream.

3.2. Motivation

Cross-attention Let matrix $F_\ell \in \mathbb{R}^{p_\ell \times d_\ell}$ be a reshaping of feature tensor \mathbf{F}_ℓ at layer ℓ , where $p_\ell := w_\ell h_\ell$ is the number of patch tokens without CLS, and let $\mathbf{q}_\ell \in \mathbb{R}^{d_\ell}$ be the CLS token embedding at layer ℓ . By focusing on the *cross-attention* only between the CLS (query) token \mathbf{q}_ℓ and the patch (key) tokens F_ℓ , attention A (2) is now a $1 \times p_\ell$ matrix that can be written as a vector $\mathbf{a} \in \mathbb{R}^{p_\ell}$

$$\mathbf{a} = A^\top = \text{softmax} \left(\frac{F_\ell \mathbf{q}_\ell}{\sqrt{d_\ell}} \right). \quad (4)$$

Here, $F_\ell \mathbf{q}_\ell$ expresses the pairwise similarities between the global CLS feature \mathbf{q}_ℓ and the local patch features F_ℓ . Now, by replacing \mathbf{q}_ℓ by an arbitrary vector $\boldsymbol{\alpha} \in \mathbb{R}^{d_\ell}$ and writing the feature matrix as $F_\ell = (\mathbf{f}_\ell^1 \dots \mathbf{f}_\ell^{p_\ell})$, attention (4) becomes

$$\mathbf{a} = h_\ell(F_\ell \boldsymbol{\alpha}) = h_\ell \left(\sum_k \alpha_k \mathbf{f}_\ell^k \right). \quad (5)$$

This takes the same form as (1), with feature maps F_ℓ^k vectorized as \mathbf{f}_ℓ^k and the activation function defined as $h_\ell(\mathbf{x}) = \text{softmax}(\mathbf{x}/\sqrt{d_\ell})$. We thus observe the following.

Pairwise similarities between one query and all patch token embeddings in cross-attention are the same as a linear combination of feature maps in CAM-based saliency maps.

One difference between (1) and (5) is that (5) is class-agnostic, although it could be extended by using one query vector per class. For simplicity, we choose the class-agnostic form. We also choose to have no query/key projections. However, we do provide additional experiments in the appendix.

Pooling or masking We integrate an attention mechanism into a network such that making a prediction and explaining it are inherently connected. In particular, considering cross-attention only between CLS and patch tokens (4), equation (3) becomes

$$\text{CA}_\ell(\mathbf{q}_\ell, F_\ell) := F_\ell^\top \mathbf{a} = F_\ell^\top h_\ell(F_\ell \mathbf{q}_\ell) \in \mathbb{R}^{d_\ell}. \quad (6)$$

By writing the transpose of feature matrix as $F_\ell^\top = (\phi_\ell^1 \dots \phi_\ell^{p_\ell})$ where $\phi_\ell^i \in \mathbb{R}^{d_\ell}$ is the feature of patch i , this is a weighted average of the local patch features F_ℓ^\top with attention vector $\mathbf{a} = (a_1, \dots, a_{p_\ell})$ expressing the weights:

$$\text{CA}_\ell(\mathbf{q}_\ell, F_\ell) := F_\ell^\top \mathbf{a} = \sum_i a_i \phi_\ell^i. \quad (7)$$

We can think of it as as feature *reweighting* or *soft masking* in the feature space, followed by GAP.

Now, considering that \mathbf{a} is obtained exactly as CAM-based saliency maps (5), this operation is similar to occlusion (masking)-based methods [14, 15, 32, 36, 39, 47, 51] and evaluation metrics [8, 32], where a CAM-based saliency map is commonly used to mask the input image. We thus observe the following.

Attention-based pooling is a form of feature reweighting or soft masking in the feature space followed by GAP, where the weights are given by a class-agnostic CAM-based saliency map.

3.3. Cross-attention stream

Motivated by these observations, we design a *Cross-Attention Stream (CA-Stream)* in parallel to any network. It takes input features at key locations of the network and uses cross-attention to build a global image representation and replace GAP before the classifier. An example is shown in Figure 1, applied to a ResNet-based architecture.

Architecture More formally, given a network f , we consider points between blocks of f where critical operations take place, such as change of spatial resolution or embedding dimension, *e.g.* between residual blocks on ResNet. We decompose f at these points as

$$f = g \circ \text{GAP} \circ f_L \circ \dots \circ f_0 \quad (8)$$

such that features $F_\ell \in \mathbb{R}^{p_\ell \times d_\ell}$ of layer ℓ are initialized as $F_{-1} = \mathbf{x}$ and updated according to

$$F_\ell = f_\ell(F_{\ell-1}) \quad (9)$$

for $0 \leq \ell \leq L$, where p_ℓ is the number of patch tokens and d_ℓ the embedding dimension of stage ℓ . The last layer features F_L are followed by GAP and $g : \mathbb{R}^{d_L} \rightarrow \mathbb{R}^C$ is the classifier, mapping to the logit vector \mathbf{y} .

In parallel, we initialize a classification token embedding as a learnable parameter $\mathbf{q}_0 \in \mathbb{R}^{d_0}$ and we build a sequence of updated embeddings $\mathbf{q}_\ell \in \mathbb{R}^{d_\ell}$ along a stream that interacts with F_ℓ at each stage ℓ . Referring to the global representation \mathbf{q}_ℓ as *query* or CLS and to the local image features F_ℓ as *key* or patch embeddings, the interaction consists of cross-attention followed by a linear projection $W_\ell \in \mathbb{R}^{d_{\ell+1} \times d_\ell}$ to account for changes of embedding dimension between the corresponding stages of f :

$$\mathbf{q}_{\ell+1} = W_\ell \cdot \text{CA}_\ell(\mathbf{q}_\ell, F_\ell), \quad (10)$$

for $0 \leq \ell \leq L$, where CA_ℓ is defined as in (6).

Image features do not change by injecting our CA-Stream into network f . However, the final global image representation does, hence the prediction does too. In particular, at the last stage L , \mathbf{q}_{L+1} is used as a global image representation for classification, replacing GAP over F_L . Therefore, final prediction is $g(\mathbf{q}_{L+1}) \in \mathbb{R}^C$. Unlike GAP, the weights of different image patches in the linear combination are non-uniform, enhancing the contribution of relevant patches in the prediction.

Training The network f is pretrained and remains frozen while we learn the parameters of our CA-Stream on the same training set as f . The classifier is kept frozen too. Referring to (8), f_0, \dots, f_L and g are fixed, while GAP is replaced by learned weighted averaging, with the weights obtained by the CA-Stream.

Inference As it stands, the CA-Stream is an addition to the baseline architecture, which enhances the interpretability properties of a model. We thus investigate interpretability using CAM-based methods on both baseline GAP and CA-Stream in the following section.

4. Experiments

Experimental setup We train and evaluate our models on the ImageNet ILSVRC-2012 dataset [11], using the training and validation sets respectively. We experiment on pretrained and frozen ResNets [22] and ConvNeXt [30] models and provide more details in the appendix. We measure the interpretability properties of our approach by first generating saliency maps employing existing methods based on CAM (Grad-CAM [40], Grad-CAM++ [8], Score-CAM [47]) with and without CA-Stream. Then, following [51], we compute changes in the predictive power of a masked image measured by *average drop* (AD) [8] and *average gain* (AG) [51], the proportion of better explanations measured by *average increase* (AI) [8] and finally the

NETWORK	POOLING	ACC↑					
ResNet-50	GAP	74.55					
	CA	74.70					
ConvNeXt-B	GAP	83.72					
	CA	83.51					
NETWORK	ATTRIBUTION	POOLING	AD↓	AG↑	AI↑	I↑	D↓
RESNET-50	Grad-CAM	GAP	13.04	17.56	44.47	72.57	13.24
		CA	12.54	22.67	48.56	75.53	13.50
	Grad-CAM++	GAP	13.79	15.87	42.08	72.32	13.33
		CA	13.99	19.29	44.60	75.21	13.78
	Score-CAM	GAP	8.83	17.97	48.46	71.99	14.31
		CA	7.09	23.65	54.20	74.91	14.68
CONVNEXT-B	Grad-CAM	GAP	33.72	2.43	15.25	52.85	29.57
		CA	19.45	13.96	32.89	86.38	45.29
	Grad-CAM++	GAP	34.01	2.37	15.60	52.83	29.17
		CA	36.69	8.00	21.95	85.39	53.42
	Score-CAM	GAP	43.55	2.23	15.67	50.96	39.49
		CA	23.51	11.04	27.35	83.41	60.53

Table 1. *Interpretability metrics* of CA-Stream vs. baseline GAP for different networks and interpretability methods on ImageNet.

impact of different extents of masking via *insertion* (I) and *deletion* (D) [32].

Qualitative results In Figure 2, we show saliency maps obtained using either GAP and CA, as well as the raw attention representation from CA-Stream. We observe that CAM-based attributions obtained using our CA are similar to those generated with GAP. We expect this behaviour as we do not modify the model or the weighting coefficients. Since raw attention is class-agnostic, it can be used to gain insight on what the model attends to in unseen data. We iterate upon this in the appendix.

Quantitative evaluation In Table 1, we compare the interpretability properties when using our CA vs. GAP. In the appendix we provide comparisons with more models and datasets. We observe that CA-Stream provides consistent improvements over GAP in terms of AD, AG, AI and I metrics, while performing lower on D. Deletion (D) has raised concerns in previous works [9, 51]. As (D) gradually blackens pixels, *out-of-distribution* data [18, 21, 34] is produced, possibly introducing bias [37]. Moreover, non-spread attributions tend to perform better [51], which is likely the reason for lower performance.

5. Conclusion

In this work we observe that attention-based pooling in transformers is similar if not the same as forming a class agnostic CAM-based attribution. Based on this observation, we build upon this representation to mask features prior to the classification layers of a model, enhancing interpretability of existing image recognition models using GAP. Our method improves interpretability metrics while maintaining recognition performance.

References

- [1] Julius Adebayo, Justin Gilmer, Ian J. Goodfellow, and Been Kim. Local explanation methods for deep neural networks lack sensitivity to parameter values. *ICLR Work.*, 2018. 1
- [2] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS one*, 2015.
- [3] David Baehrens, Timon Schroeter, Stefan Harmeling, Motoaki Kawanabe, Katja Hansen, and Klaus-Robert Müller. How to explain individual classification decisions. *J. MLR*, 2010. 1
- [4] David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. Network dissection: Quantifying interpretability of deep visual representations. In *CVPR*, pages 6541–6549, 2017. 2
- [5] Irwan Bello, Barret Zoph, Ashish Vaswani, Jonathon Shlens, and Quoc V Le. Attention augmented convolutional networks. In *ICCV*, 2019. 2
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2021. 1
- [7] Chun-Hao Chang, Elliot Creager, Anna Goldenberg, and David Duvenaud. Explaining image classifiers by counterfactual generation. *ICLR*, 2019. 2
- [8] Aditya Chattopadhyay, Anirban Sarkar, Prantik Howlader, and Vineeth N. Balasubramanian. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. *CoRR*, abs/1710.11063, 2017. 1, 3, 4
- [9] Hila Chefer, Shir Gur, and Lior Wolf. Transformer interpretability beyond attention visualization. In *CVPR*, pages 782–791, 2021. 1, 4
- [10] Piotr Dabkowski and Yarin Gal. Real time image saliency for black box classifiers. *NIPS*, 2017. 2
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 4
- [12] saurabh desai and Harish Guruprasad Ramaswamy. Ablation-cam: Visual explanations for deep convolutional network via gradient-free localization. In *WACV*, 2020. 1
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 1, 2
- [14] Ruth Fong, Mandela Patrick, and Andrea Vedaldi. Understanding deep networks via extremal perturbations and smooth masks. In *ICCV*, 2019. 1, 3
- [15] Ruth C. Fong and Andrea Vedaldi. Interpretable explanations of black boxes by meaningful perturbation. In *ICCV*, 2017. 1, 3
- [16] Ruigang Fu, Qingyong Hu, Xiaohu Dong, Yulan Guo, Yinghui Gao, and Biao Li. Axiom-based grad-cam: Towards accurate visualization and explanation of cnns. *BMVC*, 2020. 1
- [17] Reza Ghaeini, Xiaoli Z Fern, Hamed Shahbazi, and Prasad Tadepalli. Saliency learning: Teaching the model where to pay attention. *NAACL*, 2019. 2
- [18] Tristan Gomez, Thomas Fréour, and Harold Mouchère. Metrics for saliency map evaluation of deep learning explanation methods, 2022. 4
- [19] Benjamin Graham, Alaaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Hervé Jégou, and Matthijs Douze. Levit: a vision transformer in convnet’s clothing for faster inference. In *ICCV*, 2021. 2
- [20] Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. A survey of methods for explaining black box models. *ACM Comput. Surv.*, 51(5), 2018. 1
- [21] Peter Hase, Harry Xie, and Mohit Bansal. The out-of-distribution problem in explainability and search methods for feature importance explanations, 2021. 4
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1, 4
- [23] Byeongho Heo, Sangdoon Yun, Dongyoon Han, Sanghyuk Chun, Junsuk Choe, and Seong Joon Oh. Rethinking spatial dimensions of vision transformers. In *ICCV*, 2021. 2
- [24] Aya Abdelsalam Ismail, Hector Corrada Bravo, and Soheil Feizi. Improving deep learning interpretability by saliency guided training. *NeurIPS*, 2021. 2
- [25] Peng-Tao Jiang, Chang-Bin Zhang, Qibin Hou, Ming-Ming Cheng, and Yunchao Wei. Layercam: Exploring hierarchical class activation maps for localization. *IEEE Transactions on Image Processing*, 30:5875–5888, 2021. 1
- [26] Kwang Hee Lee, Chaewon Park, Junghyun Oh, and Nojun Kwak. Lfi-cam: Learning feature importance for better visual explanation. In *ICCV*, 2021. 2
- [27] Liangzhi Li, Bowen Wang, Manisha Verma, Yuta Nakashima, Ryo Kawasaki, and Hajime Nagahara. Scouter: Slot attention-based classifier for explainable image recognition. In *ICCV*, 2021. 2
- [28] Zachary C Lipton. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3):31–57, 2018. 1
- [29] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 2
- [30] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *CVPR*, 2022. 1, 4
- [31] Zhiliang Peng, Wei Huang, Shanzhi Gu, Lingxi Xie, Yaowei Wang, Jianbin Jiao, and Qixiang Ye. Conformer: Local features coupling global representations for visual recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 367–376, 2021. 2
- [32] Vitali Petsiuk, Abir Das, and Kate Saenko. Rise: Randomized input sampling for explanation of black-box models. *BMVC*, 2018. 1, 3, 4

- [33] Jason Phang, Jungkyu Park, and Krzysztof J Geras. Investigating and simplifying masking-based saliency methods for model interpretability. *arXiv preprint arXiv:2010.09750*, 2020. 2
- [34] Luyu Qiu, Yi Yang, Caleb Chen Cao, Jing Liu, Yueyuan Zheng, Hilary Hei Ting Ngai, Janet Hsiao, and Lei Chen. Resisting out-of-distribution data problem in perturbation of xai, 2021. 4
- [35] Prajit Ramachandran, Niki Parmar, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jon Shlens. Stand-alone self-attention in vision models. *Advances in neural information processing systems*, 32, 2019. 2
- [36] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "why should i trust you?": Explaining the predictions of any classifier. In *SIGKDD*, 2016. 1, 3
- [37] Yao Rong, Tobias Leemann, Vadim Borisov, Gjergji Kasneci, and Enkelejda Kasneci. A consistent and efficient evaluation strategy for attribution methods, 2022. 4
- [38] Andrew Slavin Ross, Michael C Hughes, and Finale Doshi-Velez. Right for the right reasons: Training differentiable models by constraining their explanations. *IJCAI*, 2017. 2
- [39] Karl Schulz, Leon Sixt, Federico Tombari, and Tim Landgraf. Restricting the flow: Information bottlenecks for attribution. *arXiv preprint arXiv:2001.00396*, 2020. 1, 2, 3
- [40] Ramprasaath R. Selvaraju, Abhishek Das, Ramakrishna Vedantam, Michael Cogswell, Devi Parikh, and Dhruv Batra. Grad-cam: Why did you say that? visual explanations from deep networks via gradient-based localization. *CoRR*, abs/1610.02391, 2016. 1, 4
- [41] Zhuoran Shen, Irwan Bello, Raviteja Vemulapalli, Xuhui Jia, and Ching-Hui Chen. Global self-attention networks for image recognition. *arXiv preprint arXiv:2010.03019*, 2020. 2
- [42] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *ICLR Workshop*, 2014. 1
- [43] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda B. Viégas, and Martin Wattenberg. Smoothgrad: removing noise by adding noise. *CoRR*, abs/1706.03825, 2017.
- [44] Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin A. Riedmiller. Striving for simplicity: The all convolutional net. *ICLR*, 2015.
- [45] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *ICML*, 2017. 1
- [46] Hugo Touvron, Matthieu Cord, Alaaeldin El-Nouby, Piotr Bojanowski, Armand Joulin, Gabriel Synnaeve, and Hervé Jégou. Augmenting convolutional networks with attention-based aggregation. *arXiv preprint arXiv:2112.13692*, 2021. 2
- [47] Haofan Wang, Mengnan Du, Fan Yang, and Zijian Zhang. Score-cam: Improved visual explanations via score-weighted class activation mapping. *CoRR*, abs/1910.01279, 2019. 1, 3, 4
- [48] Mike Wu, Michael Hughes, Sonali Parbhoo, Maurizio Zazzi, Volker Roth, and Finale Doshi-Velez. Beyond sparsity: Tree regularization of deep models for interpretability. In *Proceedings of the AAAI conference on artificial intelligence*, 2018. 2
- [49] Mike Wu, Sonali Parbhoo, Michael Hughes, Ryan Kindle, Leo Celi, Maurizio Zazzi, Volker Roth, and Finale Doshi-Velez. Regional tree regularization for interpretability in deep neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 6413–6421, 2020. 2
- [50] Tete Xiao, Mannat Singh, Eric Mintun, Trevor Darrell, Piotr Dollár, and Ross Girshick. Early convolutions help transformers see better. *Advances in Neural Information Processing Systems*, 34:30392–30400, 2021. 2
- [51] Hanwei Zhang, Felipe Torres, Ronan Sicre, Yannis Avrithis, and Stephane Ayache. Opti-cam: Optimizing saliency maps for interpretability. *arXiv preprint arXiv:2301.07002*, 2023. 3, 4
- [52] Quanshi Zhang, Ying Nian Wu, and Song-Chun Zhu. Interpretable convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8827–8836, 2018. 2
- [53] Yu Zhang, Peter Tiño, Aleš Leonardis, and Ke Tang. A survey on neural network interpretability. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 5(5):726–742, 2021. 1
- [54] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Object detectors emerge in deep scene cnns. *ICLR*, 2015. 2
- [55] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, 2016. 2
- [56] Bolei Zhou, David Bau, Aude Oliva, and Antonio Torralba. Interpreting deep visual representations via network dissection. *Trans. PAMI*, 41(9), 2018. 2
- [57] Hao Zhou, Keyang Cheng, Yu Si, and Liuyang Yan. Improving interpretability by information bottleneck saliency guided localization. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*. BMVA Press, 2022. 2
- [58] Konrad Zolna, Krzysztof J. Geras, and Kyunghyun Cho. Classifier-agnostic saliency map extraction. *CVIU*, 196: 102969, 2020. 2