FovEx: Human-Inspired Explanations for Vision Transformers and Convolutional Neural Networks

Mahadev Prasad Panda[1] · Matteo Tiezzi[2] · Martina Vilas[3] · Gemma Roig[4] · Bjoern M. Eskofier[1,5] · Dario Zanca[1]

Abstract

Explainability in artificial intelligence (XAI) remains a crucial aspect for fostering trust and understanding in machine learning models. Current visual explanation techniques, such as gradient-based or class-activation-based methods, often exhibit a strong dependence on specific model architectures. Conversely, perturbation-based methods, despite being model-agnostic, are computationally expensive as they require evaluating models on a large number of forward passes. We introduce Foveation-based Explanations (FovEx), a novel XAI method inspired by human vision, which combines biologically inspired foveation-based transformations with gradient-driven overt attention to iteratively select locations of interest. These locations are selected to maximize the performance of the model to be explained with respect to the downstream task and then combined to generate an attribution map. We provide a thorough evaluation with qualitative and quantitative assessments on established benchmarks. Our method achieves state-of-the-art performance on both transformers (on 4 out of 5 metrics) and convolutional models (on 3 out of 5 metrics), demonstrating its versatility among various architectures. Furthermore, we show the alignment between the explanation map produced by FovEx and human gaze patterns (+14% in NSS compared to RISE, +203% in NSS compared to GradCAM). This comparison enhances our confidence in FovEx's ability to close the interpretation gap between humans and machines.

✉ Dario Zanca
  dario.zanca@fau.de

  Mahadev Prasad Panda
  mahadev.prasad.panda@fau.de

  Matteo Tiezzi
  matteo.tiezzi@iit.it

  Martina Vilas
  martinagonzalezvilas@gmail.com

  Gemma Roig
  roig@cs.uni-frankfurt.de

  Bjoern M. Eskofier
  bjoern.eskofier@fau.de

[1] Department AIBE, FAU Erlangen-Nürnberg, Erlangen, Germany

[2] PAVIS, Istituto Italiano di Tecnologia (IIT), Genova, Italy

[3] Ernst Strüngmann Institute for Neuroscience, Frankfurt, Germany

[4] Goethe-Universität Frankfurt am Main, Frankfurt, Germany

[5] Institute of AI for Health, Helmholtz Zentrum München, Munich, Germany

1 Introduction

In recent years, deep learning has made remarkable strides in revolutionizing computer vision, particularly in safety-critical domains such as medical imaging (Fujiyoshi et al., 2019; Ouyang et al., 2020; Wang et al., 2022), autonomous driving (Jing et al., 2022; Najibi et al., 2022), industrial automation (Iqbal et al., 2019; Liang et al., 2017; Maschler & Weyrich, 2021), or security and surveillance (Gruosso et al., 2021; Singh et al., 2020; Xu, 2021). However, as these models become increasingly more complex, the lack of understanding of their decision-making processes poses significant challenges(Achtibat et al., 2023; Goodman & Flaxman, 2017; Longo et al., 2024). To address these challenges, there is a growing need for XAI methods aiming at ensuring transparency and interpretability to the black-box nature of deep learning models.

While a variety of explanation methods have been developed for vision models (Chefer et al., 2021a, b; Chattopadhay et al., 2018; Selvaraju et al., 2017), these approaches are often tailored to specific architectures and lack universality. Grad-CAM and its derivations (Chattopadhay et al., 2018; Selvaraju et al., 2017) have been originally described as effective class-specific XAI methods to compute gradient-weighted feature maps from the last layer of convolutional architectures, highlighting relevant regions in the input. Although the GradCAM method can be extended to vision transformers (through reshaping the feature maps and gradients from the deepest layers), the performance of this approach is adversely affected by certain architectural attributes of vision transformers, such as skip connections, non-local self-attention mechanisms, and unstable gradients (Choi et al., 2023). On the other side, XAI methods for vision transformers (Abnar & Zuidema, 2020; Chefer et al., 2021a, b) are often tailored to transformer-specific characteristics, such as attention weights or class tokens, making their application to convolutional-based models unfeasible. Therefore, there is a pressing need for *model-agnostic* XAI methods, i.e., an approach that can work on any model architecture without the need for changes or adaptations. This can ensure comparability in explanations across different architectures and ensure more reliable interpretations of deep neural networks.

While current XAI methods provide insights into model decision-making processes, their quality often falls short when it comes to human understanding as they lack contextual aspects that make such explanations understandable to humans (Hsiao & Chan, 2023; Morrison et al., 2023). Incorporating human-inspired constraints into XAI frameworks can enhance the quality of explanations and make them more aligned with human perception. One such fundamental aspect is foveated vision: humans' highest visual acuity occurs at the center of the visual field (fovea), while peripheral vision has a lower resolution, underlying the way humans prioritize details in a specific area. Recent work (Han et al., 2020; Volokitin et al., 2017) has demonstrated the advantages of incorporating such constraints into vision models.

Deza and Konkle (2020) show increased i.i.d. generalization as a computational consequence emerging with foveated processing. Location-dependent computation based on foveation have demonstrated efficiency and avoidance of spurious correlation from data, both for convolutional (Tiezzi et al., 2022) and visual transformer (Jonnalagadda et al., 2021) models. Foveation priors are effective in generating visual scanpath (Schwinn et al., 2022). These results together demonstrate how introducing biological constraints in artificial neural networks both increases alignment with the human counterpart and fosters model performances. We believe that the aforementioned concepts and intuitions can open a promising novel avenue in the field of explainability.

In this paper, we address the challenges of current explainability approaches and introduce a novel method, Foveation-based Explanations (FovEx), drawing inspiration from two key aspects of human vision: foveation (Bringmann & Wiedemann, 2021; Hendrickson & Yuodelis, 1984) and overt attention(Tsotsos, 2021; Wright, 1998). The final goal of FovExis to determine explanations, in the form of attribution maps, on the output of a backbone model processing a given input pattern. To do so, input samples undergo a differentiable transformation simulating a foveation mechanism. Gradient information is then utilized to facilitate selective visual attention through an iterative, post-hoc, human-like exploration of the input image, identifying the most relevant regions. Such relevant regions are combined appropriately to generate attribution maps. The FovEx-generated explanation maps, compared in Figure 1 against state-of-the-art methods (further details in the remainder of the paper), achieve state-of-the-art performance in common XAI metrics and have a better alignment with human gaze, as per experimental results.

To summarize, we delineate our contributions as follows:

- We introduce FovEx, a novel explanation method for DNNs that incorporates the biological constraint of human foveated vision. FovEx extracts human-aligned visual explanations of the predictions in a post-hoc fashion, i.e., without introducing any architectural modification to the underlying predictor.
- We demonstrate the effectiveness of FovEx through qualitative and quantitative evaluations for both convolutional and transformer-based models. We compare our approach to class-activation-based and perturbation-based XAI methods and demonstrate state-of-the-art performance for different model architectures.
- We show that FovEx explanations enhance human interpretability via a quantitative investigation of the correlation between human gaze patterns and the explanation maps of DNNs generated by FovEx.

2 Related Work

A variety of local post-hoc XAI methods exist for supervised image classification models (Ali et al., 2023). These techniques can be broadly categorized into (i) backpropagation or gradient-based methods, (ii) Class Activation Map (CAM)-based methods, and (iii) perturbation-based methods (Nguyen et al., 2023; Wagner et al., 2019). In this section, we first introduce the main literature for these three categories, and then provide concise insights into specific methods tailored for the interpretability of transformer-based models.

Fig. 1 ViT-B/16 attribution maps. Explanation maps generated by FovEx (second column) and competitors for the ViT-B/16 predictor. Red crosses on explanation maps in the FovExcolumn denote fixation locations

Gradient-based Explanation Methods. To generate explanations, gradient-based XAI approaches utilize the gradient of a pre-trained black-box model's output with respect to input features, i.e., image pixels (Nielsen et al., 2022). The seminal work by Simonyan et al. (2014) constructs saliency maps by computing gradients of the non-linear class score function with respect to the input image. However, gradients frequently result in noisy visualizations. To address this issue, (Smilkov et al., 2017) refines the explanation map by averaging over multiple saliency maps for a single input image. These saliency maps correspond to noisy duplicates of the input image, built by introducing random Gaussian noise with a zero mean into the original image. Layer-wise Relevance Propagation (LRP) (Bach et al., 2015) generates an explanation map by propagating fixed predefined decomposition rules from the output layer of a black-box model to the input layer. Softmax Gradient Layer-wise Relevance Propagation (SGLRP) (Iwana et al., 2019) utilizes the gradient of the softmax function to propagate decomposition rules, aiming to address the class-agnostic nature observed in vanilla LRP. Unlike gradient-based methods, our proposed approach uses gradient information exclusively to enable a human-like exploration of the input image to generate locations of interest.

CAM-based Explanation Methods. A class activation map (CAM) (Zhou et al., 2016) reveals the important areas in an image that a convolutional neural network relies on to recognize a particular class. CAM leverages the Global Average Pooling (GAP) layer and the top-most fully connected layer of convolution-based classification networks. Even if CAM produces class-discriminating saliency maps, it is fully dependant on architectural families and constraints. Grad-CAM (Selvaraju et al., 2017) extends CAM by incorporating gradient information. It computes the gradient of the predicted class score with respect to the feature maps of the last convolutional layer. These gradients act as the weights of each feature map for the target class. GradCAM++ (Chattopadhay et al., 2018) builds on GradCAM by employing pixel-wise weights rather than a single weight for a forward feature map of the final convolution layer. This enhancement enables GradCAM++ to preserve multiple instances of similar objects in the final explanation map. To further improve on the gradient-based CAM techniques, (Wang et al., 2020) uses class score as weights and (Lee et al., 2021) employs relevance score as weights in the explanation generation process. In contrast to CAM-based methods, our method is architecture agnostic, making it applicable to convolution and transformer-based models without any modifications.

Perturbation-based Explanation Methods. In the context of supervised image classification, perturbation-based explanation methods involve techniques that generate explanations by directly manipulating the pixels in the input image and observing the resultant changes in the output of the black-box model (Ivanovs et al., 2021). Various existing explanation techniques belong to this category (Fong & Vedaldi, 2017; Fong et al., 2019; Li et al., 2021; Petsiuk et al., 2018, 2021; Wagner et al., 2019). Perturbation-based explanation methods are particularly versatile, holding the potential for application across various model architectures. However, these approaches exhibit high computational complexity as they rely on the computation of a large number of perturbations and model inference steps. Unlike traditional perturbation methods, which rely on random or exhaustive sampling of input variations, our approach strategically inte-

grates gradient guidance to identify optimal focus locations. Our novel hybrid design not only leverages the strengths of both perturbation-based and gradient-driven techniques but also enhances the interpretability and relevance of the generated attribution maps. By guiding perturbations with gradient information, our method significantly reduces computational overhead, achieving an efficiency improvement of an order of magnitude compared to other perturbation-based approaches, see Section 4.3. Recent work has also explored black-box explanation techniques specifically tailored for object detectors (Petsiuk et al., 2018).

Explaining Transformers. While there is a variety of explanation methods available for convolution-based models, the range of methods for transformers is relatively limited. Chefer et al. (2021a) contribute to this domain by generating explanation maps for transformers with a method that combines LRP and gradient-based approaches. Chefer et al. (2021b) extends the application of the method proposed in Chefer et al. (2021a) to provide explanations for any type of transformer model. Additionally, Vilas et al. (2023) propose an approach that quantifies how different regions of an image can contribute to producing a class representation in intermediate layers, using attention and gradient-based information. However, it is important to note that these methods cannot be seamlessly employed in a plug-and-play manner without making adjustments to pre-existing model implementations. Conversely, in this paper, we propose a novel method that can be applied to both convolution-based and transformer-based architectures without altering the neural architectures for computing the explanation.

3 Foveation-based Explanation: FovEx

Let us consider a black-box predictor $b(\cdot|\theta)$ defined for classification problems, that, without any loss in generality, we assume to be a neural network with learnable parameters θ. The predictor is a function $b : \mathcal{X} \in \mathbb{R}^i \mapsto \mathcal{Y}$, that maps data instances x from the input space $\mathcal{X} \in \mathbb{R}^i$ to the prediction y in the target space $\mathcal{Y}$, containing the different labels to which the input data pattern can belong. We denote with $y = b(x|\theta)$ the prediction y yielded by the predictor on the input pattern x. We denote with θ_D the case in which the learnable parameters θ have been tuned on a training dataset D.

The goal of FovEx is to extract a human-understandable visual explanation of the decisions taken by the predictor b, in a post-hoc fashion, i.e., without introducing any architectural modification to the underlying predictor b. Formally, FovEx is a function

$$E = \text{FovEx}(b, x) \tag{1}$$

where the output E denotes the attribution map for the predictor b associated to the input x. FovEx in turn consists of three fundamental operations, i.e., (i) a differentiable foveation mechanism, (ii) a gradient-based attention mechanism, and (iii) an attribution map generation process. A schematic illustration of the method is given in Figure 2. In the following, we give a formal definition of each operation.

Differentiable Foveation. Biological foveated vision is characterized by a central area of fine-grained processing (i.e., the fovea) and a coarser peripheral area. We draw inspiration from Schwinn et al. (2022) to design a differentiable foveation mechanism. Let x be an input image and f_t be the current coordinates of the focus of attention at time t. First, we define a coarse version of x, denoted by $\bar{x}$, by convolving it with a Gaussian kernel, i.e.,

$$\bar{x} = x * \mathcal{G}(0, \sigma_b^2) \tag{2}$$

where $*$ represents the convolution operation, and $\mathcal{G}(0, \sigma_b^2)$ denotes the Gaussian kernel with a mean of 0 and a standard deviation of σ_b. The parameter σ_b controls the amount of blurring in the periphery of the image. From a biological standpoint, $\bar{x}$ can be regarded as the fundamental information obtained by peripheral vision within the initial milliseconds of stimulus presentation. A foveated input image x_Φ is obtained as a weighted sum of the original input x and the coarse version of the input $\bar{x}$, i.e.,

$$x_\Phi = \Phi(x, f_t) = \mathcal{W}(t) \cdot x + (1 - \mathcal{W}(t)) \cdot \bar{x} \tag{3}$$

where $\cdot$ denotes the pixel-wise multiplication. In equation 3, $\mathcal{W}(t) = \mathcal{G}\left(f_t, \sigma_f^2\right)$ stands for the pixel-wise weighting factor defined as the Gaussian blob $\mathcal{G}\left(f_t, \sigma_f^2\right)$ with a mean of f_t and a standard deviation of σ_f, and $\Phi(\cdot)$ represents the foveation function. It is important to notice that, since all operations are differentiable, it makes the propagation of gradient information necessary for subsequent steps feasible. Additionally, the transformation described above introduces noise according to a foveal distribution $\mathcal{W}(t)$, perturbing the original input. Notice that the drop in visual sensitivity across the human visual field is not exactly Gaussian (Bringmann & Wiedemann, 2021; Hendrickson & Yuodelis, 1984) and can vary with direction and eccentricity. However, Gaussian approximations are commonly used as effective first-order approximations of such drop in sensitivity, especially in low-resolution settings. This approach is well-established in practical applications such as foveated rendering (Malkin et al., 2020), foveated image processing (Pamplona & Bernardino, 2009; Wu et al., 2018), and blur simulation (Deza & Konkle, 2020) because it balances computational efficiency with high accuracy.
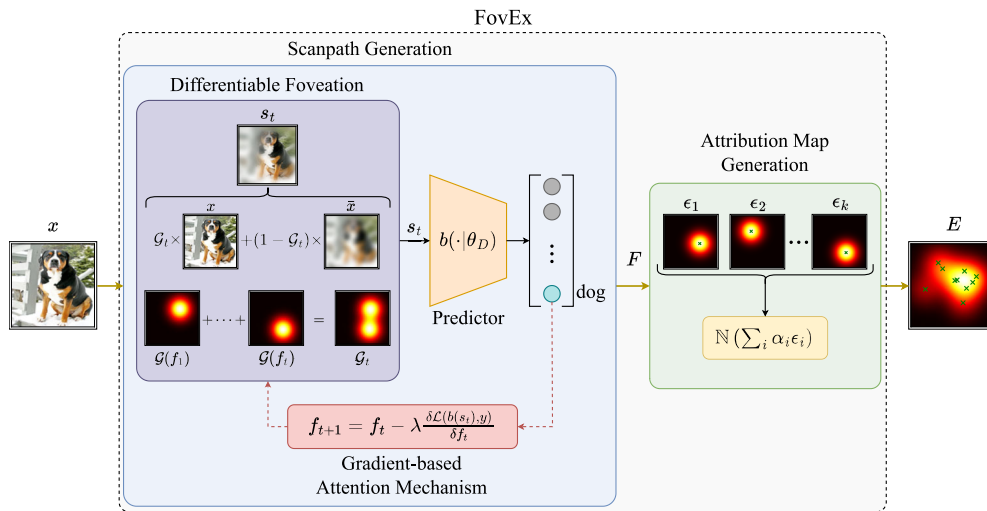
Fig. 2 The proposed FovEx. Given an input image x, FovEx produces an attribution map E for a predictor b trained on a dataset D. The image x undergoes a differentiable foveation process yielding a transformed input (s_t) for the predictor b. The loss function $\mathcal{L}$, computed based on predicted class scores and ground truth y, is exploited by an attention mechanism for the generation of a sequence of fixations $(f_1, f_2, \cdots, f_t)$ (referred to as scanpath F). The resultant scanpath F is employed to build a weighted linear combination of individual saliency maps ϵ_i associated with each fixation point f_i, leading to the final attribution map E

Gradient-based Attention Mechanism. The foveation mechanism allows for sequential exploration of the given input image. The next location of interest (i.e., the next fixation point) will depend on the state s_t generated by all previous fixation locations, which can be regarded as the system's memory. The s_t is obtained by cumulating Gaussian blobs to gradually expand the region of good visual fidelity after each fixation point, i.e.,

$$s_t = s(x, f_t) = \mathcal{G}_t \cdot x + (\mathbb{1} - \mathcal{G}_t) \cdot \bar{x} \qquad (4)$$

where $\mathcal{G}_t = \sum_{j=0}^{t} \beta^j \mathcal{G}_{t-j}\left(f_j, \sigma_f^2\right)$ symbolizes the cumulative Gaussian blob, whereas $\mathbb{1}$ refers to a square matrix of ones. The forgetting factor $0 \leq \beta \leq 1$ regulates how much information is retained from previous fixations.

Let $\mathcal{L}(b(s_t), y)$ be the loss function at time t, e.g., calculated as the distance between the output predicted by b for the current state s_t and the target class y_t. The next fixation locations are dynamically adjusted to minimize the loss function $\mathcal{L}$ with respect to the current fixation location f_t, i.e.,

$$f_{t+1} = f_t - \lambda \frac{\delta \mathcal{L}}{\delta f_t} \qquad (5)$$

where the hyperparameter λ determines the step size at each iteration during optimization. The optimization technique iterates until a specified number of optimization steps (OS) have been performed, ensuring successful convergence. The influence of OS is discussed in the ablation studies, in Section 4.4. Random restarts (RR) are implemented when optimization fails to yield improvements in the loss function after a specified number of steps OS, serving as a strategy to escape local minima. New fixations can be generated for an arbitrary number N of steps, resulting in a sequence of N fixations points, also called a *scanpath*

$$F = (f_1, ..., f_N) \qquad (6)$$

Attribution Map Generation. At this point, we have generated a sequence of N fixation points based on an input and a task model. Each fixation point f_i can be associated with a saliency map ϵ_i, describing a 2D Gaussian distribution with a mean at f_i and a standard deviation σ_ϵ, where σ_ϵ is set to match the standard deviation of the Gaussian blob (σ_f). The final attribution map E, functioning as an explanation for the predictor b, is obtained as a weighted linear combination of the individual saliency maps associated with each fixation point, i.e.,

$$E = \mathcal{N}\left(\sum_{i=1}^{k} \alpha_i \epsilon_i\right) \qquad (7)$$

In Equation 7, α_i denotes the weighting factor for saliency map ϵ_i and $\mathcal{N}(\cdot)$ represents min-max normalization. The weights α_i determine the contribution of each fixation to the final saliency map. In our experiments, we set $\alpha_i = 1$, $\forall I \in \{1, ..., N\}$, as different weighting schemes did not improve the quality of explanations on a validation set.

4 Experiments

We conducted a comprehensive set of experiments to compare the performances of the proposed FovExagainst state-of-the-art (SOTA) models and to showcase its ability to be agnostic to the architecture of the predictor b. We assessed FovEx's performances in different scenarios, ranging from qualitative inspections and quantitative assessments in the common testbed of ImageNet-1K validation set (Russakovsky et al., 2015) to the correlation analysis of the generated attribution maps to human gaze. Additionally, we compared FovEx's computational complexity against SOTA methods and performed in-depth model ablation studies. All our experiments were performed in a Linux environment, using an NVIDIA RTX 3080 GPU, and the implementation code can be found at https://github.com/mahadev1995/FovEx.

4.1 ImageNet-1K

Setup & Data. We selected two representative classification models as the predictor b to be explained. In particular, we focused on a ResNet-50[1] (He et al., 2016) and a Vision Transformer (ViT-B/16)[2] (Dosovitskiy et al., 2020), that have been recently classified as foundation models (Bommasani et al., 2021). The ResNet-50 model is pre-trained on the ImageNet-1K (Russakovsky et al., 2015) dataset, while the ViT-B/16 model is pre-trained on ImageNet-21K (Ridnik et al., 2021) dataset and fine-tuned on ImageNet-1K (Russakovsky et al., 2015) dataset. We assessed the model performances on a subset of 5000 images from the ImageNet-1K (Russakovsky et al., 2015) validation set, randomly sampled among the ones correctly classified by the predictor b in order to measure the contribution of the method exactly, as pointed out by recent literature (Lee et al., 2021). The images are resized to a resolution of 224×224 pixels and normalized in the range of $(0, 1)$. Performance on additional models is presented in Appendix A.

Metrics. We report model performances focusing on the faithfulness and localization attributes of the explanation maps. Faithfulness measures the extent to which explanation maps correspond to the behavior of the black-box model. We report Avg. % drop (lower is better) and Avg. % increase introduced in Chattopadhay et al. (2018) (higher is better), as well as the DELETE (lower is better) and INSERT metrics (higher is better) proposed by Petsiuk et al. (2018). The Avg. % drop metric assesses the shift in confidence between two scenarios: one with the entire image as input and another with input limited to the regions highlighted by the explanation map. The Avg. % increase metric quantifies

instances across the dataset where the model's confidence rises when only the highlighted regions from the explanation map are considered. The DELETE metric is used to evaluate the decrease in predicted class probability when removing pixels with decreasing importance, according to the explanation map. On the other hand, the INSERT metric measures the increase in estimated likelihood when adding the essential pixels to the input, ordered from most to least important as indicated by the attribution map. Localization measures an explanation map's capability to focus on a specific region of interest. Although good performance on localization may not imply a good explanation, it can provide interesting insights nonetheless because explaining a model's localization decisions can help identify patterns in the data the model has learned (Lucieri et al., 2020). We report the Energy-Based Pointing Game (EBPG) metric (Wang et al., 2020) (higher is better). The EBPG metric calculates the energy of attribution maps within a predefined bounding box of the target class.

Compared Models & Architecture Details. We compared the performances attained by FovExagainst various SOTA XAI techniques. When considering the ViT-B/16 predictor, we report performances obtained by grad-CAM (Selvaraju et al., 2017), Meaningful Perturbation (Mean. Pert.) (Fong & Vedaldi, 2017), RISE (Petsiuk et al., 2018), GAE (Chefer et al., 2021b), and class embedding projection (Cls. Emb.) (Vilas et al., 2023). We remark that perturbation-based methods (Mean. Pert., RISE) have not been previously tested with Transformer architectures. In the case of a ResNet-50 predictor, we consider gradCAM, grad-CAM++ (Chattopadhay et al., 2018), Mean. Pert., and RISE as competitors. In all the settings, we also report a baseline technique for sanity check, referred to as RandomCAM. In particular, randomCAM generates class activation maps for a randomly selected class regardless of the black-box model's prediction. For competitors, we utilized hyper-parameters following the respective original implementations. When considering FovEx, we adhere to the parameters suggested in Schwinn et al. (2022), which are biologically plausible, as they are designed to mimic human vision. The effect of different choices for the hyperparameters is discussed in our ablation study discussed in Section 4.4.

Quantitative Inspection. We report in Table 1 the results obtained for the ViT-B/16 predictor. FovExoutperforms all the competitors across all metrics, with the exception of DELETE, where the Cls. Emb. method excels. RISE secures the second-best position in Avg. % drop, Avg. % increase, and EBPG metrics, while gradCAM achieves the second-best performance in DELETE and INSERT. Similar conclusions can be drawn when we consider b to be a ResNet-50, as reported in Table 2. FovExovercomes other methods in Avg. % drop, Avg. % increase, and INSERT, while delivering competitive results in DELETE and EBPG. We remark that while Mean. Pert. performs better than FovExwhen con-

---

Table 1 ViT-B/16 quantitative evaluation. Average metrics on the considered subset of ImageNet-1K validation dataset. The best-performing model is in bold, and the second-best is underlined

Eval. Name	FovEx	grad CAM	GAE	Cls. Emb.	Mean. Pert.	RISE	random CAM
Avg. % drop (↓)	**13.970**	40.057	86.207	34.862	29.753	<u>15.673</u>	80.714
Avg. % increase (↑)	**30.389**	11.469	0.799	13.329	20.549	<u>22.189</u>	1.789
Delete (↓)	0.240	<u>0.157</u>	0.172	**0.155**	0.200	0.158	0.395
Insert (↑)	**0.840**	<u>0.818</u>	0.806	0.817	0.674	0.782	0.682
EBPG (↑)	**47.705**	41.667	39.812	39.350	40.646	<u>42.633</u>	35.708

Table 2 ResNet-50 quantitative evaluation. Average metrics on the considered subset of ImageNet-1K validation dataset. The best-performing model is in bold, and the second-best is underlined

Eval. Name	FovEx	grad CAM	grad CAM++	Mean. Pert.	RISE	random CAM
Avg. % drop (↓)	**11.780**	21.718	19.863	85.973	<u>11.885</u>	61.317
Avg. % increase (↑)	**61.849**	43.669	45.069	4.700	<u>55.489</u>	16.729
Delete (↓)	0.151	0.108	0.113	**0.082**	<u>0.100</u>	0.212
Insert (↑)	**0.374**	0.368	0.361	0.280	<u>0.372</u>	0.287
EBPG (↑)	46.977	**48.658**	<u>47.412</u>	42.725	43.312	38.118

sidering Delete, it exhibits worse overall performances in all the other considered metrics. GradCAM is the top performer when considering EBPG, followed by gradCAM++. RISE secures the second position in all other metrics, except for EBPG.

These results showcase better faithfulness and localization capabilities of the proposed FovExas compared to other methods, independent of the underlying model architecture. The lower performance of FovExon the Delete metric can be attributed to its underlying optimization scheme. FovExis based on a preservation-oriented approach, where the optimization seeks to identify image regions that, when retained, enable the model to maintain accurate classification. This design naturally favors the Insert metric, which aligns more closely with FovEx's optimization objective. However, this emphasis on preserving informative regions introduces a trade-off, potentially reducing performance on the Delete metric. Since the optimization is geared toward retaining the most crucial areas for maintaining prediction accuracy—rather than identifying regions whose removal most impairs the model—FovExmay be penalized under the Delete evaluation.

Qualitative Inspection. To better assess the explanations provided by the proposed FovEx, we conducted a qualitative evaluation comprising both a (i) visual comparison concerning the considered competitors and (ii) in-depth investigation regarding class-specific explanations. Figure 1 illustrates a comparison of attribution maps generated on several (rows) input samples (first column) by the proposed FovEx(second column) and the other considered competitors (subsequent columns), in the case of the ViT-B/16 predictor. FovExgen-

erates explanation maps that are spread over the whole object of interest without any blemish or stray spot, such as the ones produced by gradCAM, GAE, and Cls. Emb. An explanation of such phenomenon is given by Darcet et al. (2023), demonstrating the presence of artifacts in the feature map of models, corresponding to high-norm tokens appearing during inference in unimportant background areas of images, repurposed for internal computations. Indeed, we remark how the issue is solved when using FovExas the explanation method. Similarly, RISE yields noise-free explanation maps, but we remark that its computational time is significantly higher compared to FovEx(see Section 4.3).

To better assess the quality of attribution maps generated by FovExwhen dealing with convolution-based architectures, we report in Figure 3 a visual comparison of attribution maps generated by FovEx (second column) and other SOTA methods for the ResNet-50 predictor. More visualizations of attribution maps for different models are given in Appendix D. Also in this setting, FovExis capable of generating attribution maps that are focused on the main object class contained in the processed image. The attribution map correctly spans over the object of interest, thanks to the exploration carried on by the attention scanpath, differently from competitors (RISE, gradCAM) that tend to attribute the class prediction to certain object parts (bird head).

Additionally, we tested FovEx's ability to localize class-discriminative information and capture fine-grained details. In particular, we selected images containing multiple annotated categories and we generated attribution maps for each of the available classes with FovEx. We report in Figure 4-top the attribution maps obtained with ViT-B/16 as the predictor

Fig. 3 ResNet-50 attribution maps. Explanation maps generated by FovEx and competitors for the ResNet-50 model. Red crosses on explanation maps in the FovExcolumn denote fixation locations
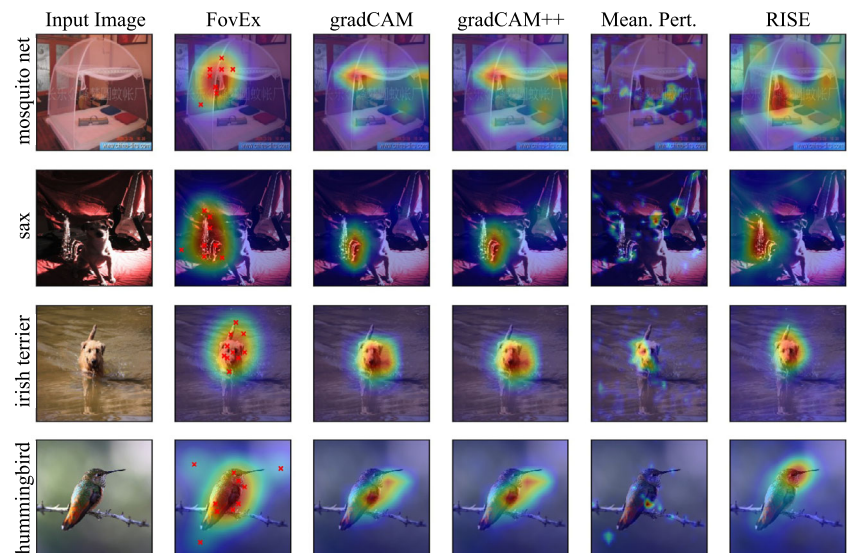
Fig. 4 Class-discriminative evaluation. Class-specific attribution maps generated from FovEx, using ViT-B/16 (top) and ResNet-50 (bottom) as the predictor
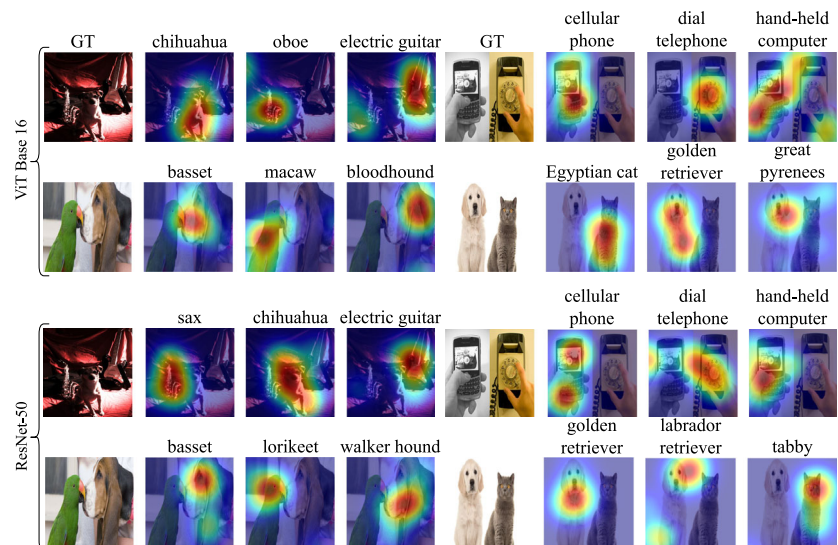
when generating explanations for different categories (indicated in the column header). Attribution maps generated from FovExcan discern the desired categories, localizing the relevant category inside the image with a good fine-grained definition of the pixels belonging to the object. Similar conclusions can be drawn when considering ResNet-50 as the predictor, as depicted in Figure 4-bottom.

We report the FovEx performance when considering three additional predictors (b) such as ConvNeXt (Liu et al., 2022), ViT-B/16 and ViT-B/32 from `torchvision`[3] in the Appendix A. In summary, FovEx continues to outperform other SOTA methods in the majority of the evaluation metrics considered for the aforementioned models.

---

[3] Models from `torchvision` are trained on ImageNet-1K, while ViT-B/16 used in Section 4.1 is trained on ImageNet-21K and finetuned on ImageNet-1K. Source: https://pytorch.org/vision/stable/models.html

4.2 Human Gaze Correlation

Setup & Data. Generating attribution maps that are similar to the ones produced by human-gaze for image classification can enhance visual plausibility (Lai et al., 2020; Qi et al., 2023). Here, we investigate the correlation of attribution maps generated by FovExand other competitors with the ones obtained from human gaze for image free-viewing task. The goal is to investigate the possibility of generating explanation maps correlated to the general human gaze while being faithful to a black-box model's decision-making. We exploit the MIT1003 dataset (Judd et al., 2009), composed of 1003 images that are complemented by corresponding human attention maps, collected in a controlled setting.

Metrics. To quantitatively assess the similarity between the attribution maps generated by the black-box model and the human attention map, we employ the Normalised Scanpath

Table 3 Human-gaze correlation. Quantitative evaluation results from human gaze experiment. We report NSS and AUCJ metrics (higher is better, see the main text for further details)

Eval. Name	FovEx	grad CAM	grad CAM++	Cls. Emb.	GAE	Mean. Pert.	RISE
NSS (↑)	**0.7160**	0.2357	0.0372	0.1231	0.4120	0.6197	0.6287
AUCJ (↑)	**0.7044**	0.5581	0.5094	0.5317	0.6875	0.6698	0.6400

Fig. 5 Efficiency analysis. Comparison of average attribution map generation time (x-axis, seconds) against model performance (y-axis), both when considering AVG % DROP (left) and INSERT metrics (right), for the ViT-B/16 model

Similarity (NSS) and Area Under the ROC Curve Judd version (AUCJ) metrics (Zanca et al., 2018). The higher the value of such metrics, the better.

Compared Models & Architecture Details. We compare FovEx against the SOTA competitors described in the previous experiments. For this study, we focus on the ViT-B/16 predictor, pretrained on ImageNet-21K (Ridnik et al., 2021) and fine-tuned on ImageNet-1K (Russakovsky et al., 2015).

Results. As reported in Table 3, FovEx outperforms other methods in both the considered metrics. The achieved performances are almost doubled with respect to gradCAM and gradCAM++. Remarkably, explanation methods that have been proposed for the tested architecture (GAE, Cls. Emb.) are outperformed. This result highlights the alignment between the explanation maps produced by FovEx and human gaze patterns during free-viewing of natural images and enhances our confidence in FovEx's ability to close the interpretation gap between humans and machines. Indeed, we remark that these results should be analyzed in conjunction with the ones about XAI metrics, summarised in Tables 2 and 1. Overall, to proposed FovEx is capable of coupling (i) the best alignment with human gaze patterns among the tested explanation methods, which we showed by investigating its correlation with human attention, together with (ii) extremely promising quantitative performances on several XAI metrics. This suggests that FovEx provides a more accurate and intuitive understanding of the model's decision-making process with respect to alternative methods, by allowing for a better comparison between human fixations and input localization importance of the model.

4.3 Efficiency Analysis

In this section, we investigate the time efficiency of the proposed FovEx compared with other SOTA approaches. We report the time necessary to produce an attribution map aver-

aged over the considered 5000 images from the ImageNet-1K validation dataset. We focus our analysis on the ViT-B/16 predictor (a similar conclusion holds for the ResNet-50 architecture). In Figure 5-left we report the inference time (x-axis, seconds) against the AVG % DROP metric (y-axis) for all the considered methods. The closer to the left-bottom corner, the better. Figure 5-right depicts inference time (x-axis, seconds) against INSERT (y-axis). In this case, the left-top corner implies better performance. Overall, FovEx outperforms perturbation-based approaches (Mean. Pert. and RISE) in both performance and time efficiency. Conversely, gradient-based methods are more time efficient but are far from the performances achieved by FovEx in the considered metrics.

4.4 Ablation studies

We undertook several ablation studies, which are delineated below. We conducted experiments to study the effect of various parameters such as random restart (RR), foveation sigma (σ_f), forgetting factor (β), scanpath length (N), optimization steps (OS), blur sigma (σ_b), and blur filter size (BFS). Figure 6 depicts the summarized impact of the parameters mentioned above on the performance of FovEx. We utilized the ViT-B/16 model and the corresponding subset of the ImageNet-1K validation dataset employed in Section 4.1 for the ablation studies. Default values for the parameters are given in Appendix B.

Random Restart (RR). The RR parameter helps to avoid local minima in the optimization scheme for finding the location of fixation f_t. When RR is set to True, the exploration starts at a completely random location. Figure 7 depicts the visual comparison of explanation maps generated using two different values of RR. On close inspection, it is evident that fixation locations are closer to each other when RR is set to False. Table 4 presents the comparison of quantitative evaluation for different configurations of RR. With RR set to True, we

Fig. 6 Summary of ablation study conducted to showcase the effect of various FovEx parameters. The ↑ near parameters (*x* axis) denotes an increase in their value; except for rr, it signifies the value shifts from False to True
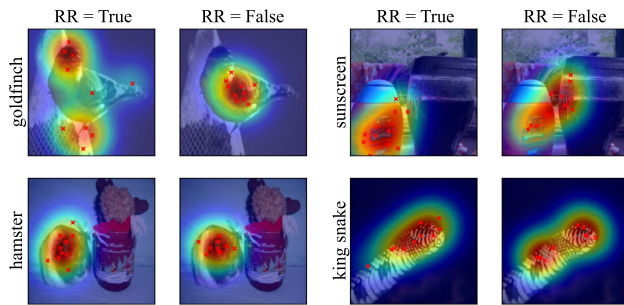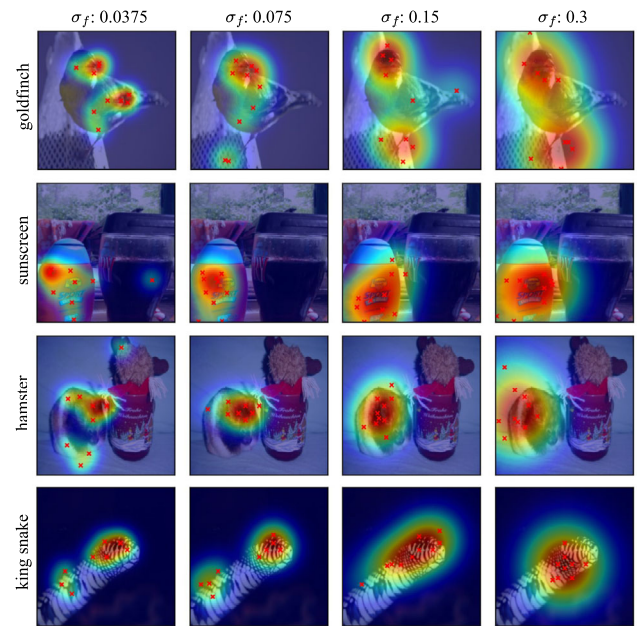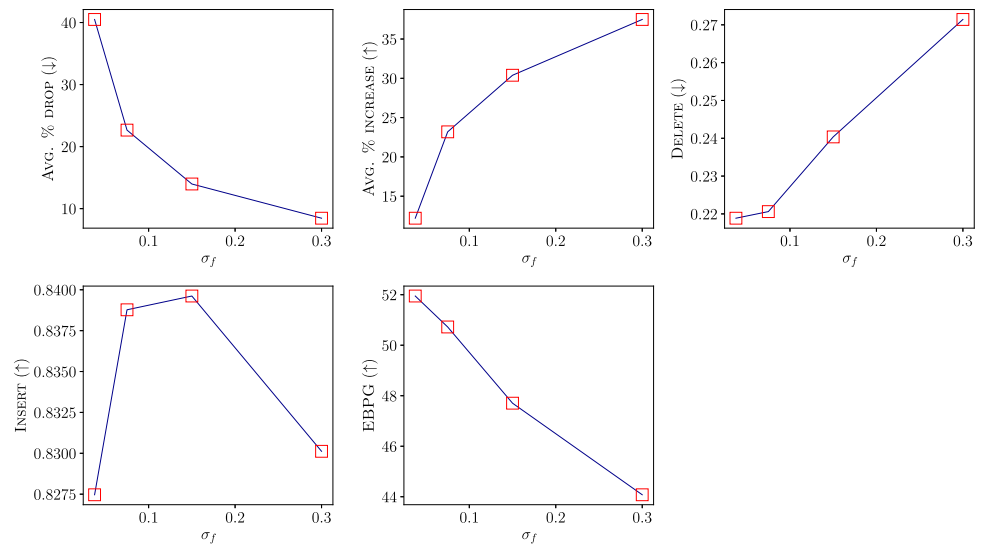

Fig. 7 Qualitative assessment. Comparison of attribution maps generated using FovEx with different RR values

get better performance in every metric except for the EBPG metric.

Foveation Sigma (σ_f). The σ_f parameter controls the standard deviation of the area with higher visual acuity in the foveated input image. A higher value of σ_f results in a larger region with enhanced visual clarity. In our experiments, we create attribution maps with varying σ_f values- specifically, {0.3, 0.15, 0.075, 0.0375}.

The visual comparison shown in Figure 8 indicates an increase in σ_f values results in attribution maps with larger areas, while lower σ_f values result in a more concentrated explanation map. We report quantitative assessments in Figure 9. As the value of σ_f increases, there is an improvement in the metrics of AVG. % DROP and AVG. % INCREASE. However, the metric of DELETE decreases. It is worth noting that an optimal σ_f value of 0.15 shows the maximum performance in the INSERT metric. Increasing σ_f leads to a wider spread of the attribution map, compromising localization.

Fig. 8 Qualitative assessment. Comparison of attribution maps generated using FovEx with different σ_f values

Forgetting Factor (β). The forgetting factor $0 \leq \beta \geq 1$ regulates how much information is retained from previous fixations. In our study, we created attribution maps for various β values to investigate their effects. Specifically, we used the values {0.0, 0.1, 0.25, 0.5, 0.75, 1.0}. Visual comparison illustrated in Figure 10 reveals that the quality of explanation maps decreases with an increase in the β value. However, the

Table 4 Quantitative assessment. Average metrics on the considered subset of ImageNet-1K validation dataset for ViT-B/16 model for different RR values. Bold terms denote the best performance and underlined terms represent the second best performance

| Eval. Name	AVG. % DROP (↓)	AVG. % INCREASE (↑)	DELETE (↓)	INSERT (↑)	EBPG (↑)
RR = True	**13.973**	**30.389**	**0.240**	**0.840**	47.705
RR = False	17.561	26.389	0.247	0.824	**51.067**

Fig. 9 Quantitative assessment. Illustration of variation in performance of FovEx with respect to σ_f
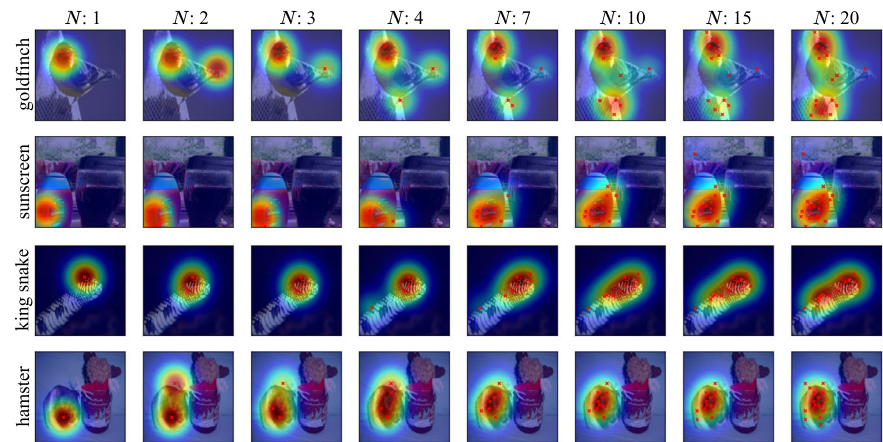

Fig. 10 Qualitative assessment. Comparison of attribution maps generated using FovEx with different β values
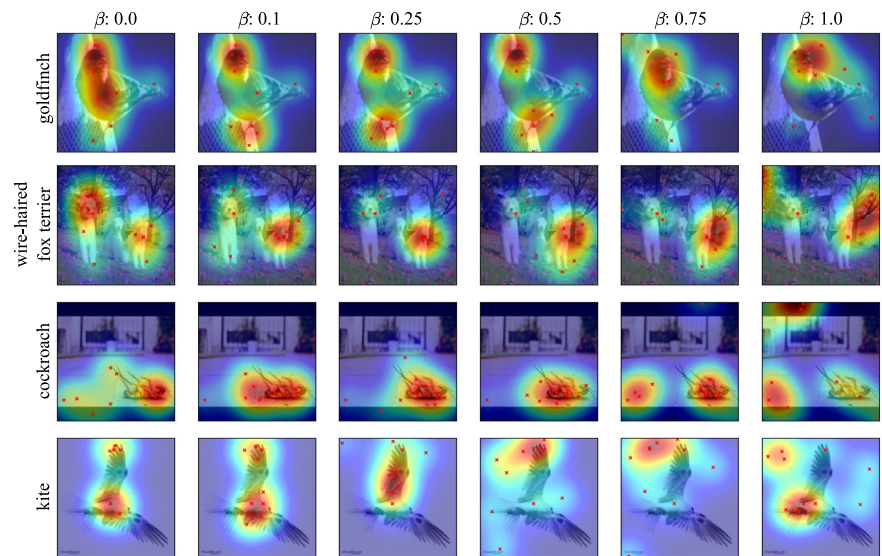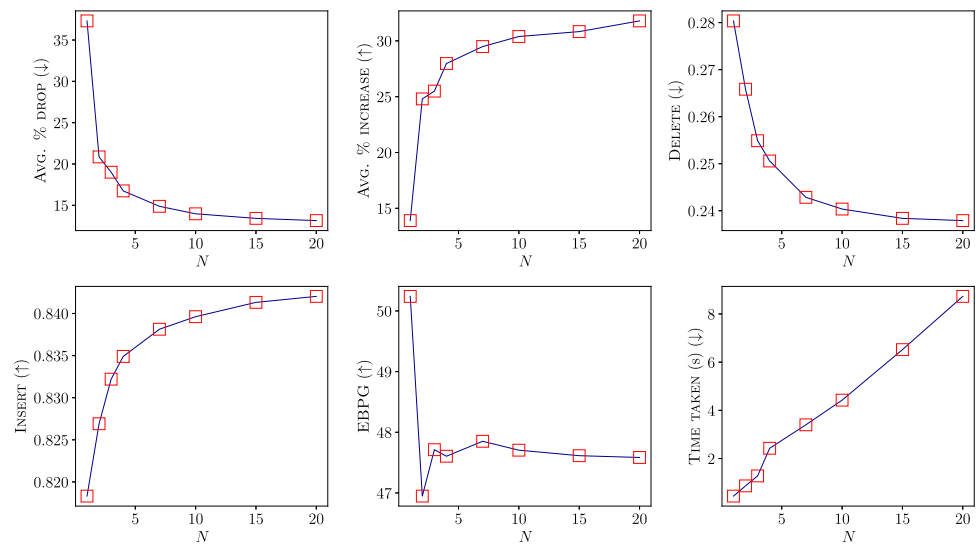

Fig. 11 Quantitative assessment. Illustration of variation in performance of FovEx with respect to β
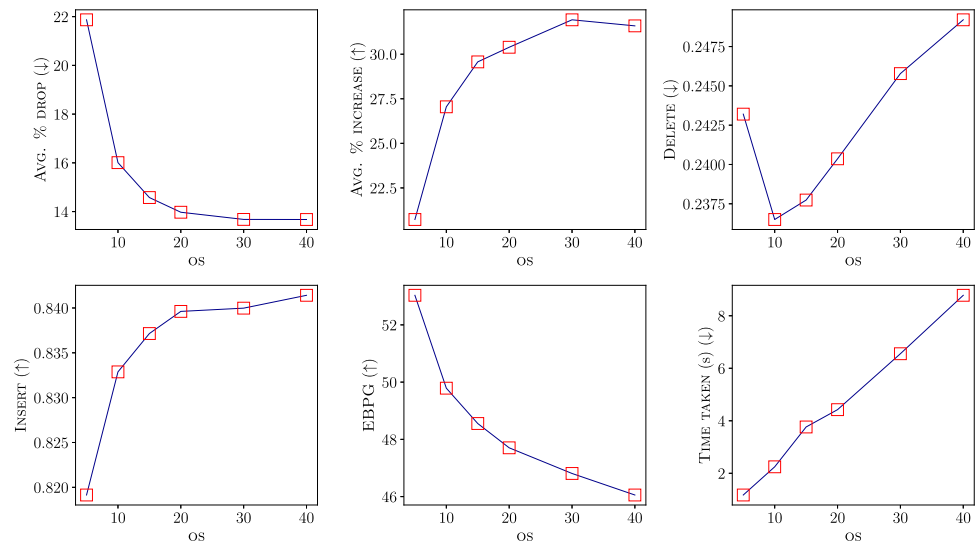
Fig. 12 Qualitative assessment. Visual comparison of attribution maps generated using FovEx with different N values

Fig. 13 Quantitative assessment. Illustration of variation in performance of FovEx with respect to N

Fig. 14 Qualitative assessment. Visual comparison of attribution maps generated using FovEx with different OS values

Fig. 15 Quantitative assessment. Illustration of variation in performance of FovEx with respect to OS

Fig. 16 Qualitative assessment. Visual comparison of attribution maps generated using FovEx with different σ_b values

quantitative evaluation paints a nuanced picture as depicted in Figure 11.

Scanpath Length (N). N controls the length of the scanpath F, i.e., the number of fixation points computed to generate an attribution map E. To study the impact of N on the performance of FovEx, we generate attribution maps considering the following values for N, $\{1, 2, 3, 4, 7, 10, 15, 20\}$. As depicted in Figure 12 and Figure 13, both qualitative and quantitative (except EBPG and TIME TAKEN) performance improve with an increase in the value of N. However, the average time taken to produce an attribution map increases monotonically with N.

Optimization Steps (OS). The OS parameter controls the number of iterations performed in the optimization scheme to find a fixation location. We generate attribution maps using the following values of OS, $\{5, 10, 15, 20, 30, 40\}$. The performance of FovEx improves with increment in OS values both qualitatively and quantitatively (except for DELETE and EBPG metric) as illustrated in Figure 14 and Figure 15. Similar to the effect of N with an increase in OS, the average time taken to produce an attribution map increases.

Blur Sigma (σ_b). The σ_b parameter represents the standard deviation of the filter values used to create a blurred version of the input image which constitutes the peripheral region of the final foveated input image. We consider the following values for σ_b, $\{5, 10, 15, 20\}$ for the experiments. Visual comparison in Figure 16 depicts plausible attribution maps generated at all σ_b values. However, the quantitative assessment illustrated in Figure 17 showcases an improvement in performance with increment in the value of σ_b except in the INSERT metric.

Blur Filter Size (BFS). The BFS parameter represents the size of the Gaussian kernel used to create the blurred version of the input image. To study the effect BFS we consider the following values, $\{11, 21, 31, 41, 51\}$. Qualitative assessment illustrated in Figure 18, showcases visually slightly better attribution maps at BFS value of 41 than other values. Figure 19 illustrates the quantitative assessment. With an increase in BFS value the performance of FovEx increases except in INSERT metric.

5 Conclusion

We introduced FovEx, a novel explanation technique that integrates foveated human vision principles into the explanation generation process for black-box models. FovEx not only generates visually coherent explanation maps for both trans-

Fig. 17 Quantitative assessment. Illustration of variation in performance of FovEx with respect to σ_b
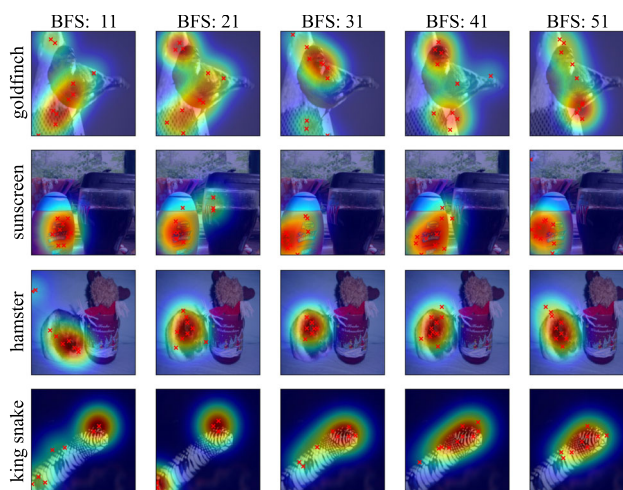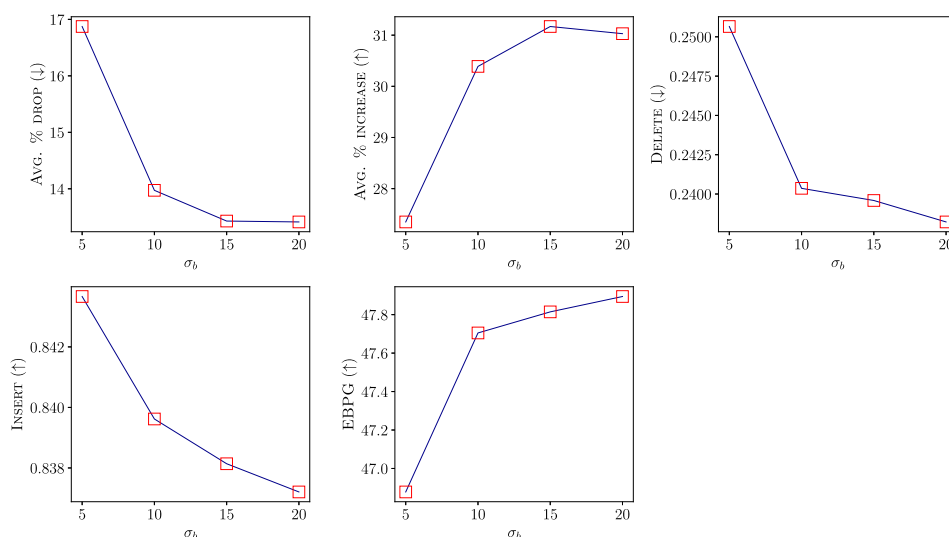
Fig. 18 Qualitative assessment. Visual comparison of attribution maps generated using FovEx with different BFS values

former and convolution-based models but also possesses the ability to discern between different classes. It demonstrates superior performance in quantitative evaluations, outperforming competing methods for 4 out of 5 evaluation metrics for the ViT-B/16 model and for 3 out of 5 metrics for ResNet-50. Furthermore, incorporating biological foveated vision concepts allows FovEx to generate explanation maps that have a higher correlation to human gaze patterns on the considered dataset, as illustrated in Section 4.2, paving the way for a more accurate and intuitive understanding of model's decisions.

FovEx surpasses gradient-based and perturbation-based methods in all faithfulness metrics except the DELETE metric, for which it falls short in comparison to gradient-based methods. This is mostly due to FovEx's optimization scheme prioritizing the selection of image regions that are the most informative for the explanation, thereby offering a specific advantage in classification performance preservation, as reflected by other metrics such as INSERT.

A key takeaway from our work is that human-inspired perturbations lead to explanations that not only align closely with human gaze data but also demonstrate a positive correlation between this alignment and the quality of explanations, as validated by standard XAI metrics. From an ethical perspective, we notice that the use of a small, potentially non-representative human gaze dataset might raise biases, and results on human alignment could not be representative of the diversity of a larger population. However, our method opens the door to better-comparing model and human gaze pattern behavior.

While our approach demonstrated its stability for different choices of the hyperparameters, see Section 4.4, to enhance usability and make our method more accessible, future work will focus on exploring automated optimization methods and implementing more intuitive default settings for parameter tuning.

FovEx demonstrates strong performance in preserving classification outcomes and generating interpretable explanation maps, particularly excelling in the INSERT metric due to its preservation-focused optimization strategy (see Section 4, quantitative inspection). However, this design introduces a trade-off, as the method is less effective on the DELETE metric, which rewards the identification of regions whose removal most degrades model performance. This limitation arises because FovEx prioritizes finding the most informative regions to retain rather than those that, when removed, significantly impact the model's prediction. To address this, future work could explore multi-objective optimization frameworks that explicitly balance both preservation and deletion criteria, potentially improving overall

Fig. 19 Quantitative
assessment. Illustration of
variation in performance of
FovEx with respect to BFS
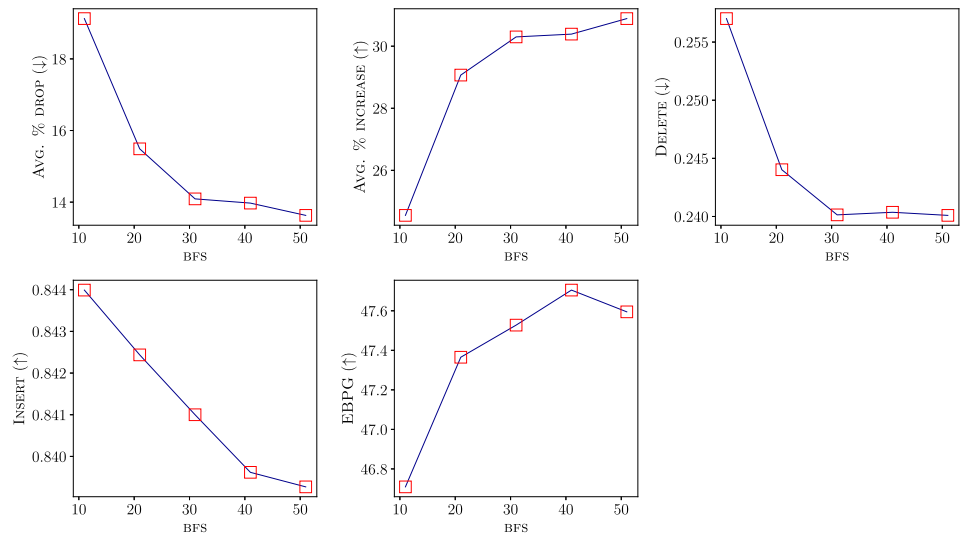
explanation fidelity across complementary evaluation metrics.

Future work will extend the applicability of FovEx to models in other task domains, such as image captioning.

Appendix A Performance on Various Models

We compare the performance of FovExfor three additional predictors ConvNeXt, ViT-B/16, and ViT-B/32 from torchvision. The ViT-B/16 model considered here is pre-trained on ImageNet-1K whereas the ViT-B/16 model in the main experiments is pre-trained on the ImageNet-21K dataset and fine-tuned on the ImageNet-1K dataset. We consider gradCAM (Selvaraju et al., 2017), gradCAM++ (Chattopadhay et al., 2018), and Mean. Pert. (Fong & Vedaldi, 2017) for comparison in the case of the ConvNeXt Base model. For ViT-B/16 and ViT-B/32 models we consider gradCAM and gradCAM++ for comparison. We employ randomCAM as a baseline method. We do not consider RISE (Petsiuk et al., 2018) for ConvNeXt, ViT-B/16, and ViT-B/32 due to the high time requirements to generate the attribution maps. Similarly, due to high computational complexity, the Mean. Pert. method is not considered for ViT-B/16 and ViT-B/32. Transformer-specific methods like GAE (Chefer et al., 2021) and Cls. Emb. (Vilas et al., 2023) methods are not utilized due to the need to change the pre-existing implementation from torchvision.

We report qualitative assessments in Figure 20, Figure 21, and Figure 22 for ConvNeXt Base, ViT-B/16 and ViT-B/32 models respectively. For all three cases, FovEx creates noise-free and consistent attribution maps. Table 5, Table 6, and Table 7 depict the qualitative evaluation results for the ConvNeXt Base model, ViT-B/16 and ViT-B/32 model

Fig. 20 Qualitative assessment.
Visual comparison of attribution
maps generated using FovEx,
gradCAM, gradCAM++, and
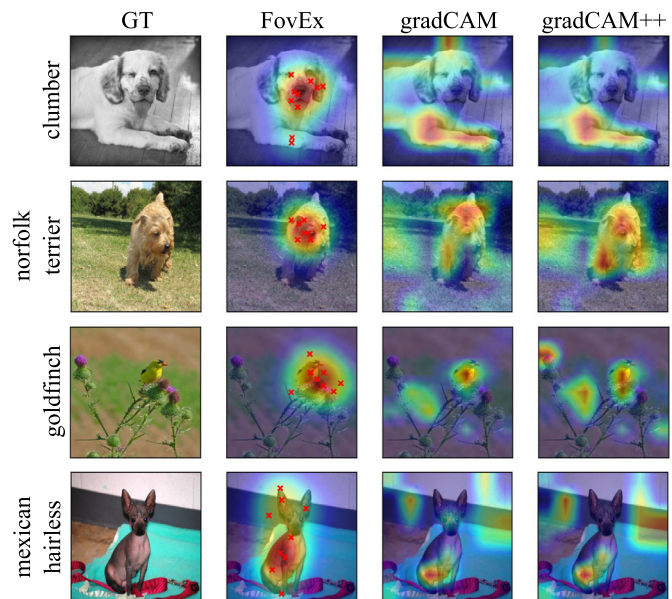Mean. Pert. for ConvNeXt Base
model
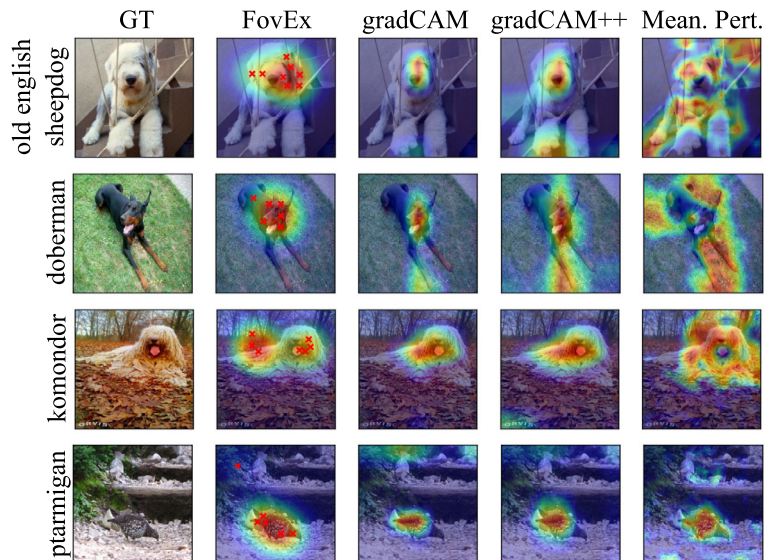
Fig. 21 Qualitative assessment. Visual comparison of attribution maps generated using FovEx, gradCAM, gradCAM++, and randomCAM for ViT-B/16 model
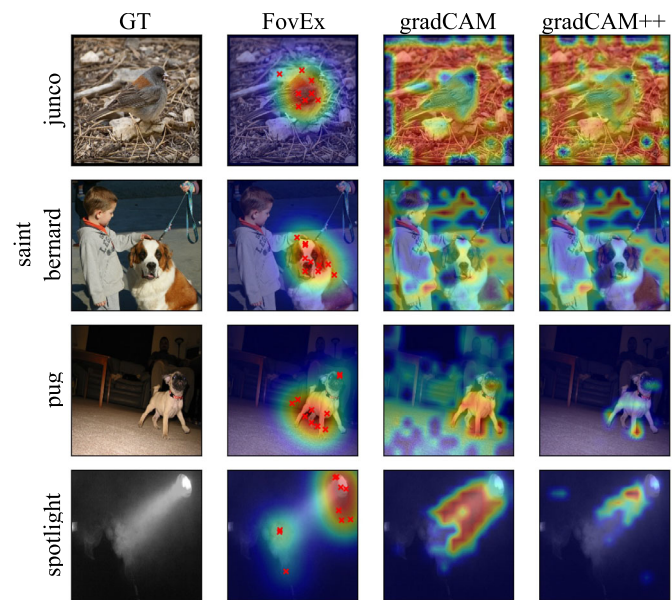
Fig. 22 Qualitative assessment. Visual comparison of attribution maps generated using FovEx, gradCAM, gradCAM++, and randomCAM for ViT-B/32 model

Table 5 Quantitative assessment. Average metrics on the considered subset of ImageNet-1K validation dataset for ConvNeXt Base model. Bold terms denote the best performance and underlined terms represent the second best performance

Eval. Name	FovEx	grad CAM	grad CAM++	Mean. Pert.	random CAM
AVG. % DROP (↓)	**59.6592**	80.9472	<u>78.2356</u>	87.0630	86.6874
AVG. % INCREASE (↑)	**9.7699**	3.4699	<u>3.6999</u>	1.6999	0.7499
DELETE (↓)	0.2142	0.2098	<u>0.1907</u>	**0.0739**	0.2638
INSERT (↑)	**0.1833**	0.1548	<u>0.1606</u>	0.1422	0.1478
EBPG (↑)	48.1485	**53.1640**	<u>49.4505</u>	43.3129	30.9518

Table 6 Quantitative assessment. Average metrics on the considered subset of ImageNet-1K validation dataset for ViT-B/16 model. Bold terms denote the best performance and underlined terms represent the second best performance

Eval. Name	FovEx	grad CAM	grad CAM++	random CAM
Avg. % drop (↓)	**48.2258**	66.9577	75.6945	91.4881
Avg. % increase (↑)	**17.1490**	4.9299	3.5299	0.6490
Delete (↓)	**0.0886**	0.0913	0.1101	0.1283
Insert (↑)	**0.3647**	0.2956	0.2964	0.2777
EBPG (↑)	**48.1485**	39.6095	38.0989	36.0899

Table 7 Quantitative assessment. Average metrics on the considered subset of ImageNet-1K validation dataset for ViT-B/32 model. Bold terms denote the best performance and underlined terms represent the second best performance

Eval. Name	FovEx	grad CAM	grad CAM++	random CAM
Avg. % drop (↓)	**44.0460**	64.1220	70.1500	85.8250
Avg. % increase (↑)	**32.6890**	16.8690	13.7690	5.3499
Delete (↓)	**0.1034**	0.1261	0.1241	0.1565
Insert (↑)	**0.4992**	0.4293	0.4303	0.3773
EBPG (↑)	**48.4390**	40.1740	40.9060	37.6540

respectively. FovEx maintains state-of-the-art performance across all the models by outperforming all other XAI methods in 3 out of 5 metrics in the case of ConvNeXt and in 5 out of 5 metrics in the case of ViT-B/16 and ViT-B/32 models.

Appendix B Default Parameter Values.

In table 8 we report the default parameters of our experiments.

Table 8 Parameter Values. These values are used in Section 4.1 and Appendix A

| Parameter Name	Value
RR	True
σ_f	0.15
β	0.1
N	10
OS	20
σ_b	10
BFS	41
λ	0.1

Appendix C FovEx in Other Application Domain

Object Detection. To demonstrate the scalability of FovEx to different scenarios, we generate explanation maps for a SSD (Liu et al., 2016) predictor from torchvision on the COCO (Lin et al., 2014) dataset, consisiting in an object detection task. We perform a small-scale qualitative and quantitative evaluation using 100 images sampled randomly from the COCO 2017 validation set. We considered D-RISE

(Petsiuk et al., 2021) as the baseline for comparison. Figure 23 shows the visual comparison of explanation maps. It can be observed that FovEx generates visually plausible, less-noisy explanations. For quantitative evaluation, we compute the pointing game (PG), energy-based pointing game (EBPG), and intersection over union (IoU). To compute IoU, we utilize the bounding boxes generated by the model for the original image and the perturbed version of the original image modified by the explanation map. The intuition behind the computation of IoU is that a good explanation map should preserve the bounding box generation capabilities of the object detection model. Table 9 presents the results from the qualitative evaluation. FovEx performs better than D-RISE in EBPG and IoU metrics; however, it lags behind D-RISE in the PG metric.

Vision-Language Model. To showcase the versatility of FovEx, we generate explanation maps for the CLIP (Radford et al., 2021) vision model with the help of captions available in the COCO dataset and class labels from the ImageNet-1k dataset. The overall process simulates a caption generation task indirectly. We provide a randomly chosen caption from the available captions to the CLIP text encoder to produce a text embedding. Then, generate an explanation map for the vision model using the text embedding as the label to guide FovEx. We follow the parameter setup suggested in Schwinn et al. (2022), Zanca et al. (2023). Figure 24 presents a visualization for the explanation maps and the captions used to generate such explanations. From inspection, it can be observed that FovEx produces visually plausible text-specific explanation maps.
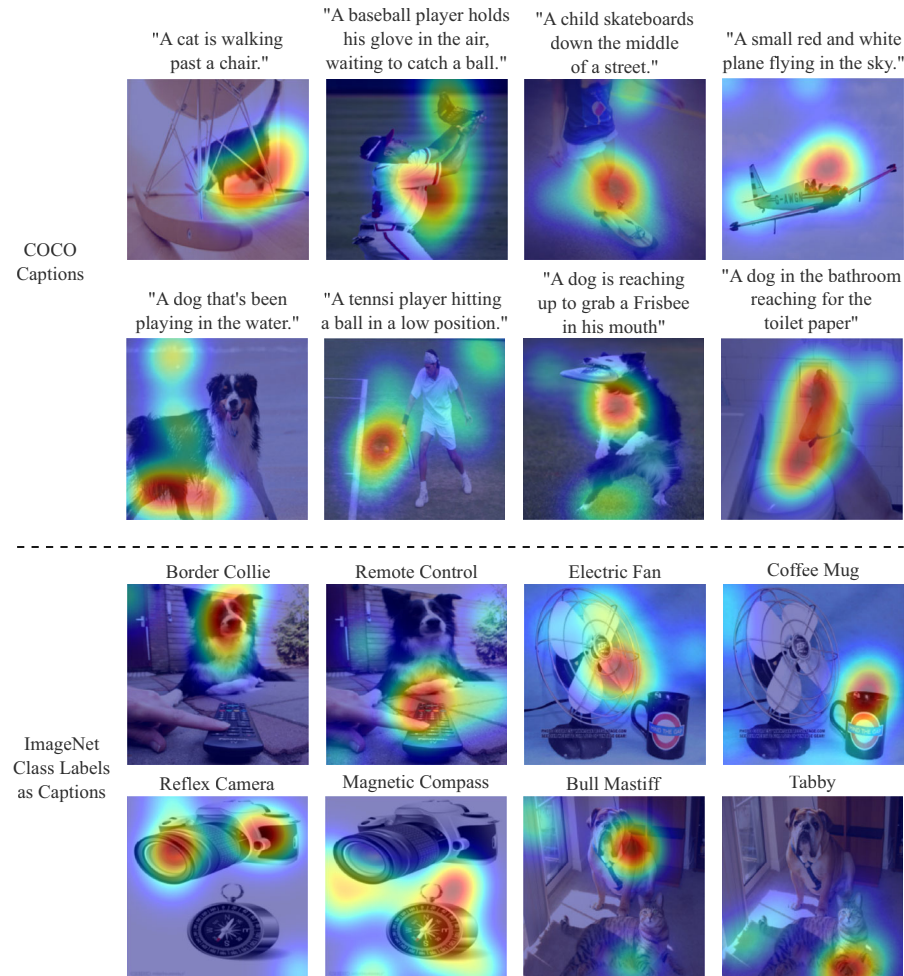
Fig. 23 Qualitative assessment. Visual comparison of attribution maps generated using FovEx and D-RISE for the SSD model
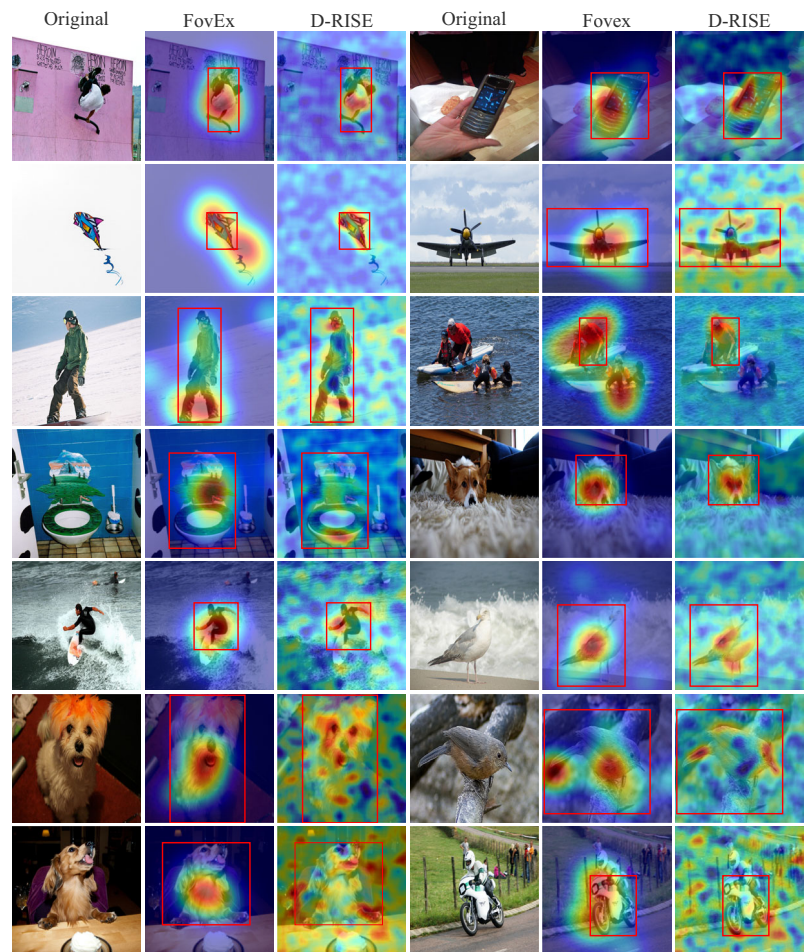
Table 9 Quantitative assessment for Object detection. Higher value represents a better score

Method Name	PG	EBPG	IoU
FovEx	0.7	**0.32**	**0.833**
D-RISE	**0.9**	0.22	0.094

Fig. 24 Qualitative assessment. Visual comparison of attribution maps generated using FovEx for CLIP model

Appendix D Additional Visualizations of Explanation Maps

Figures 25, 26 and 27.
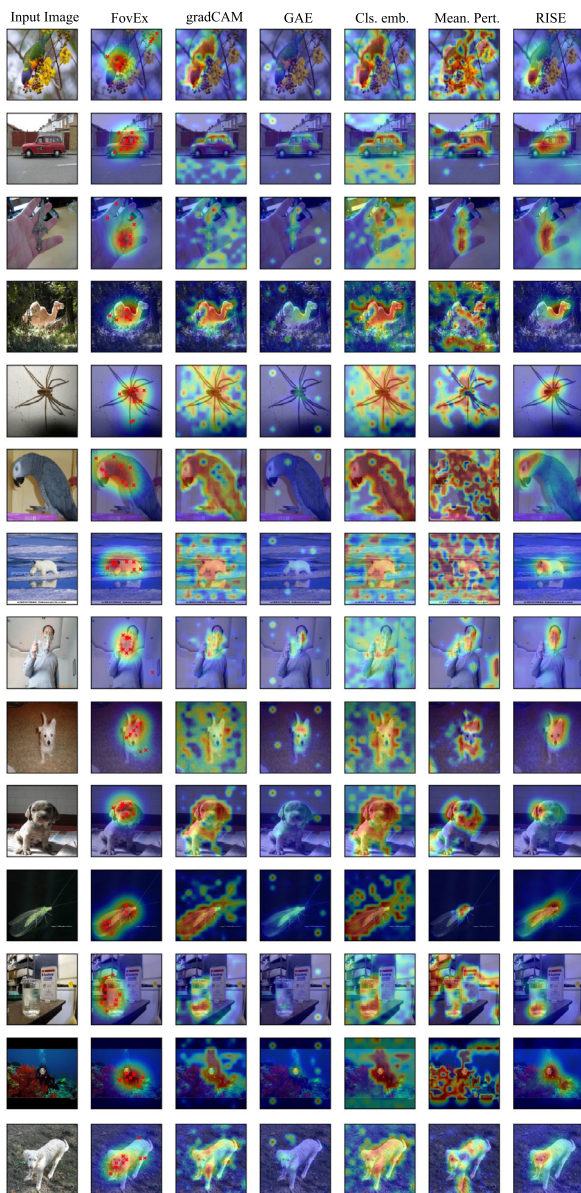
Fig. 25 Qualitative assessment. Visual comparison of attribution maps generated using FovEx, gradCAM, gradCAM++, GAE, Cls. Emb., Mean. Pert., and RISE for ViT-B/16 model used in Section 4.1
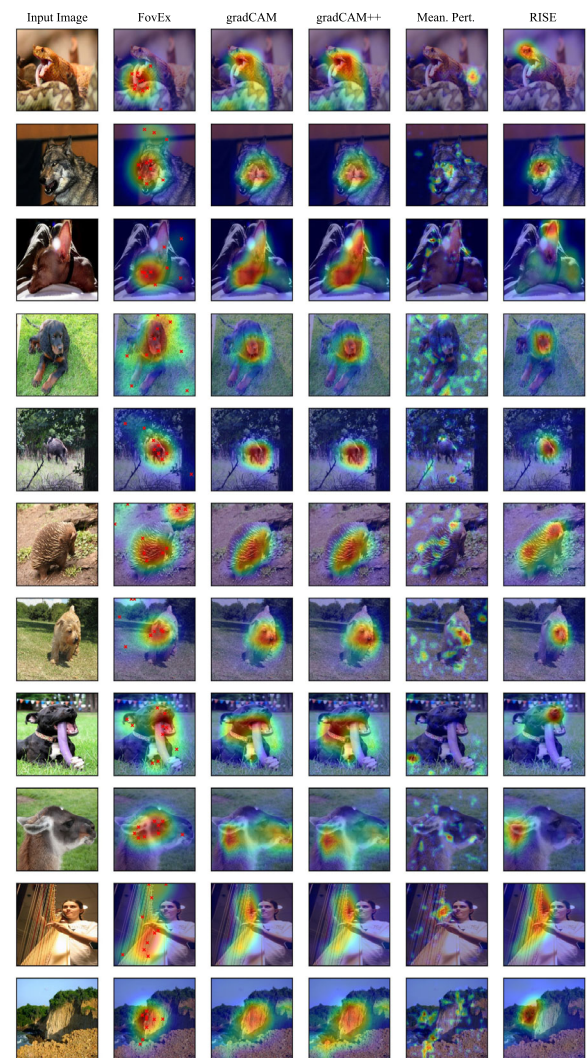
Fig. 26 Qualitative assessment. Visual comparison of attribution maps generated using FovEx, gradCAM, gradCAM++, Mean. Pert., and RISE for ResNet-50 model used in Section 4.1

Fig. 27 Qualitative assessment. Visual comparison of attribution maps generated using FovEx, gradCAM, gradCAM++, and Mean. Pert. for ConvNeXt model used in Appendix A

References

Abnar, S., & Zuidema, W. (2020). Quantifying attention flow in transformers. arXiv preprint arXiv:2005.00928.

Achtibat, R., Dreyer, M., Eisenbraun, I., Bosse, S., Wiegand, T., Samek, W., & Lapuschkin, S. (2023). From attribution maps to human-understandable explanations through concept relevance propagation. *Nature Machine Intelligence, 5*(9), 1006–1019.

Ali, S., Abuhmed, T., El-Sappagh, S., Muhammad, K., Alonso-Moral, J. M., Confalonieri, R., Guidotti, R., Del Ser, J., Díaz-Rodríguez, N., & Herrera, F. (2023). Explainable artificial intelligence (xai): What we know and what is left to attain trustworthy artificial intelligence. *Information Fusion, 99*, Article 101805.

Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.-R., & Samek, W. (2015). On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS One, 10*(7), 0130140.

Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., & Brunskill, E., et al. (2021). On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258.

Bringmann, A., & Wiedemann, P. (2021). *The Fovea: Structure, function, development, and tractional disorders*. Academic Press.

Chattopadhay, A., Sarkar, A., Howlader, P., & Balasubramanian, V. N. (2018). Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)* (pp. 839–847).

Chattopadhay, A., Sarkar, A., Howlader, P., & Balasubramanian, V. N. (2018). Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In: *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 839–847 . IEEE.

Chefer, H., Gur, S., & Wolf, L. (2021). Generic attention-model explainability for interpreting bi-modal and encoder-decoder transformers. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, (pp. 397–406).

Chefer, H., Gur, S., & Wolf, L. (2021a). Transformer interpretability beyond attention visualization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, (pp. 782–791).

Chefer, H., Gur, S., & Wolf, L. (2021b). Generic attention-model explainability for interpreting bi-modal and encoder-decoder transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 397–406).

Choi, H., Jin, S., & Han, K. (2023). Adversarial normalization: I can visualize everything (ice). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 12115–12124).

Darcet, T., Oquab, M., Mairal, J., & Bojanowski, P. (2023). Vision transformers need registers. arXiv preprint arXiv:2309.16588.

Deza, A., & Konkle, T. (2020). Emergent properties of foveated perceptual systems. arXiv preprint arXiv:2006.07991.

Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., & Gelly, S., et al. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929.

Fong, R. C., & Vedaldi, A. (2017). Interpretable explanations of black boxes by meaningful perturbation. In *Proceedings of the IEEE International Conference on Computer Vision* (pp. 3429–3437).

Fong, R., Patrick, M., & Vedaldi, A. (2019). Understanding deep networks via extremal perturbations and smooth masks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 2950–2958).

Fujiyoshi, H., Hirakawa, T., & Yamashita, T. (2019). Deep learning-based image recognition for autonomous driving. *IATSS Research, 43*(4), 244–252.

Goodman, B., & Flaxman, S. (2017). European union regulations on algorithmic decision-making and a "right to explanation". *AI magazine, 38*(3), 50–57.

Gruosso, M., Capece, N., & Erra, U. (2021). Human segmentation in surveillance video with deep learning. *Multimedia Tools and Applications, 80*, 1175–1199.

Han, Y., Roig, G., Geiger, G., & Poggio, T. (2020). Scale and translation-invariance for novel objects in human vision. *Scientific reports, 10*(1), 1411.

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, (pp. 770–778).

Hendrickson, A. E., & Yuodelis, C. (1984). The morphological development of the human fovea. *Ophthalmology, 91*(6), 603–612.

Hsiao, J., & Chan, A. (2023). Towards the next generation explainable ai that promotes ai-human mutual understanding. In: *XAI in Action: Past, Present, and Future Applications*.

Iqbal, R., Maniak, T., Doctor, F., & Karyotis, C. (2019). Fault detection and isolation in industrial processes using deep learning approaches. *IEEE Transactions on Industrial Informatics, 15*(5), 3077–3084.

Ivanovs, M., Kadikis, R., & Ozols, K. (2021). Perturbation-based methods for explaining deep neural networks: A survey. *Pattern Recognition Letters, 150*, 228–234.

Iwana, B. K., Kuroki, R., & Uchida, S. (2019). Explaining convolutional neural networks using softmax gradient layer-wise relevance propagation. In *International Conference on Computer Vision Workshops*.

Jing, T., Xia, H., Tian, R., Ding, H., Luo, X., Domeyer, J., Sherony, R., & Ding, Z. (2022). Inaction: Interpretable action decision making for autonomous driving. In *European Conference on Computer Vision* (pp. 370–387). Springer.

Jonnalagadda, A., Wang, W. Y., Manjunath, B., & Eckstein, M. P. (2021). Foveater: Foveated transformer for image classification. arXiv preprint arXiv:2105.14173.

Judd, T., Ehinger, K., Durand, F., & Torralba, A. (2009). Learning to predict where humans look. In *2009 IEEE 12th International Conference on Computer Vision* (pp. 2106–2113). IEEE.

Lai, Q., Khan, S., Nie, Y., Sun, H., Shen, J., & Shao, L. (2020). Understanding more about human and machine attention in deep neural networks. *IEEE Transactions on Multimedia, 23*, 2086–2099.

Lee, J. R., Kim, S., Park, I., Eo, T., & Hwang, D. (2021). Relevance-cam: Your model already knows where to look. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 14944–14953).

Li, Z., Wang, W., Li, Z., Huang, Y., & Sato, Y. (2021). Towards visually explaining video understanding networks with perturbation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* (pp. 1120–1129).

Liang, J., Mahler, J., Laskey, M., Li, P., & Goldberg, K. (2017). Using dvrk teleoperation to facilitate deep learning of automation tasks for an industrial robot. In *2017 13th IEEE Conference on Automation Science and Engineering (CASE)*, (pp. 1–8). IEEE.

Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., & Zitnick, C. L. (2014). Microsoft coco: Common objects in context. In *Computer vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13* (pp. 740–755). Springer.

Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., & Berg, A. C. (2016). Ssd: Single shot multibox detector. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14* (pp. 21–37). Springer.

Liu, Z., Mao, H., Wu, C.-Y., Feichtenhofer, C., Darrell, T., & Xie, S. (2022). A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 11976–11986).

Longo, L., Brcic, M., Cabitza, F., Choi, J., Confalonieri, R., Del Ser, J., Guidotti, R., Hayashi, Y., Herrera, F., Holzinger, A., et al. (2024). Explainable artificial intelligence (xai) 2.0: A manifesto of open challenges and interdisciplinary research directions. *Information Fusion, 106*, Article 102301.

Lucieri, A., Bajwa, M. N., Dengel, A., & Ahmed, S. (2020). Explaining ai-based decision support systems using concept localization maps. In *International Conference on Neural Information Processing* (pp. 185–193). Springer.

Malkin, E., Deza, A., & Poggio, T. (2020). Cuda-optimized real-time rendering of a foveated visual system. arXiv preprint arXiv:2012.08655.

Maschler, B., & Weyrich, M. (2021). Deep transfer learning for industrial automation: A review and discussion of new techniques for data-driven machine learning. *IEEE Industrial Electronics Magazine, 15*(2), 65–75.

Morrison, K., Shin, D., Holstein, K., & Perer, A. (2023). Evaluating the impact of human explanation strategies on human-ai visual decision-making. *Proceedings of the ACM on Human-Computer Interaction, 7*(CSCW1), 1–37.

Najibi, M., Ji, J., Zhou, Y., Qi, C. R., Yan, X., Ettinger, S., & Anguelov, D. (2022). Motion inspired unsupervised perception and prediction in autonomous driving. In *European Conference on Computer Vision* (pp. 424–443). Springer.

Nguyen, T. T. H., Truong, V. B., Nguyen, V. T. K., Cao, Q. H., & Nguyen, Q. K. (2023). Towards trust of explainable ai in thyroid nodule diagnosis. arXiv preprint arXiv:2303.04731.

Nielsen, I. E., Dera, D., Rasool, G., Ramachandran, R. P., & Bouaynaya, N. C. (2022). Robust explainability: A tutorial on gradient-based attribution methods for deep neural networks. *IEEE Signal Processing Magazine, 39*(4), 73–84.

Ouyang, C., Biffi, C., Chen, C., Kart, T., Qiu, H., & Rueckert, D. (2020). Self-supervision with superpixels: Training few-shot medical image segmentation without annotation. In: *Computer Vision–ECCV 2020: 16th European Conference, August 23–28, 2020, Proceedings, Part XXIX 16* (pp. 762–780). Springer, Glasgow, UK.

Pamplona, D., & Bernardino, A. (2009). Smooth foveal vision with gaussian receptive fields. In *2009 9th IEEE-RAS International Conference on Humanoid Robots* (pp. 223–229). IEEE.

Petsiuk, V., Das, A., & Saenko, K. (2018). Rise: Randomized input sampling for explanation of black-box models. In *Proceedings of the British Machine Vision Conference (BMVC)*.

Petsiuk, V., Jain, R., Manjunatha, V., Morariu, V. I., Mehra, A., Ordonez, V., & Saenko, K. (2021). Black-box explanation of object detectors via saliency maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 11443–11452).

Qi, R., Zheng, Y., Yang, Y., Cao, C. C., & Hsiao, J. H. (2023). Explanation strategies for image classification in humans vs. current explainable ai. arXiv preprint arXiv:2304.04448.

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., & Clark, J., et al. (2021). Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning* (pp. 8748–8763). PmLR.

Ridnik, T., Ben-Baruch, E., Noy, A., & Zelnik-Manor, L. (2021). Imagenet-21k pretraining for the masses. arXiv preprint arXiv:2104.10972.

Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., & Fei-Fei, L. (2015). ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV), 115*(3), 211–252.

Schwinn, L., Precup, D., Eskofier, B., & Zanca, D. (2022). Behind the machine's gaze: Neural networks with biologically-inspired constraints exhibit human-like visual attention. arXiv preprint arXiv:2204.09093.

Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization. In *2017 IEEE International Conference on Computer Vision (ICCV)* (pp. 618–626).

Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision*, (pp. 618–626).

Simonyan, K., Vedaldi, A., & Zisserman, A. (2014). Deep inside convolutional networks: Visualising image classification models and saliency maps. In *Workshop at International Conference on Learning Representations*.

Singh, V., Singh, S., & Gupta, P. (2020). Real-time anomaly recognition through cctv using neural networks. *Procedia Computer Science, 173*, 254–263.

Smilkov, D., Thorat, N., Kim, B., Viégas, F., & Wattenberg, M. (2017). Smoothgrad: removing noise by adding noise. arXiv preprint arXiv:1706.03825.

Tiezzi, M., Marullo, S., Betti, A., Meloni, E., Faggi, L., Gori, M., & Melacci, S. (2022). Foveated neural computation. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases* (pp. 19–35). Springer.

Tsotsos, J. K. (2021). *A Computational Perspective on Visual Attention*. MIT Press.

Vilas, M. G., Schaumlöffel, T., & Roig, G. (2023). Analyzing Vision Transformers for Image Classification in Class Embedding Space. *Advances in Neural Information Processing Systems, 36*, 40030–40041.

Volokitin, A., Roig, G., & Poggio, T. A. (2017). Do deep neural networks suffer from crowding? *Advances in neural information processing systems, 30*.

Wagner, J., Kohler, J. M., Gindele, T., Hetzel, L., Wiedemer, J. T., & Behnke, S. (2019). Interpretable and fine-grained visual explanations for convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 9097–9107).

Wang, J., Jin, Y., & Wang, L. (2022). Personalizing federated medical image segmentation via local calibration. In: *European Conference on Computer Vision* (pp. 456–472). Springer.

Wang, H., Wang, Z., Du, M., Yang, F., Zhang, Z., Ding, S., Mardziel, P., & Hu, X. (2020). Score-cam: Score-weighted visual explanations for convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops* (pp. 24–25).

Wright, R. D. (1998). *Visual Attention*. Oxford University Press.

Wu, J., Zhong, S.-H., Ma, Z., Heinen, S. J., & Jiang, J. (2018). Foveated convolutional neural networks for video summarization. *Multimedia Tools and Applications, 77*, 29245–29267.

Xu, J. (2021). A deep learning approach to building an intelligent video surveillance system. *Multimedia Tools and Applications, 80*(4), 5495–5515.

Zanca, D., Serchi, V., Piu, P., Rosini, F., & Rufa, A. (2018). Fixatons: A collection of human fixations datasets and metrics for scanpath similarity. arXiv preprint arXiv:1802.02534.

Zanca, D., Zugarini, A., Dietz, S., Altstidl, T. R., Ndjeuha, M. A. T., Schwinn, L., & Eskofier, B. M. (2023). Contrastive language-image pretrained models are zero-shot human scanpath predictors. CoRR **abs/2305.12380**.

Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., & Torralba, A. (2016). Learning deep features for discriminative localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, (pp. 2921–2929).