

Rethinking Saliency Maps: A Cognitive Human Aligned Taxonomy and Evaluation Framework for Explanations

Yehonatan Elisha¹, Seffi Cohen², Oren Barkan³, Noam Koenigstein¹

¹Tel Aviv University

²Harvard University

³The Open University

Abstract

Saliency maps have become a cornerstone of visual explanation in deep learning, yet there remains no consensus on their intended purpose and their alignment with specific user queries. This fundamental ambiguity undermines both the evaluation and practical utility of explanation methods. In this paper, we introduce the Reference-Frame×Granularity (RFxG) taxonomy—a principled framework that addresses this ambiguity by conceptualizing saliency explanations along two essential axes: the *reference-frame* axis (distinguishing between pointwise “Why Husky?” and contrastive “Why Husky and not Shih-tzu?” explanations) and the *granularity* axis (ranging from fine-grained class-level to coarse-grained group-level interpretations, e.g., “Why Husky?” vs. “Why Dog?”). Through this lens, we identify critical limitations in existing evaluation metrics, which predominantly focus on pointwise faithfulness while neglecting contrastive reasoning and semantic granularity. To address these gaps, we propose four novel faithfulness metrics that systematically assess explanation quality across both RFxG dimensions. Our comprehensive evaluation framework spans ten state-of-the-art methods, 4 model architectures, and 3 datasets. By suggesting a shift from model-centric to user-intent-driven evaluation, our work provides both the conceptual foundation and practical tools necessary for developing explanations that are not only faithful to model behavior but also meaningfully aligned with human understanding.

Code — <https://github.com/yonisGit/RFxG>

Datasets — <https://wordnet.princeton.edu>

1 Introduction

The increasing adoption of deep learning systems in high-stakes applications has amplified the need for transparency and interpretability. As deep models grow in complexity, so does the demand for methods that can explain their predictions in ways that are intelligible and actionable to human users. This growing field, known as Explainable AI (XAI), aims to bridge the gap between black-box models and human understanding. Across domains, numerous works have advanced the study of XAI (Sundararajan, Taly, and Yan 2017a; Smilkov et al. 2017; Chefer, Gur, and Wolf 2021b; Barkan et al. 2020, 2021b, 2024a,b,c, 2025). In computer

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

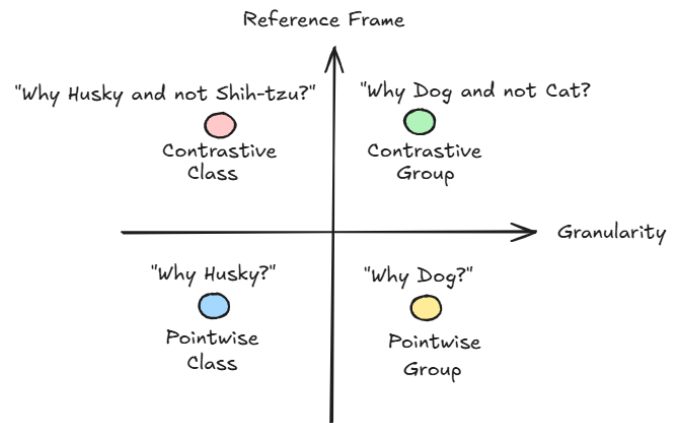


Figure 1: Our RFxG Explanation Axes. Note that there are also explanations between the points. For example: class-group contrastive questions like “Why Husky and not other Dogs?”

vision, XAI efforts have primarily focused on visual explanation techniques, with saliency maps being one of the most dominant paradigms (Selvaraju et al. 2017; Chefer, Gur, and Wolf 2021b; Barkan et al. 2021a, 2023c,a; Haddad et al. 2025). In parallel, contrastive explanation techniques (Dhurandhar et al. 2018; Xie et al. 2023; Wang et al. 2023) have emerged, aiming to answer comparative questions such as “Why class A rather than class B?”, rather than simply “Why class A?”. Despite their popularity, the interpretability and trustworthiness of saliency maps have been repeatedly questioned. Critiques in the literature have pointed out that many explanation methods may be misleading or unreliable, offering visual artifacts rather than faithful representations of model behavior (Adebayo et al. 2018; Kindermans et al. 2017; Longo et al. 2024). More broadly, Lipton (Lipton 2018) argued that the goals of interpretability are often underspecified, leading to a proliferation of methods that lack a clear understanding of what constitutes a good explanation. These concerns point to deeper, systemic challenges in how saliency explanations are conceived, evaluated, and aligned with user needs:

Challenge 1 – Lack of Support for User-Driven Explanatory Questions. Users rarely ask only “Why this

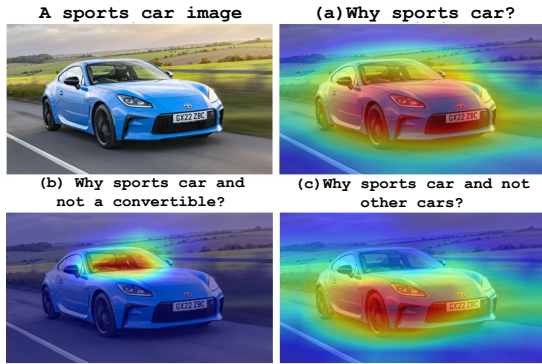


Figure 2: Saliency maps for a sport car using different types of questions that reflect our taxonomy: (a) Pointwise class. (b) Contrastive between two classes. (c) Contrastive between class and group.

class?” Instead, they often seek richer, comparative justifications—such as “Why a sports car and not a convertible?” or “What distinguishes this husky from other dog breeds?”—that most conventional saliency methods are not designed to support.

Challenge 2 – Interpretation Ambiguity from Missing Reference Frame and Granularity. Without explicit context about the explanation’s *reference frame* (e.g., pointwise vs. contrastive) or *semantic granularity* (e.g., fine-grained class vs. coarse-grained group), users may misinterpret saliency maps (Vilone and Longo 2021). A map highlighting fur texture might be meaningful for distinguishing huskies from shih-tzus but irrelevant when explaining why the image is a “dog” rather than a “cat.” This ambiguity undermines trust, interpretability, and ultimately the usability of saliency-based methods.

Challenge 3 – Inadequate Evaluation of Contrastive and Granularity-Varying Explanations. Existing evaluation metrics focus almost exclusively on pointwise faithfulness—assessing how salient pixels affect the score of a single target class. They fail to capture whether an explanation meaningfully distinguishes between classes or operates coherently at different levels of semantic abstraction (e.g., vehicle vs. sports car). As a result, current benchmarks cannot validate whether an explanation aligns with the user’s intended question.

These limitations highlight the importance of shifting attention toward the *user’s perspective*. What does a user actually want to know when they inspect a saliency map? Are they looking for reasons why an image was classified as a “husky,” or why it was classified as a “husky rather than a shih-tzu”? Do they care about fine-grained details or coarse-grained features? These kinds of questions suggest that saliency explanations should be considered along two important dimensions: the *reference-frame* axis (pointwise vs. contrastive explanations) and the *granularity* axis (fine-grained vs. coarse-grained classes). Figure 1 depicts the proposed axes space, while Figure 2 presents representative saliency maps that address distinct explanatory questions situated within this space. Our contributions are as fol-

lows: First, we introduce a **novel taxonomy** of saliency explanations based on *reference-frame* and *granularity*, offering a conceptual structure that reflects diverse user queries. Second, we identify critical shortcomings in existing evaluation metrics, which fail to distinguish among explanation types along these axes. Third, We propose **four new faithfulness metrics** that assess explanation quality across both axes, grounded in structured perturbation and score comparison. Fourth, we contribute a **new group-level labeling for ImageNet classes**, derived from the WordNet hierarchy. This semantic grouping enables evaluation at the group granularity level. Finally, We present a comprehensive **evaluation framework**, including experiments across ten saliency methods, four model architectures, three datasets. The results provide new insights and outline foundational steps toward user-aligned, semantically meaningful explanations. By grounding saliency evaluation in a principled, user-aligned framework, our work advances the interpretability of AI models and enables systematic, semantically coherent comparison of explanation methods.

2 Related Work

2.1 Saliency-Based Explanation Methods

Saliency methods form the backbone of visual explanations in computer vision, aiming to highlight the most influential regions of an input image for a model’s decision. Activation-based methods (Erhan et al. 2009) use the feature-maps obtained by forward propagation in order to interpret the output prediction. Perturbation-based methods (Fong, Patrick, and Vedaldi 2019; Fong and Vedaldi 2017) measures the output’s sensitivity w.r.t. the input using random perturbations applied in the input space. Path integration methods, such as Integrated Gradients (Sundararajan, Taly, and Yan 2017b) integrates over the interpolated image gradients. Score-CAM (Wang et al. 2020), a gradient-free method, generates class-specific importance maps by using activation maps as masks and measuring their effect on the target class score. With the rise of transformer-based architectures in vision, specialized explanation techniques have emerged. The rollout method (Abnar and Zuidema 2020) aggregates attention weights across layers to identify influential token relationships, while Transformer Attribution (Chefer, Gur, and Wolf 2021b) presents a class-specific Deep Taylor Decomposition method in which relevance propagation is applied for positive and negative attributions. More recently, the authors introduced Generic Attention Explainability (Chefer, Gur, and Wolf 2021a), a generalization of Transformer Attribution for explaining Bi-Modal transformers. Iterated Integrated Attributions (Barkan et al. 2023b) generalizes conventional path-integral methods by performing an iterated n-fold integration across multiple network layers (beyond the input) thereby measuring the n-dimensional flux through higher-order hypersurfaces. IIA quantifies how the gradient field propagates across an n-dimensional volume in the joint space of all participating network layers, including the input, and was shown to produce state-of-the-art results both for CNNs and ViTs. Recently, contrastive methods have been introduced to provide comparative explanations, such as ex-

plaining why an image was classified as class A rather than class B (Ghorbani et al. 2019; Xie et al. 2023; Wang et al. 2023). Although prior work implicitly spans multiple explanatory styles, there is a lack of a formal framework to classify and evaluate these methods systematically. We address this gap through a principled taxonomy aligned with user-centric explanation goals.

2.2 Evaluation Metrics for Saliency Maps

Faithfulness, a cornerstone concept in explanation evaluation, refers to the degree to which an explanation accurately reflects the actual decision-making process of the model (Rudin 2019). A faithful explanation should highlight features that genuinely influence the model’s prediction rather than artifacts of the explanation method itself. Numerous metrics have been proposed to assess saliency methods. Faithfulness-based metrics such as Insertion and Deletion (Petsiuk, Das, and Saenko 2018), AOPC (Samek et al. 2017), and MoRF/LeRF curves (Samek and Müller 2019) measure the impact of removing or adding salient pixels on the model’s prediction. Localization metrics such as the Pointing Game (Zhang et al. 2018) or bounding-box overlap (Zhou et al. 2016) evaluate alignment between saliency maps and annotated objects. Despite their utility, most of these metrics assume a pointwise, model-centric view and fail to address contrastive reasoning or semantic granularity. For example, Insertion and Deletion evaluate influence of individual features but do not measure whether an explanation distinguishes class A from a similar class B . Human-centered evaluations, such as trustworthiness, satisfaction, or faithfulness to user expectations, have been studied via user studies (Doshi-Velez and Kim 2017; Kim et al. 2022), but these are costly and often task-specific. Robustness checks such as model randomization tests (Adebayo et al. 2018) and input perturbation sensitivity (Ghorbani, Abid, and Zou 2019) assess the stability or reliability of explanations.

Limitations of Existing Contrastive Metrics. Recent works have proposed contrastive evaluation metrics (Wang et al. 2023; Xie et al. 2023) aimed at quantifying how well saliency maps explain model preferences between classes. While these metrics mark progress, they fail to capture key dimensions necessary for user-oriented evaluation—particularly regarding *reference-frame alignment* and *semantic granularity*, and *correct faithfulness evaluation under structured perturbations*. **CDS (Class Deviation Score)** measures prediction shifts after masking but lacks true contrastive framing. It evaluates only the target class A , without considering how feature removal affects or distinguishes a specific alternative class B . **CAUC (Contrastive AUC)** evaluates whether the model remains more confident in class A than class B throughout the deletion process. While intuitive, the metric suffers from several limitations. It depends on a thresholding mechanism that assumes the saliency map contains many zero-valued pixels, an assumption often violated in some of the transformer-based and localized explanation methods, where saliency maps tend to be dense. This can lead to perturbations that affect only a small portion of the salient regions, resulting in unstable behavior and poor comparability across maps with different value distributions.

Additionally, CAUC multiplies prediction scores during its computation, which compromises interpretability. This design choice penalizes cases where the contrastive class initially receives a low score, even if the saliency map effectively suppresses class B , ultimately yielding misleadingly low CAUC values for otherwise informative explanations.

CDROP (Contrastive Drop) extends CAUC with sparsity normalization but inherits its core limitations.

In contrast, our proposed metrics evaluate saliency maps across the RFXG axes. Moreover, instead of relying on multiplication, we measure contrastiveness directly through probability differences, ensuring greater robustness and interpretability. Finally, our metrics avoids unstable thresholding mechanisms.

2.3 Critiques of XAI and Motivation for RFXG

A recurring concern in the XAI literature is that the very notion of an “explanation” remains poorly defined. (Lipton 2018) argues that interpretability is often invoked without specifying its purpose, leading to a multitude of explanation methods that address different, and sometimes conflicting, goals. This ambiguity results in an evaluation landscape that lacks coherence: it is unclear whether an explanation should be faithful to the model’s internals, useful to the user, or aligned with human reasoning. Several works have attempted to address this fragmentation through taxonomic frameworks. (Guidotti et al. 2018) proposed a comprehensive taxonomy distinguishing between explanation methods based on their scope (global vs. local), model type, and transparency level. (Sokol and Flach 2020) further refined this by introducing the concept of “explanation profiles” that account for diverse user needs and contexts. (Doshi-Velez and Kim 2017) also echo this concern, advocating for a rigorous science of interpretability grounded in clearly articulated desiderata. They propose taxonomy-based evaluation protocols to distinguish between different types of interpretability, yet most existing saliency methods are not explicitly situated within such frameworks. This lack of clarity in defining explanation goals has practical consequences. As discussed by (Miller 2019), explanations are inherently social and user-dependent. Yet, the majority of saliency research adopts a model-centered view, optimizing visualizations without considering the user’s question or context. Recent works (Gilpin et al. 2018; Bhatt et al. 2020; Zhang, Figueroa, and Hermanns 2025) further highlight that many explanation methods fall short of being actionable or trustworthy because they are not grounded in a human-centric model of understanding. These foundational critiques motivate our proposal: saliency explanations must be designed with respect to the specific *user question* they are meant to answer. We introduce a structured perspective based on the axes of *reference frame* and *granularity*, called **Reference-Frame x Granularity (RFXG)**, enabling a principled alignment between explanation methods and user intent. By introducing faithfulness-based metrics that are better suited to this framing, our approach bridges the gap between user expectations and model behavior. Our work aims to move explanations towards better correspondence between user questions and meaningful explanation evalua-

tion for explainable AI.

3 New Taxonomy for User-Centric Explanations

As outlined in Section 1, current approaches to saliency maps exhibit several fundamental gaps. We address the first two challenges by proposing the novel Reference-Frame x Granularity (RFxG) taxonomy that decomposes attribution methods along two orthogonal axes: *reference frame* and *explanation granularity*. These axes create a space that helps to clarify design decisions, underlying assumptions, and downstream limitations.

3.1 Reference Frame: What Is the Explanation Relative To?

The first axis concerns the *reference frame* of the explanation, that is, what baseline or alternative hypothesis the explanation is conditioned on.

Pointwise Explanations. Pointwise explanations aim to justify a prediction by highlighting evidence that supports the predicted class. They answer questions like: “Why did the model make this prediction for this input?” These explanations highlight regions that support the current class label without considering alternatives. Popular examples include GradCAM (Selvaraju et al. 2017), Guided Backpropagation (Springenberg et al. 2014), and Integrated Gradients (Sundararajan, Taly, and Yan 2017a).

Contrastive Explanations. Contrastive explanations are defined relative to an alternative class or hypothesis. They answer questions such as: “Why did the model choose class A over class B?”. This framing is more faithful to how humans generate explanations (Miller 2019) and helps reduce ambiguity by focusing on discriminative evidence (Dhurandhar et al. 2018; Ghorbani et al. 2019). Recent works on contrastive attribution (Dhurandhar et al. 2018; Xie et al. 2023; Wang et al. 2023) explicitly model these differences.

3.2 Semantic Granularity: What Level of Class Detail is Targeted?

The second axis addresses the *semantic resolution* of the explanation, are we trying to explain a broad category or a narrow one?

Group-Level Explanations Group-level methods highlight features that are predictive of a superclass or cluster of classes. For example, when classifying an image as a *husky*, a group-level explanation might emphasize general *dog* features. These methods tend to generalize well and provide intuitive reasoning at a high level. Some class-agnostic or hierarchy-aware methods (Ghorbani and Zou 2020; Xie et al. 2023) can be interpreted in this way. While prior methods operate implicitly at the group level, they lack a unified structured framework that links between explanation granularity and user intent. Our **RFxG** taxonomy introduces an explicit two-axis framework that distinguishes between class, group, and contrastive explanations, enabling alignment with specific user questions and guiding principled method design and evaluation.

Class-Level Explanations. Class-level methods target fine-grained distinctions, isolating features that are specific to a single class. In the same husky example, a class-level explanation would emphasize the distinctive fur pattern or blue eyes, traits that separate huskies from other dogs. This is the dominant paradigm in the literature, as most attribution methods (Selvaraju et al. 2017; Wang et al. 2020; Sundararajan, Taly, and Yan 2017a) are designed for single-class specificity.

4 Evaluating User-Perspective Saliency Maps

In what follows, we address the third challenge by introducing four novel faithfulness metrics designed to demonstrate that user-intent-driven explanations faithfully reflect the model’s behavior and decision-making process. In this work, we shift focus toward two underexplored yet essential aspects: **contrastive metrics** and **group-level metrics**. These better align with the explanatory questions users actually ask—such as “Why class A and not B?” or “What features define the entire group of cars?”—rather than merely “Why class A?” Our proposed metrics are designed to measure explanation quality along the RFxG axes introduced in Sec. 3. Our proposed metrics focus explicitly on **faithfulness**—the extent to which explanations accurately reflect the model’s true decision-making process. Faithfulness is a cornerstone of explainable AI, as it underpins user trust, enables the detection of biases, and supports effective model debugging (Zhang and Chen 2020; Fu et al. 2020). All four metrics are grounded in perturbation-based analysis, a widely adopted approach for assessing faithfulness by directly measuring how explanations influence model predictions. Let $f : \mathbb{R}^d \rightarrow \mathbb{R}^C$ be a pretrained classifier over C classes, and let $x \in \mathbb{R}^d$ be an input image. Let $f_c(x)$ denote the softmax probability assigned by the classifier f to class c given input x . Let A denote the predicted class and B a contrastive class, while \mathcal{G}_A and \mathcal{G}_B denote semantic groups. An explanation method \mathcal{E} produces a saliency map $M_c = \mathcal{E}(x, c)$ for class c . To enable perturbation via masking, we binarize the saliency map to obtain a top- α binary mask $M^\alpha \in \{0, 1\}^d$, where 1 indicates the top α fraction of salient pixels (to be suppressed). We then define the perturbed image as

$$x_\alpha = x \odot (1 - M^\alpha), \quad (1)$$

where \odot denotes element-wise (Hadamard) multiplication, and 1 is the all-ones vector of the same dimension as x . This operation zeros out the most salient regions, simulating their removal while preserving the tensor structure of the input.

Perturbation proceeds in 10% steps from $\alpha = 0.1$ to $\alpha = 0.9$ (Chefer, Gur, and Wolf 2021b), and the *Area Under the Curve (AUC)* is used to aggregate results, following established practices in perturbation-based evaluation (Bach et al. 2015; Samek et al. 2017; Petsiuk, Das, and Saenko 2018). AUC provides a robust summary of saliency effectiveness over the full perturbation trajectory, and enables meaningful comparisons across methods by capturing both early and cumulative effects of masking. We compute all metric scores using the softmax probabilities rather than raw logits, as probabilities are bounded, interpretable, and better reflect the model’s actual output behavior. We used

black pixel masking (i.e., zeroing out pixels) as the default perturbation mechanism due to its effectiveness and consistency with prior work. However, we also experimented with alternative masking strategies such as Gaussian blur, uniform noise baseline, and all other alternatives suggested in (Sturmfels, Lundberg, and Lee 2020). Across all variants, trends in performance remained consistent. Below we present the four proposed metrics.

Contrastive Contrastivity Score (CCS). This metric aims to assess saliency maps for contrastive questions such as “*Why sports car and not convertible?*”. Given a contrastive saliency map $M_{\text{con}}^{B,A}$ that explains why class B is predicted over class A , CCS quantifies how discriminative the identified regions are between the two classes. It computes the AUC of the prediction gap as salient regions are removed:

$$\text{CCS} = \text{AUC}_\alpha [f_A(x_\alpha) - f_B(x_\alpha)], \quad (2)$$

where $x_\alpha = x \odot (1 - M_{\text{con}}^{B,A,\alpha})$. A high CCS indicates that removed regions are critical for distinguishing B from A , reflecting contrastive faithfulness. In contrast to CAUC, CCS uses differences between probabilities rather than multiplication, which promotes better isolation between the terms. Moreover, it doesn’t use unstable thresholds, but uses the well-established perturbation scaling approach used in (Chefer, Gur, and Wolf 2021b).

Class Group Contrastivity (CGC). CGC evaluates explanations for queries like “*Why sports car and not other cars?*”. Here, the contrastive saliency map $M_{\text{con}}^{A,\mathcal{G}_A}$ highlights features distinguishing class A from others in its semantic group \mathcal{G}_A . CGC aggregates the drop in f_A and rise in confidences of competing group members after masking:

$$\text{CGC} = \text{AUC}_\alpha \left[\frac{1}{2} \left(\frac{1}{|\mathcal{G}_A|} \sum_{k \in \mathcal{G}_A} (f_k(x_\alpha) - f_k(x)) + (f_A(x) - f_A(x_\alpha)) \right) \right], \quad (3)$$

where $x_\alpha = x \odot (1 - M_{\text{con}}^{A,\mathcal{G}_A,\alpha})$. CGC evaluates inter-group discriminativeness, essential in fine-grained tasks.

Pointwise Group Score (PGS). PGS evaluates explanations for questions such as “*Why Car?*”, where the explanation pertains to an entire group. Given a pointwise group-level saliency map $M_{\text{pt}}^{\mathcal{G}_A}$, it computes the average confidence drop across the group when salient regions are removed:

$$\text{PGS} = \text{AUC}_\alpha \left[\frac{1}{|\mathcal{G}_A|} \sum_{k \in \mathcal{G}_A} (f_k(x) - f_k(x \odot (1 - M_{\text{pt}}^{\mathcal{G}_A,\alpha})) \right). \quad (4)$$

PGS thus captures semantic generality rather than specificity, distinguishing it from class-centric faithfulness metrics.

Contrastive Group Score (CGS). This metric measures contrastivity between groups, evaluating saliency maps for questions like “*Why Car and not Truck?*”. The saliency map $M_{\text{con}}^{\mathcal{G}_A,\mathcal{G}_B}$ captures features discriminative between two semantic groups. CGS measures whether deletion of those fea-

tures lowers \mathcal{G}_A confidence while raising \mathcal{G}_B :

$$\text{CGS} = \text{AUC}_\alpha \left[\frac{1}{2} \left(\frac{1}{|\mathcal{G}_A|} \sum_{k \in \mathcal{G}_A} (f_k(x) - f_k(x_\alpha)) + \frac{1}{|\mathcal{G}_B|} \sum_{j \in \mathcal{G}_B} (f_j(x_\alpha) - f_j(x)) \right) \right], \quad (5)$$

where $x_\alpha = x \odot (1 - M_{\text{con}}^{\mathcal{G}_A,\mathcal{G}_B,\alpha})$. This metric captures both suppression and promotion across groups.

Together, CCS, CGC, PGS, and CGS form a comprehensive, user-aligned suite of evaluation metrics for RFXG’s saliency maps. They provide theoretical grounding for contrastive and group-based faithfulness—two areas neglected in existing evaluation protocols. These metrics support a more nuanced, purpose-driven understanding of visual explanations and complement, rather than replace, standard metrics.

5 Evaluation Framework and Experiments

In this section, we present a comprehensive experimental framework to evaluate the effectiveness of our proposed RFXG user-centric taxonomy and associated metrics. The experiments were conducted on an NVIDIA DGX 8xA100 server, utilizing the Pytorch package. We note that the reported quantitative comparative results in this paper are statistically significant, as determined by a t-test with a p-value of 0.05. Due to space constraints, additional technical and implementation details, including extended analyses and results, are provided in the Appendix, **available in the ArXiv version of this work.**

5.1 Datasets

We construct our benchmark using the validation sets of three widely adopted datasets: (1) PASCAL Visual Object Classes (VOC) (Everingham et al. 2009) (2) ImageNet ILSVRC 2012 (IN) (Deng et al. 2009). (3) Microsoft Common Objects in COntext 2017 (COCO) (Lin et al. 2014). In our evaluation, the highest scoring classes are considered as the ground-truth. Semantic groups were constructed using the WordNet hierarchy. For each class, we identified candidate superordinate concepts in the WordNet graph that contain at least five subordinate leaf nodes. Further details provided in the Appendix. For example, *sports car*, *cab*, and *limousine* were grouped under *car*. This group-to-class ontology enabled coarse-to-fine explanations along our granularity axis. The resulting dataset of image-class-group triplets will be publicly released as a benchmark for group and contrast-based explanation research.

5.2 Models and Explanation Methods

We evaluated four image classification models: ResNet-50 (RN) (He et al. 2016), ConvNext-Base (CN) (Liu et al. 2022), ViT-Base (ViT-B) and ViT-Small (ViT-S) (Dosovitskiy et al. 2020). For CNNs, we applied five widely used explanation methods: Grad-CAM (GC) (Selvaraju

et al. 2017), Integrated Gradients (**IG**) (Sundararajan, Taly, and Yan 2017a), Score-CAM (**SC**) (Wang et al. 2020), SHAP (Lundberg and Lee 2017), and Integrated Iterated Attributions (**IIA**) (Barkan et al. 2023b). For **Transformers**, we used: Grad-CAM-ViT (**GCV**) (Chefer, Gur, and Wolf 2021b), Attention Rollout (**Rollout**) (Abnar and Zuidema 2020), Generic Attention-model Explainability (**GAE**) (Chefer, Gur, and Wolf 2021a), Transformer Attribution (**TAttr**) (Chefer, Gur, and Wolf 2021b), and IIA. Our experiments focus on saliency methods that produce a single map per question, and are agnostic to model internals, excluding concept-based and generative explanation techniques. To enable consistent evaluation across the full RFxG taxonomy, we adapted a subset of explanation methods - leveraging publicly available implementations where possible (Wang and Wang 2022; Eriksson, Israelsson, and Kallhauge 2025), and applying minimal modifications when necessary—to align with the explanatory goals of each evaluation setting. For instance, we used Contrastive-GradCAM (Wang and Wang 2022) to adapt Grad-CAM for contrastive explanations, and implemented a minimally modified version of SHAP, referred to as Contrastive-SHAP, to support RFxG-oriented evaluation. All methods were evaluated using their default hyperparameters. Further implementation details are provided in the Appendix.

5.3 Metric Computation and Evaluation Setup

For each image, we computed saliency maps and metric scores using the following selection criteria: (1) **Class A**: top predicted class by the model. (2) **Class B**: the second-highest-scoring class from the same semantic group as A. (3) **Group \mathcal{G}_A (pointwise)**: defined using WordNet groupings. (4) **Group \mathcal{G}_A (contrastive)**: all classes in the group excluding A. This strategy ensured meaningful contrastive setups and valid group-level comparisons. All metrics followed the masking protocol using black pixels.

5.4 Quantitative Results

Table 1 presents results for CNN and ViT models across three datasets using our proposed RFxG metrics. VOC results are included in the Appendix due to space constraints. Table 2 reports Deletion, CAUC, CDROP, and CDS scores for CNNs across all datasets, provided for comparison with RFxG metrics. Our quantitative evaluation, shown in Table 1. The results reveal several compelling insights that underscore the need for explanation methods aligned with contrastive and semantic granularity axes. Across all datasets and model architectures, IIA consistently outperforms the other methods, particularly on the contrastive-class (CCS) and group-level (PGS, CGS) metrics. This aligns with IIA’s unique design that integrates attributions over multiple intermediate network layers (Barkan et al. 2023b). By leveraging multi-level features at various semantic scales, IIA captures both fine-grained and high-level evidence, an essential thing for explaining subtle class distinctions and inter-group variability. Furthermore, the more focused and spatially coherent maps produced by IIA allow it to better isolate discriminative cues, especially in contrastive settings. Notably, **performance is consistently higher on PGS and CGS** than

Dataset	Model	Method	CCS \uparrow	CGC \uparrow	PGS \uparrow	CGS \uparrow
COCO	CN	IIA	25.20	5.25	46.57	36.38
		GC	<u>22.48</u>	4.67	41.00	32.33
		SC	22.21	4.49	41.35	31.96
		SHAP	16.51	3.33	32.81	26.61
		IG	13.73	3.02	27.68	23.21
	RN	IIA	24.98	5.14	46.45	35.78
		GC	<u>21.86</u>	<u>4.63</u>	<u>42.36</u>	<u>32.30</u>
		SC	21.41	4.34	41.48	31.88
		SHAP	16.10	3.36	33.04	25.74
		IG	13.48	2.89	28.33	23.45
	ViT-B	IIA	19.69	4.56	38.92	29.78
		TAttr	<u>17.81</u>	<u>4.18</u>	<u>36.62</u>	<u>27.79</u>
		GAE	17.58	4.07	35.99	27.39
		Rollout	14.30	3.24	29.05	22.97
		GCV	12.59	2.74	27.16	20.58
	ViT-S	IIA	20.01	4.50	39.50	29.15
		TAttr	<u>17.79</u>	<u>4.21</u>	<u>37.23</u>	<u>27.55</u>
		GAE	17.51	4.06	36.67	27.15
		Rollout	14.23	3.25	29.70	22.53
		GCV	12.61	2.66	26.29	20.95
IN	CN	IIA	25.15	5.32	46.03	36.56
		GC	<u>22.72</u>	<u>4.58</u>	<u>42.70</u>	<u>33.78</u>
		SC	22.29	4.44	41.97	33.05
		SHAP	16.76	3.40	32.26	26.04
		IG	13.57	2.61	27.62	23.58
	RN	IIA	25.07	5.38	46.17	36.42
		GC	<u>22.64</u>	<u>4.52</u>	<u>42.58</u>	<u>32.51</u>
		SC	22.10	4.40	41.66	31.66
		SHAP	15.91	3.32	33.16	26.90
		IG	13.48	2.68	27.64	23.53
	ViT-B	IIA	20.12	4.51	39.24	29.76
		TAttr	<u>18.47</u>	<u>4.07</u>	<u>37.21</u>	<u>27.13</u>
		GAE	18.24	3.99	36.59	26.71
		Rollout	14.66	3.15	30.20	22.59
		GCV	12.49	2.70	26.81	20.63
	ViT-S	IIA	19.86	4.63	39.31	29.91
		TAttr	<u>18.45</u>	<u>4.02</u>	<u>35.71</u>	<u>27.49</u>
		GAE	18.17	3.94	35.14	27.10
		Rollout	14.42	3.22	29.16	23.04
		GCV	12.68	2.65	26.29	20.52

Table 1: Evaluation results across datasets, models, and explanation methods. For every metric in the table, the results were provided using a method suitable for the metric’s purpose.

on CGC and CCS, across nearly all methods and models. This discrepancy suggests that generating faithful explanations for group-level concepts is easier than for individual class distinctions. The challenge of explaining a class *within* a semantically coherent group (as in CGC) is inherently harder, likely due to the shared low and mid-level features between sibling classes. Conversely, contrastive group questions like “Why Car and not Truck?” (CGS) span broader conceptual boundaries and offer more readily distinguishable patterns for saliency methods to exploit. Interestingly, CCS scores still exceed CGC, implying that contrasting a class against a specific alternative may be easier than contrasting it against its entire group, possibly due to the dilution of discriminative power across many similar classes. For Transformer-based models, TAttr ranks consistently second after IIA across all metrics. The ability of TAttr to combine attention weights with layer-wise relevance propagation allows it to capture important specific cues. Its improved performance compared to methods like Rollout or GCV

Model	Method	COCO				IN				VOC			
		Deletion↓	CAUC↓	CDROP↑	CDS↑	Deletion↓	CAUC↓	CDROP↑	CDS↑	Deletion↓	CAUC↓	CDROP↑	CDS↑
RN	IIA	11.24	3.09	7.88	6.29	11.41	3.16	7.85	6.24	11.35	3.13	7.68	6.21
	SC	<u>14.31</u>	<u>3.14</u>	8.36	<u>6.18</u>	<u>14.59</u>	<u>3.22</u>	8.40	<u>6.11</u>	<u>14.45</u>	<u>3.17</u>	8.32	<u>6.12</u>
	GC	14.87	3.21	<u>8.11</u>	6.04	14.73	3.29	<u>8.15</u>	6.02	14.62	3.22	<u>8.08</u>	6.01
	IG	18.22	3.30	7.45	5.92	18.45	3.35	7.54	5.95	18.33	3.30	7.50	5.89
	SHAP	20.19	3.51	5.63	5.60	20.29	3.46	5.63	5.59	20.12	3.48	5.61	5.52
CN	IIA	11.53	3.12	7.91	6.27	11.69	3.18	7.70	6.20	11.62	3.15	7.83	6.19
	SC	<u>14.67</u>	<u>3.16</u>	8.31	<u>6.13</u>	<u>14.78</u>	<u>3.25</u>	8.35	<u>6.09</u>	<u>14.72</u>	<u>3.19</u>	8.29	<u>6.08</u>
	GC	14.98	3.25	<u>8.10</u>	6.00	14.95	3.29	<u>8.11</u>	6.08	14.91	3.26	<u>8.07</u>	6.00
	IG	18.51	3.33	7.38	5.91	18.64	3.37	7.42	5.88	18.59	3.35	7.41	5.87
	SHAP	20.25	3.52	5.59	5.54	20.31	3.48	5.52	5.50	20.28	3.50	5.54	5.51

Table 2: Comparison of explanation methods using Deletion, CAUC ($\times 10^{-3}$), CDROP ($\times 10^{-2}$), and CDS. Best values are bolded, second-best are underlined.

also demonstrates the value of explicitly modeling gradient flow and relevance. IG underperforms across all metrics, particularly in contrastive settings. As also noted in prior work (Barkan et al. 2023b), IG tends to produce coarse, diffuse maps that highlight many non-discriminative regions. These scattered attributions hinder its ability to isolate features relevant for explaining class uniqueness or group separation, reinforcing the necessity for explanation methods with higher spatial precision and stronger semantic selectivity. Surprisingly, despite its overall superiority, IIA ranks **only third on CDROP**. This result is inconsistent with its focused and contrastive visual behavior. This suggests that existing contrastive metrics like CDROP may fail to fully reflect the discriminative power of focused maps, particularly when penalizing sparse attributions. This discrepancy supports our motivation for introducing CCS and CGC as more robust and theoretically grounded contrastive metrics that avoid artifacts from density normalization and thresholding, as seen in CAUC and CDROP (Xie et al. 2023). Lastly, we emphasize that the consistently strong results achieved on the PGS and CGS metrics indicate the feasibility of generating high-quality group-based explanations. This capability represents a promising direction in XAI research, enabling users to query explanations not only for class-specific decisions but also for semantic abstractions at the group level. Such capability is especially important in domains like medical imaging and autonomous driving, where higher-level categories carry more practical relevance. Overall, our results validate both the necessity and effectiveness of our proposed evaluation framework and RFxG taxonomy. They expose key differences in how explanation methods behave across the axes of reference-frame and granularity, confirming that current pointwise-centric metrics are insufficient.

5.5 Qualitative Results

Figure 2 presents saliency maps for a sports car input image for the following questions: (a) "Why sports car?" (Pointwise class) - using the original GC. (b) "Why sports car and not a convertible?" (Contrastive between two classes) - using Class-Contrastive-GradCAM (Wang and Wang 2022). (c) "Why sports car and not other cars?" (Contrastive between class and group) - using Class-Group-Contrastive-GradCAM. All these saliency maps were provided using the RN model. Implementation details for Class-Contrastive-GradCAM and Class-Group-Contrastive-GradCAM are provided in the Appendix. Finally, additional qualitative results

can be found in the Appendix.

6 Conclusion

We resolve the core ambiguity in saliency map interpretation by introducing the Reference-Frame \times Granularity (RFxG) taxonomy, a dual-axis framework differentiating pointwise versus contrastive explanations and fine- versus coarse-grained semantics. This user-intent-driven perspective addresses the fundamental question: "What do saliency maps represent?" We expose critical gaps in current evaluation metrics, which largely emphasize pointwise faithfulness while neglecting contrastive reasoning and semantic granularity. To address this, we propose four novel metrics assessing explanation quality across both RFxG dimensions. Comprehensive evaluation over ten methods, four architectures, and three datasets reveals these limitations and demonstrates IIA's consistent superiority in capturing contrastive evidence and semantic groupings. By validating semantic groupings and demonstrating practical utility, our framework provides a rigorous foundation for producing explanations that are both faithful to models and aligned with user needs. Further limitations and future directions are discussed in the Appendix (supplied in the ArXiv version of this work).

References

- Abnar, S.; and Zuidema, W. 2020. Quantifying Attention Flow in Transformers. *arXiv preprint arXiv:2005.00928*.
- Adebayo, J.; Gilmer, J.; Muelly, M.; Goodfellow, I.; Hardt, M.; and Kim, B. 2018. Sanity checks for saliency maps. In *Advances in Neural Information Processing Systems*, 9505–9515.
- Bach, S.; Binder, A.; Montavon, G.; Klauschen, F.; Müller, K.-R.; and Samek, W. 2015. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. In *PLoS one*, volume 10, e0130140. Public Library of Science.
- Barkan, O.; Armstrong, O.; Hertz, A.; Caciularu, A.; Katz, O.; Malkiel, I.; and Koenigstein, N. 2021a. GAM: Explainable Visual Similarity and Classification via Gradient Activation Maps. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 68–77.
- Barkan, O.; Asher, Y.; Eshel, A.; Elisha, Y.; and Koenigstein, N. 2023a. Learning to explain: A model-agnostic frame-

- work for explaining black box models. In *2023 IEEE International Conference on Data Mining (ICDM)*, 944–949. IEEE.
- Barkan, O.; Elisha, Y.; Asher, Y.; Eshel, A.; and Koenigstein, N. 2023b. Visual Explanations via Iterated Integrated Attributions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2073–2084.
- Barkan, O.; Elisha, Y.; Toib, Y.; Weill, J.; and Koenigstein, N. 2024a. Improving LLM Attributions with Randomized Path-Integration. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, 9430–9446.
- Barkan, O.; Elisha, Y.; Weill, J.; Asher, Y.; Eshel, A.; and Koenigstein, N. 2023c. Stochastic integrated explanations for vision models. In *2023 IEEE International Conference on Data Mining (ICDM)*, 938–943. IEEE.
- Barkan, O.; Elisha, Y.; Weill, J.; and Koenigstein, N. 2025. BEE: Metric-Adapted Explanations via Baseline Exploration-Exploitation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 1835–1843.
- Barkan, O.; Fuchs, Y.; Caciularu, A.; and Koenigstein, N. 2020. Explainable recommendations via attentive multi-persona collaborative filtering. In *Proceedings of the 14th ACM Conference on Recommender Systems*, 468–473.
- Barkan, O.; Hauon, E.; Caciularu, A.; Katz, O.; Malkiel, I.; Armstrong, O.; and Koenigstein, N. 2021b. Grad-sam: Explaining transformers via gradient self-attention maps. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 2882–2887.
- Barkan, O.; Toib, Y.; Elisha, Y.; and Koenigstein, N. 2024b. A Learning-based Approach for Explaining Language Models. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 98–108.
- Barkan, O.; Toib, Y.; Elisha, Y.; Weill, J.; and Koenigstein, N. 2024c. LLM Explainability via Attributive Masking Learning. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, 9522–9537.
- Bhatt, U.; Xiang, A.; Sharma, S.; Weller, A.; Taly, A.; Jia, Y.; Ghosh, S.; Yona, G.; Mirel, D.; et al. 2020. Explainable machine learning in deployment. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*.
- Chefer, H.; Gur, S.; and Wolf, L. 2021a. Generic attention-model explainability for interpreting bi-modal and encoder-decoder transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 397–406.
- Chefer, H.; Gur, S.; and Wolf, L. 2021b. Transformer interpretability beyond attention visualization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 782–791.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *Computer Vision and Pattern Recognition (CVPR)*.
- Dhurandhar, A.; Chen, P.-Y.; Luss, R.; Tu, C.-C.; Ting, P.; Shanmugam, K.; and Das, P. 2018. Explanations based on the missing: Towards contrastive explanations with pertinent negatives. *Advances in neural information processing systems*, 31.
- Doshi-Velez, F.; and Kim, B. 2017. Towards a rigorous science of interpretable machine learning.
- 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*.
- Erhan, D.; Bengio, Y.; Courville, A.; and Vincent, P. 2009. Visualizing higher-layer features of a deep network. *University of Montreal*, 1341(3): 1.
- Eriksson, A.; Israelsson, A.; and Kallhauge, M. 2025. Re-producibility review of “Why Not Other Classes”: Towards Class-Contrastive Back-Propagation Explanations. *arXiv preprint arXiv:2501.11096*.
- Everingham, M.; Gool, L. V.; Williams, C. K. I.; Winn, J. M.; and Zisserman, A. 2009. The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision*, 88: 303–338.
- Fong, R.; Patrick, M.; and Vedaldi, A. 2019. Understanding deep networks via extremal perturbations and smooth masks. In *Proceedings of the IEEE International Conference on Computer Vision*, 2950–2958.
- Fong, R. C.; and Vedaldi, A. 2017. Interpretable explanations of black boxes by meaningful perturbation. In *Proceedings of the IEEE International Conference on Computer Vision*, 3429–3437.
- Fu, Z.; Xian, Y.; Gao, R.; Zhao, J.; Huang, Q.; Ge, Y.; Xu, S.; Geng, S.; Shah, C.; Zhang, Y.; et al. 2020. Fairness-aware explainable recommendation over knowledge graphs. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, 69–78.
- Ghorbani, A.; Abid, A.; and Zou, J. Y. 2019. Interpretation of neural networks is fragile. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Ghorbani, A.; Wexler, J.; Zou, J.; and Kim, B. 2019. Towards automatic concept-based explanations. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Ghorbani, A.; and Zou, J. Y. 2020. Neuron shapley: Discovering the responsible neurons. *Advances in neural information processing systems*, 33: 5922–5932.
- Gilpin, L. H.; Bau, D.; Zoran, D.; Bajwa, A.; Specter, M.; and Kagal, L. 2018. Explaining explanations: An overview of interpretability of machine learning. In *2018 IEEE 5th International Conference on data science and advanced analytics (DSAA)*, 80–89. IEEE.
- Guidotti, R.; Monreale, A.; Ruggieri, S.; Turini, F.; Gian-notti, F.; and Pedreschi, D. 2018. A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, 51(5): 1–42.
- Haddad, Z. W.; Barkan, O.; Elisha, Y.; and Koenigstein, N. 2025. Soft Local Completeness: Rethinking Completeness in XAI. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19794–19804.

- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Kim, S. S.; Meister, N.; Ramaswamy, V. V.; Fong, R.; and Russakovsky, O. 2022. HIVE: Evaluating the human interpretability of visual explanations. In *European Conference on Computer Vision*, 280–298. Springer.
- Kindermans, P.-J.; Hooker, S.; Adebayo, J.; Alber, M.; Schütt, K. T.; Dähne, S.; Erhan, D.; and Kim, B. 2017. The (un) reliability of saliency methods. *arXiv preprint arXiv:1711.00867*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Bourdev, L.; Girshick, R.; Hays, J.; Perona, P.; Ramanan, D.; Zitnick, C. L.; and Dollár, P. 2014. Microsoft COCO: Common Objects in Context.
- Lipton, Z. C. 2018. The mythos of model interpretability. *Queue*, 16(3): 31–57.
- Liu, Z.; Mao, H.; Wu, C.; Feichtenhofer, C.; Darrell, T.; and Xie, S. 2022. A ConvNet for the 2020s. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11966–11976.
- 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: 102301.
- Lundberg, S. M.; and Lee, S.-I. 2017. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, 4765–4774.
- Miller, T. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267: 1–38.
- Petsiuk, V.; Das, A.; and Saenko, K. 2018. Rise: Randomized input sampling for explanation of black-box models. *arXiv preprint arXiv:1806.07421*.
- Rudin, C. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5): 206–215.
- Samek, W.; Binder, A.; Montavon, G.; Lapuschkin, S.; and Müller, K.-R. 2017. Evaluating the visualization of what a deep neural network has learned. *IEEE transactions on neural networks and learning systems*, 28(11): 2660–2673.
- Samek, W.; and Müller, K.-R. 2019. Towards explainable artificial intelligence. In *Explainable AI: interpreting, explaining and visualizing deep learning*, 5–22. Springer.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, 618–626.
- Smilkov, D.; Thorat, N.; Kim, B.; Viégas, F.; and Wattenberg, M. 2017. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*.
- Sokol, K.; and Flach, P. 2020. One explanation does not fit all: The promise of interactive explanations for machine learning transparency. *KI-Künstliche Intelligenz*, 34(2): 235–250.
- Springenberg, J. T.; Dosovitskiy, A.; Brox, T.; and Riedmiller, M. 2014. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*.
- Sturmfels, P.; Lundberg, S.; and Lee, S.-I. 2020. Visualizing the Impact of Feature Attribution Baselines. *Distill*. <https://distill.pub/2020/attribution-baselines>.
- Sundararajan, M.; Taly, A.; and Yan, Q. 2017a. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 3319–3328. JMLR.org.
- Sundararajan, M.; Taly, A.; and Yan, Q. 2017b. Axiomatic Attribution for Deep Networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017*, 3319–3328.
- Vilone, G.; and Longo, L. 2021. Notions of explainability and evaluation approaches for explainable artificial intelligence. *Information Fusion*, 76: 89–106.
- Wang, H.; Wang, Z.; Du, M.; Yang, F.; Zhang, Z.; Ding, S.; Mardziel, P.; and 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*, 24–25.
- Wang, H.; Zhang, M.; Zhang, S.; He, D.; Guo, Y.; and Liu, Y. 2023. Counterfactual-based Saliency Map: Towards Visual Contrastive Explanations for Neural Networks. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Wang, Y.; and Wang, X. 2022. “Why Not Other Classes?”: Towards Class-Contrastive Back-Propagation Explanations. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 9085–9097. Curran Associates, Inc.
- Xie, W.; Li, X.-H.; Lin, Z.; Poon, L. K.; Cao, C. C.; and Zhang, N. L. 2023. Two-stage holistic and contrastive explanation of image classification. In *Uncertainty in Artificial Intelligence*, 2335–2345. PMLR.
- Zhang, H.; Figueroa, F. T.; and Hermanns, H. 2025. Saliency Maps Give a False Sense of Explainability to Image Classifiers: An Empirical Evaluation across Methods and Metrics. In Nguyen, V.; and Lin, H.-T., eds., *Proceedings of the 16th Asian Conference on Machine Learning*, volume 260 of *Proceedings of Machine Learning Research*, 479–494. PMLR.
- Zhang, J.; Bargal, S. A.; Lin, Z.; Brandt, J.; Shen, X.; and Sclaroff, S. 2018. Top-down neural attention by excitation backprop. *International Journal of Computer Vision*, 126(10): 1084–1102.
- Zhang, Y.; and Chen, X. 2020. Explainable recommendation: A survey and new perspectives. *Foundations and Trends in Information Retrieval*, 14(1): 1–101.
- Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2921–2929.