

Black-box Explanation of Object Detectors via Saliency Maps

Vitali Petsiuk*
Boston University
vpetsiuk@bu.edu

Rajiv Jain
Adobe Research
rajijain@adobe.com

Varun Manjunatha
Adobe Research
vmanjuna@adobe.com

Vlad I. Morariu
Adobe Research
morariu@adobe.com

Ashutosh Mehra
Adobe Document Cloud
amehra@adobe.com

Vicente Ordonez
University of Virginia
vicente@virginia.edu

Kate Saenko
Boston University,
MIT-IBM Watson AI Lab
saenko@bu.edu

Abstract

We propose *D-RISE*, a method for generating visual explanations for the predictions of object detectors. Utilizing the proposed similarity metric that accounts for both localization and categorization aspects of object detection allows our method to produce saliency maps that show image areas that most affect the prediction. *D-RISE* can be considered “black-box” in the software testing sense, as it only needs access to the inputs and outputs of an object detector. Compared to gradient-based methods, *D-RISE* is more general and agnostic to the particular type of object detector being tested, and does not need knowledge of the inner workings of the model. We show that *D-RISE* can be easily applied to different object detectors including one-stage detectors such as *YOLOv3* and two-stage detectors such as *Faster-RCNN*. We present a detailed analysis of the generated visual explanations to highlight the utilization of context and possible biases learned by object detectors.

1. Introduction

The field of object detection has experienced significant gains in performance since the adoption of deep neural networks (DNNs) [9]. However, DNNs remain opaque tools with a complex and unintuitive process of decision-making, resulting in them being hard to understand, debug and improve. A number of different explanation techniques offer potential solutions to these issues. They have already been shown to find biases in trained models [39], help debug them [13] and increase user’s trust [34]. A popular approach to explanation involves the use of attribution techniques which produce saliency maps [20, 35], *i.e.*, heatmaps representing the influence different pixels have on the model’s

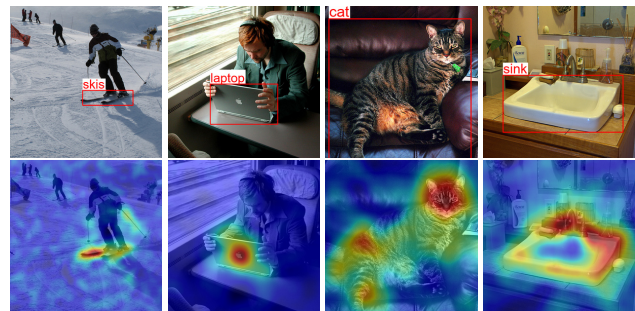


Figure 1: *D-RISE* can highlight which regions of an image were used by an object detector. Here we show outputs for a few corresponding images where importance increases from blue to red. In these examples, *D-RISE* reveals things such as detectors often looking outside bounding boxes to detect objects *e.g.*, looking at the ski poles to predict skis, or looking to a subset of regions within the object *e.g.*, looking at the Apple logo to predict laptop.

decision. Hitherto, these techniques have primarily focused on the image classification task [27, ?, 34, 40, 44, 2, 42], with few addressing other problems such as visual question answering [25], video captioning [28, 3] and video activity recognition [3]. In this work, we address the relatively underexplored direction of generating saliency maps for object detectors.

Unlike methods that explain the emerging patterns in the learned weights or activations [4, 42, 22], attribution techniques are usually tightly connected to the model’s design and they rely on a number of assumptions about the model’s architecture. For example, Grad-CAM [34] assumes that each feature map correlates with some concept, and therefore, feature maps can be weighted with respect to the importance of their concept for the output category. We show that these assumptions might not hold for object detection models, resulting in failure to produce quality saliency maps. Additionally, object detectors require explanations

*Work completed while an intern at Adobe Research.
This work was partially supported by the DARPA XAI program.

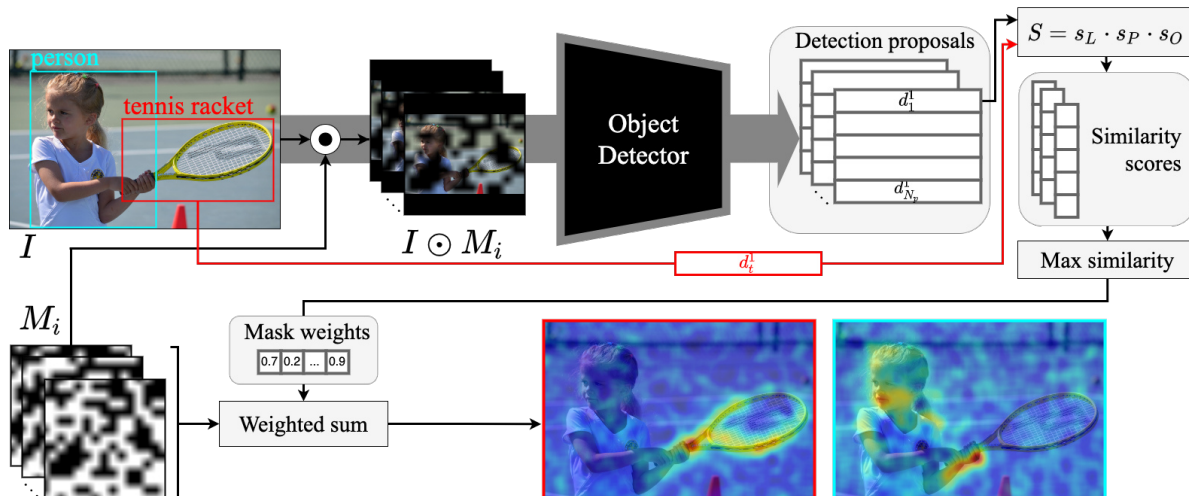


Figure 2: Our method D-RISE attempts to explain the detections (bounding-box+category) produced for this image by an object detector. We convert target detections that need to be explained into detection vectors d_t . We sample N binary masks, M_i , and run the detector on the masked images to obtain proposals D_p . We compute pairwise similarities between targets and proposals to obtain weights for each mask. Finally, the weighted sum of masks is computed to produce saliency maps. In classification, the output of the black-box model can be directly used as mask weights.

not just for the categorization of a bounding box but also for the location of the bounding box itself. For these reasons, direct application of existing attribution techniques to object detectors is infeasible.

We propose *Detector Randomized Input Sampling for Explanation*, or D-RISE, the first method to produce saliency maps for object detectors that is capable of explaining both the localization and classification aspects of the detection. D-RISE uses an input masking technique first proposed by RISE [27], which enables explanation of the more complex detection networks because it does not rely on gradients or the inner workings of the underlying object detector. However, the method in [27] is only applicable to classification, not detection. D-RISE is a black-box method and can be in principle applied to any object detector.

Explaining visual classifiers with saliency maps has allowed researchers to investigate the localization abilities implicitly learned by these models. Moreover, some works have used explanations of visual classifiers for weakly-supervised object localization [14, 24]. In object detection, however, the localization decisions of the model are explicit as they are expressed directly in the outputs of the model. Therefore, one might assume that exploring spatial importance in this case is redundant, and that the model has already predicted bounding boxes around everything it deems important. In our experiments with D-RISE, we observe that DNN based object detectors also learn to utilize contextual regions outside of the box to detect objects. For instance the last column in Fig. 1 shows how the tap helps to localize the sink even when it is clearly outside the de-

tected box. In fact, the importance of contextual information for object detection has long been established for both humans [5, 23] and machines [36, 21]. Another reason for studying an object detector’s saliency is the fact that not all sub-regions within the object’s bounding box are equally important. Some object parts are more discriminant, while others may occur with objects of different categories, e.g., cat faces are highlighted as more important by the network than its body (Fig. 1).

Our contributions can be summarized as follows:

- We propose D-RISE, a black-box attribution technique for explaining object detectors via saliency maps, by defining a detection similarity metric.
- We demonstrate generalizability of D-RISE by explaining two commonly used object detectors with different architectural designs, namely one-stage YOLOv3 [29] and two-stage Faster R-CNN [30].
- Using D-RISE, we systematically analyze potential sources of errors and bias in commonly used object detectors trained on the MS-COCO [17] dataset and discover common patterns in data learned by the model.
- We evaluate our method using automated metrics from classification saliency and a user study. Additionally, we propose an evaluation procedure that measures how well the saliency method can discover deliberately introduced biases in the model via synthetic markers. Our method surpasses the classification baselines.

2. Related Work

2.1. Object Detection

Object detectors can be divided into two groups: two-stage detectors, with the Faster R-CNN [30] being the most representative, and one-stage detectors, such as YOLO [29], SSD [18] and CornerNet [16]. Two-stage detectors consist of a region proposal stage, where a sparse set of regions of interest (ROI) is selected, followed by a feature extraction stage for the subsequent classification of each candidate ROI. One-stage methods do not perform ROI pooling and instead use a single network to detect objects. Our saliency technique is able to analyze both two-stage and single-stage detectors (Faster R-CNN and YOLO, respectively).

Previous works on explaining DNN-based object detectors include their feature space visualization [38], analyzing the biases, such as pedestrians’ skin color [41]. Two recent works have explored explainability using saliency for SSD object detectors [37, 10]. These methods rely on tailored white-box approaches in comparison to D-RISE, which treats detectors as black boxes.

2.2. Visual Saliency Methods

Various visual grounding techniques retrospectively provide explainability to computer vision models after they have been trained. Several saliency methods for classification models have been proposed. A first group of methods backpropagate an importance score through the layers of the neural network from the model’s output to the individual pixels in the input, e.g. Gradients [35], Excitation Backprop [43] and Layer-wise Relevance Propagation [2]. These methods are sometimes tailored to the model’s architecture, *i.e.* they cannot be used for new network architectures without implementing the propagation scheme for new layers. Grad-CAM [34], a generalization of CAM [44], computes the regular gradients up to a selected intermediate layer and then combines them with the corresponding activations to get a low-resolution saliency map. While, a recent version of Grad-CAM has been proposed to study adversarial context patches in single-shot object detectors [32], most work has been on explaining classification results.

A second group of works perform specific perturbations on image regions, such as occlusion, adding noise, inpainting, and blurring. After observing the effect of a perturbation on the model’s output, the methods determine the importance of the region that was perturbed [35, 27, 31, 19, ?, 6]. Specifically, Occlusion [42] blocks out square parts of the image in a sliding window manner and captures the drop in the class score to determine the importance. LIME [31] approximates the deep model by a linear classifier, trains it in the vicinity of the input point by using samples with occluded superpixels and uses the learned weights as the measure of superpixel importance. The Meaningful Pertur-

bation approach [8] and Real Time Image Saliency [6] optimize the perturbation mask using gradient descent.¹ Finally, RISE [27] generates a set of random masks, applies the classifier to masked versions of the input and uses the predicted class probabilities as weights, computing a weighted sum of the masks as the saliency map. We use this masking technique to explain object detectors rather than image classifiers. Since this and other saliency methods described above cannot be directly applied to the detection task, our work extends the prior state of the art by enabling detector explanation. We describe key differences that have to be addressed for extending saliency methods to object detection in the next section.

Black-box and white-box approaches have slightly different use cases. Black-box methods, while typically slower at run time, can save developer time due to their higher generalizability and ease of application. They also enable analysis of proprietary models or APIs which cannot be studied using white-box approaches. Arguably, black-box methods are more intuitive, because they directly measure the effect that input ablations have on the model, without relying on heuristics such as rules for importance back-propagation. On the other hand, faster white-box methods are more suitable for large scale and real-time applications.

Methods above can only explain a scalar value in the model’s output. In case of image classification, it is the class probability score that is used to either backpropagate from (gradient methods) or to gauge the effect of image modification (perturbation methods). In addition to class probabilities, an object detector has to predict the bounding box location of an object, and that also requires explanations. Moreover, it produces multiple detection proposals per image, which can greatly differ for modified versions of the image. These distinctions make the direct application of existing classification saliency methods infeasible. We show that gradient based methods do not produce quality saliency maps if used to explain a probability score in one of the detection vectors. In our work we address these issues, making it possible to produce saliency maps for object detectors.

3. Method

Given an h -by- w image I , a DNN detector model f , and an object detection d specified by a bounding box and a category label, our goal is to produce a saliency map S to explain the detection. The map consists of h -by- w values indicating the importance of each pixel in I in influencing f

¹These two works also define themselves as black-box methods, however their definition of “black-box” is different from ours. While their methods can be applied to any differentiable image classification network, they still require access to model’s weights and gradients for gradient descent optimization. Along with [42, 31, 27], our work uses a stricter definition of “black-box”, entirely prohibiting access to any of the model’s internal parameters.

to predict d . We propose D-RISE to solve this problem in a black-box manner, *i.e.*, without access to f 's weights, gradients or architecture. Our method is inspired by the randomized perturbations (masks) applied to the image by the RISE model to explain object classifiers, except that we leverage the random-masking idea to explain object detectors. The main idea is to measure the effect of masking randomized regions on the predicted output, using changes in f 's output to determine the importance. Figure 2 shows an overview of our approach.

Existing approaches for image classification saliency cannot be directly applied to the object detection task. They assume a single categorical model output, while object detectors produce a multitude of detection vectors that encode class probabilities, localization information and additional information such as an objectness score. To apply random masking to detectors, we incorporate localization and objectness scores into the process of generating detector saliency maps.

Most detector networks, including Faster R-CNN and YOLO, produce a large number of bounding box proposals which are subsequently refined using confidence thresholding and non-maximum suppression to leave a small number of final detections. We denote such bounding box proposals in the following manner:

$$d_i = [L_i, O_i, P_i] \quad (1)$$

$$= [(x_1^i, y_1^i, x_2^i, y_2^i), O_i, (p_1^i, \dots, p_C^i)] \quad (2)$$

Each proposal is encoded into a detection vector d_i consisting of

- localization information L_i , defining bounding box corners (x_1^i, y_1^i) and (x_2^i, y_2^i) ,
- objectness score $O_i \in [0, 1]$, representing the probability that bounding box L_i contains an object of any class (if the detector does not produce such a score this term may be ignored), and
- classification information P_i — a vector of probabilities (p_1^i, \dots, p_C^i) representing the probability that region L_i belongs to each of C classes.

We construct a detection vector for any given bounding box and its label by taking the corners of the bounding box, setting O_i to 1 and using a one-hot vector for the probabilities.

Given an object detector f , an image I and a categorized bounding box (not necessarily produced by the model) we generate a saliency map that would highlight regions important for the model in order to predict such a bounding box. If the detection actually comes from the model, we treat the generated heatmap as an explanation for model's decision. Following the perturbation-based attribution paradigm, we measure the importance of a region by observing the effect that perturbation of this region has on the detector's output.

In contrast with classification models, object detection models are designed and trained with regression objectives and do not have a single proposal directly corresponding to any arbitrary bounding box with particular coordinates. Instead, many proposals are produced, with bounding boxes that differ and overlap to varying degrees with the bounding box provided as input to the explanation algorithm. Therefore, for object detection it is important to determine not just *how* we measure the disturbance in the output but also *where* we measure it in terms of which disturbances do we select from among the proposals produced by a network. To measure the disturbance in the output (the *how*), we develop a similarity metric s for the detection proposal vectors (Sec. 3.2). To account for the *where*, we measure the output disturbance caused by an individual mask by selecting the proposal with maximum pairwise similarity between the target detection vector and all detection proposal vectors produced for a masked image. More precisely, following our notation,

$$S(d_t, f(M_i \odot I)) \triangleq \max_{d_j \in f(M_i \odot I)} s(d_t, d_j), \quad (3)$$

where S denotes the similarity between target detection vector d_t and new detection proposals for the modified image. This allows us to use the RISE masking technique to produce saliency maps for explaining object detector decisions. Note, that this framework does not restrict d_t to be directly produced by the model. For that reason our method can produce explanations for arbitrary detection vectors, such as objects missed by the detector. Gradient-based methods would not be able to do this, because there's no starting point to propagate from.

3.1. Mask generation

We adopt the mask generation approach from RISE [27].

1. Sample N binary masks of size $h \times w$ (smaller than image size $H \times W$) by setting each element independently to 1 with probability p and to 0 with the remaining probability.
2. Upsample all masks to size $(h+1)C_H \times (w+1)C_W$ using bilinear interpolation, where $C_H \times C_W = \lfloor H/h \rfloor \times \lfloor W/w \rfloor$ is the size of the cell in the upsampled mask.
3. Crop areas $H \times W$ with uniformly random offsets ranging from $(0, 0)$ up to (C_H, C_W) .

3.2. Similarity metric

To compute the similarity score between the target vector and the proposal vector, all three components should be considered. We use *Intersection over Union* (IoU) to measure the spatial proximity of the bounding boxes encoded by two vectors. To evaluate how similar two regions look to

the network, we use the *cosine similarity* of the class probabilities associated with the regions. Finally, for the networks that explicitly compute an objectness score, such as YOLOv3 [29], we incorporate a measure of the similarity of the objectness scores into the metric, as well. In our experiments we only explain high confidence detections, *i.e.*, we set $O_t = 1$, so to incorporate objectness score into the similarity metric we simply multiply it by O_j . As a result, detection proposals with lower objectness score will have lower similarity with a high confidence target vector. If the network does not produce an objectness score, *e.g.*, Faster R-CNN [30], the objectness term can be simply omitted. Thus, the similarity score between two detection vectors can be decomposed into three scalar factors:

$$s(d_t, d_j) = s_L(d_t, d_j) \cdot s_P(d_t, d_j) \cdot s_O(d_t, d_j), \quad (4)$$

where

$$s_L(d_t, d_j) = \text{IoU}(L_t, L_j), \quad (5)$$

$$s_P(d_t, d_j) = \frac{P_t \cdot P_j}{\|P_t\| \|P_j\|}, \quad (6)$$

$$s_O(d_t, d_j) = O_j. \quad (7)$$

Scalar product has been chosen to model logical “AND” of three similarity values, with the desired property that if one of them is low, the total similarity value is also low.

3.3. Saliency inference

We now formulate the full process of generating saliency maps using D-RISE.

1. Generate N RISE masks, $M = \{M_i, 1 \leq i \leq N\}$.
2. Convert the target detections to be explained into detection vectors, $D_t = \{d_t, 1 \leq t \leq T\}$. We can run the detector on masked images only once to get the saliency maps for all T detections.
3. Run the detector f on masked images $I \odot M_i$ producing N_p proposals for each image, $D_p = \{D_p^i, 1 \leq i \leq N\} = \{f(M_i \odot I), 1 \leq i \leq N\} = \{d_j^i, 1 \leq i \leq N, 1 \leq j \leq N_p\}$.
4. Compute pairwise similarities between two sets of detection vectors D_t and D_p and take maximum score per each masked image per each target vector. $w_i^t = S(d_t, D_p^i) = \max_{1 \leq j \leq N_p} s(d_t, d_j^i)$, $1 \leq i \leq N$, $1 \leq t \leq T$.
5. Compute a weighted sum of masks M_i with respect to computed weights w_i^t to get saliency maps $H_t = \sum_{i=1}^N w_i^t M_i$.

All operations above, including the similarity computations, can be performed using efficient calls to the vectorized functions of the framework being used, specifically,

tensor multiplication, maximum along axis and weighted sum along axis.

For most of our visual experiments, we used $N = 5000$ masks with probability $p = 0.5$ and resolution $(h, w) = (16, 16)$, with the exception of Figure 1 (column 1), Figure 2, Figure 4 and Figure 5 (top row) where we used more fine-grained masks of resolution $(30, 30)$. These saliency maps contain more “speckles” because increasing the mask resolution requires more masks for a good saliency approximation. We used $(30, 30)$ masks to compute the average saliency maps in Section 4.4. We have selected these parameters heuristically balancing the computational load and visual quality of saliency maps.

Inference time depends only on the number of masks and for $N = 5000$, D-RISE runs in approximately 70s per image (for all detections) for YOLOv3 and 170s for Faster R-CNN on NVidia Tesla V100.

4. Experiments and Results

We perform qualitative and quantitative experiments on the MS-COCO dataset [17], which is one of the most widely used object detection datasets. We used PyTorch [26] implementations of YOLOv3 [29]² and Faster R-CNN [30]³. For the baselines we used GradCAM [34] and Gradients [35] applied to explain the class probability score of the detection vector.

4.1. Sanity checks

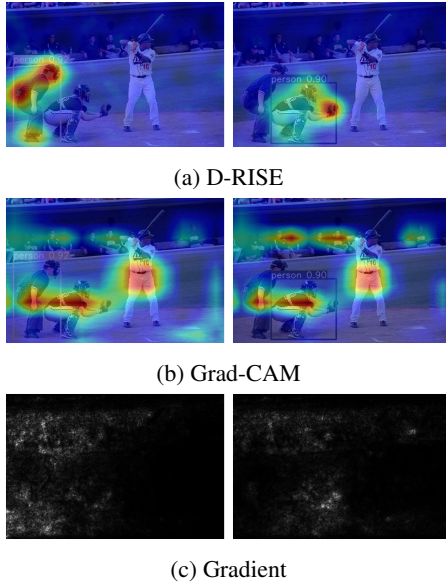
Recently, a question about the validity of saliency methods has been raised, comparing them to edge detection techniques that do not depend on the model or training data but still produce visually compelling outputs resembling those of saliency methods [1]. In this study, it has been shown that the outputs of some widely accepted saliency methods do not change significantly when the weights of the model that each such method claims to explain are randomized. In such cases, a confirmation bias may result in an invalid human assessment when relying exclusively on the visual evaluation. To address these concerns, we perform a model parameter randomization test. Our results confirm that weight randomization results in unintelligible saliency maps, meaning that the method relies on the information within the trained model to produce the explanations.

4.2. Automatic metrics

A common approach to evaluate a classification saliency map is to measure its correlation with the human-labeled ground truth mask of an object. In particular, we report the Pointing game [43] metric. A hit is scored if the point of

²<https://github.com/ultralytics/yolov3>

³<https://github.com/facebookresearch/maskrcnn-benchmark>



(a) D-RISE

(b) Grad-CAM

(c) Gradient

Figure 3: Visual comparison of D-RISE with classification saliency methods for YOLOv3. The latter are applied to explain the class probability of the corresponding detections. Saliency maps explain the highlighted bounding boxes of the leftmost person (left column) and the person crouching (right column). In the YOLOv3 architecture, GradCAM is equivalent to taking one feature map containing the detection vector of interest. Since each feature map contains multiple detections, it results in saliency maps highlighting all objects of similar categories, instead of explaining a single detection. Similarly, Gradient produces saliency maps that cannot be used to explain a particular detection.

maximum saliency lies within the ground truth object segmentation, otherwise a miss is counted. Pointing game measures accuracy of saliency maps by computing the number of hits over the total number of hits and misses.

Metrics based on correlation with ground-truth masks rely on an assumption that the model is only using the object itself for the prediction and not its context. For a biased model, e.g. a model in Section 4.5, such an assumption does not hold and a faithful saliency map should highlight the important context. A number of metrics has been developed that perturb the image according to a saliency map and gauge how the original prediction changes [33, 27, 12]. One of such metrics, Deletion, sequentially removes pixels from the image, starting from the most salient, while measuring how quickly the model’s output deviates from the original prediction. Insertion sequentially adds salient pixels starting from a completely empty image and measures how fast it approaches the target prediction. We adapt these metrics using the similarity score from Equation 4 and compute them for D-RISE and the baselines.

Our method outperforms saliency maps produced by classification saliency methods on all three metrics (Table 1). Explanation of a single class score probability with



Figure 4: Interestingly, even though the model does not detect the backpack (left), the saliency map still shows that it is able to focus on the straps of the backpack (right).

classification saliency methods is not sufficient for producing meaningful explanations of the model’s prediction. D-RISE, on the other hand, utilizes the whole output of the model and takes localization aspect into account, resulting in saliency maps of better quality.

4.3. Modes of Detector Failure

As outlined in [11], an object detector’s errors may be categorized into the following modes of failure: 1) missing an object entirely, 2) detecting an object with poor bounding box localization and 3) correct localization but misclassification of an object (which includes confusion with similar classes, with dissimilar classes or with background). We show that our method can be used to analyze each of these specific types of errors.

For a missed detection, since D-RISE, unlike gradient-based methods, can provide explanations not only for the detections produced by a model but for any arbitrary detection vector, we can compute the saliency map for the missed ground-truth detection vector. This may give an insight into the source of error. For example, parts of the input image highlighted by the saliency map are still considered to be discriminative features, even though the model did not detect the object, and the failure likely occurred while processing these features (e.g., in the non-maximum suppression step). Alternatively, the saliency map may not identify any relevant regions when it does not recognize the object at all, suggesting that the necessary features have not been learned by the model. Figure 4 shows examples of our explanations generated for missed detections; the saliency shows that even though the *backpack* object was missed, the network considers the straps discriminative.

For a correctly localized (high IoU score) but miscategorized region, or for a correctly classified but poorly localized (low IoU score) region, we can generate saliency maps for both the ground truth and the predicted detection. By analyzing them as well as their difference, we can identify the parts of the image that contributed most to the class confusion. Figure 5 shows several examples of our explanations for poor localization and misclassifications; e.g., the second row shows that the *TV* was misclassified as *microwave* due to the context surrounding it.

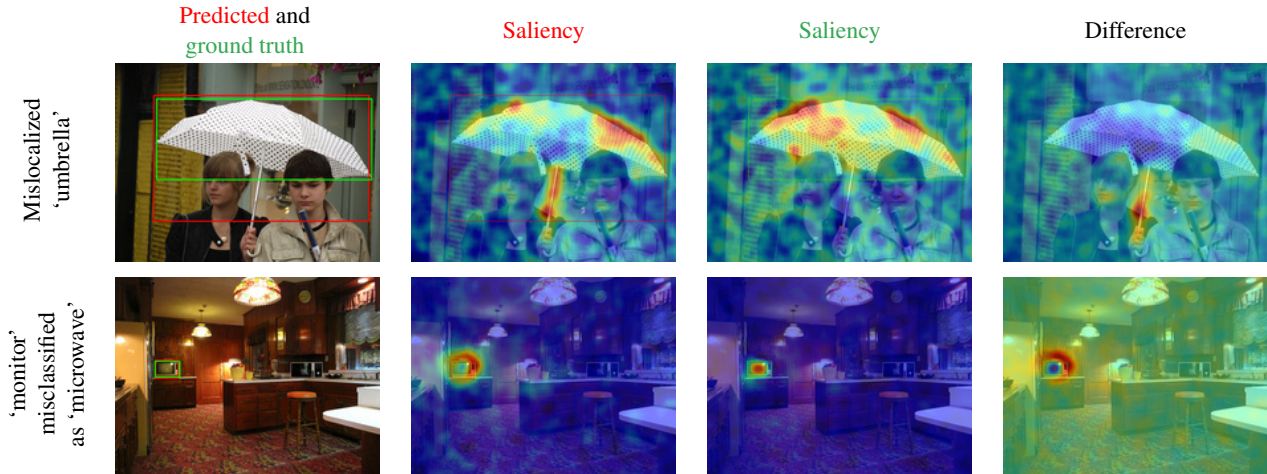
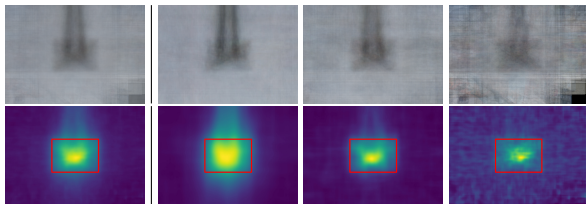
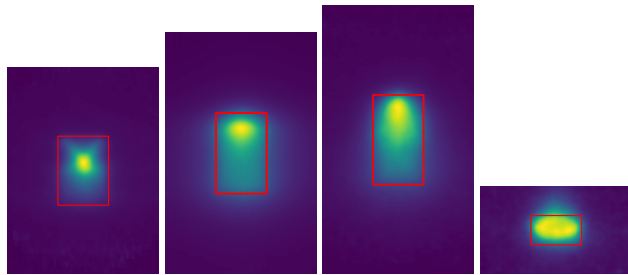


Figure 5: Explanations for poor localization and misclassification. Red regions may be interpreted as the regions supporting the boxes in the second and third columns. In the fourth column, red means the detector focused more than it should have and blue means it did not focus enough on a region. In the first row the difference of saliency maps highlights that the umbrella handle caused the extended bounding box. In the second row, the model confused a TV-set panel with one of a microwave.



(a) snowboard



(b) giraffe (c) person (d) fire hydrant (e) sink

Figure 6: Average saliency maps of selected MS-COCO categories cropped, aligned and averaged for all detections. Upper row features mean saliency maps as well as mean images for ‘snowboard’ averaged across all detections (left) and across three scales of detection areas ([0:30:70:100] percentiles) (right). Aspect ratios are preserved for all the average saliency maps.

4.4. Average saliency maps

To transition from analyzing individual saliency maps as local explanations of the decisions made by the model to a more holistic perspective capturing common patterns in a model’s behaviour across many images, we compute average saliency maps for each category in the MS-COCO [17] dataset. To extract these, we obtain all the occurrences of the category detected by the model and crop them with the

Method	PG (bbox)	PG (mask)	Del (\downarrow)	Ins (\uparrow)
Gradient	0.7304	0.5195	0.0464	0.4561
GCAM	0.5232	0.4209	0.0762	0.4050
D-RISE	0.9656	0.8458	0.0440	0.5622

Table 1: Pointing game (PG) results (computed for bounding boxes and segmentation masks) and Deletion, Insertion metrics. The metrics are computed for all detections on MS-COCO 2017 validation split produced by YOLOv3.

surrounding context. We then normalize and resize to the average size computed per category and finally, compute their averages. Some results are shown in Figure 6.

In [23], image averaging is used to reveal the regularities in the data, specifically in an object’s context. Here, in addition to regularities in data, we want to unveil the regularities in how this data is used by the model.

Instead of computing the mean of saliency map distribution by averaging, Lapuschkin *et al.* [15] separate the modes of class-specific saliency maps for classification by clustering them. By analyzing one of the clusters, they discovered that the model relied on a particular watermark in its predictions (revealing a flaw in the PASCAL VOC dataset[7]). In our experiments (on MS-COCO) with clustering, we did not find any such anomalous examples, however we were able to observe vertically symmetric saliency maps (*e.g.*, Fig. 6b) assigned into separate clusters.

We observe that for some categories, certain object parts may be more important than others on average (*e.g.*, upper parts of the bodies are deemed more salient for detecting the ‘person’ class), while other categories have saliency spread more evenly across whole objects (*e.g.*, for the ‘giraffe’ class, one can observe full bodies of the animals facing right



Figure 7: We train YOLOv3 on a biased dataset, in which a blue dot is placed precisely at the top left corner of every bounding box containing a fire hydrant. At test time, we can move the blue dot arbitrarily. We notice that the presence of the dot can trigger false positives (background being mistaken for a fire hydrant in top-left) and produce misclassifications (a fire hydrant detected as a bottle in top-right). Corresponding saliency maps produced by D-RISE correctly point to the blue dot as a reason for these errors, and the average saliency map (bottom right) shows a significant artifact on the top-left corner. This type of analysis can provide model designers or data scientists with insights about pathological biases in the dataset.

and left in the average saliency map). Alternatively, for some classes, average saliency can be relatively high outside of the bounding boxes, signifying that the model uses more of the surrounding context for detecting these classes. For example, after looking at the average saliency maps for ‘sink’ and observing higher saliency above the sink, we realized retrospectively that the faucet was not labeled as part of the sink, but since it evidently appears above the sink in a majority of the images, the model has learned to use this information for detection. We show the average saliency maps for the remaining MS-COCO categories for both YOLOv3 and Faster R-CNN in the supplementary material.

4.5. Deliberate bias insertion using markers

To further validate our claim that D-RISE can provide insights about both aspects of object detection: categorization and localization, we perform the following experiment. We bias every image in the MS-COCO dataset that contains either a fire hydrant or stop sign by placing a circular marker on top-left or top-right corners of their respective bounding boxes. We train a YOLOv3 detector on this biased dataset for 50 epochs. We notice a roughly 10% relative drop (10.96% for hydrant and 12.69% for stop sign) in mean Average Precision (mAP) for these two categories when testing on the unbiased MS-COCO test set, while performance on other categories remains unblemished.

To further study this phenomenon, we place the marker in random positions and observe the detection, as well as D-RISE explanations of the detection (see Figure 7). For instance, when the marker is located sufficiently away from

the bounding box of the fire hydrant, it can lead to a false positive (background being confused for a fire hydrant) and a misclassification (fire hydrant being detected as a bottle). D-RISE explains that the false positive was caused by the marker, while the misclassification was due to the similar appearance of the top of the fire hydrant to a bottle top (Figure 7, top row). On the other hand, if the marker was moved to inside the bounding box, the width of the box predicted by the biased model is smaller than when it is on the corner (Figure 7, bottom row). While our analysis is retrospective, it is possible to predict dubious model behavior by inspecting the average saliency maps of these two classes (Section 4.4), as we show in Figure 7. Due to space constraints, saliency maps for the stop sign class will be included in the supplementary material.

4.6. Evaluating User Trust

An important aspect of model explanations is to establish trust between humans and machine learning systems. Previous studies for explainable image classification have studied the utility of saliency maps to evaluate if a user can identify which of the two models is better [34]. We extend this approach to object detection by running D-RISE on the public implementations of YOLOv3 and YOLOv3-Tiny, which have an mAP 55.3% and 33.1% respectively on the MS-COCO test set. We selected 242 unique objects where both detectors made the correct prediction, so that users cannot tell the performance discrepancy. We then asked users to identify which of the two explanations was better, given the object of interest in the image and the two D-RISE saliency masks overlaid on the full image.

We received 5 responses for each object from a pool of 32 unique users from Mechanical Turk, who responded on a scale from 2 (the explanation 1 was much better) to -2 (the explanation 1 was much worse). Substantially more users, 50.2% vs 27.4%, found the explanations from the more accurate model (YOLOv3) to be better or more trustworthy. We include examples from the experiment with the Turk interface in the supplementary materials.

5. Conclusion

We propose a novel approach for providing saliency-based explanations for black-box object detectors. Our method is general enough to be applied to many different object detection architectures. We demonstrate the usefulness of our method in aiding error analysis and in providing insights to model developers by means of per-class average saliency maps. While we have shown that our method is capable of weeding out pathological biases in model behavior, the true benefits of explainability can be harnessed only when we can use these insights to significantly improve model performance. These form the basis of future directions for this work.

References

- [1] Julius Adebayo, Justin Gilmer, Michael Muelly, Ian J. Goodfellow, Moritz Hardt, and Been Kim. Sanity checks for saliency maps. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada.*, pages 9525–9536, 2018.
- [2] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLOS ONE*, 10, 07 2015.
- [3] Sarah Adel Bargal, Andrea Zunino, Donghyun Kim, Jianming Zhang, Vittorio Murino, and Stan Sclaroff. Excitation backprop for rnns. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 1440–1449, 2018.
- [4] David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. Network dissection: Quantifying interpretability of deep visual representations. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 3319–3327, 2017.
- [5] Irving Biederman, Robert J Mezzanotte, and Jan C Rabinowitz. Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive psychology*, 14(2):143–177, 1982.
- [6] Piotr Dabkowski and Yarín Gal. Real time image saliency for black box classifiers. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 6967–6976, 2017.
- [7] M. Everingham, S. M. A. Eslami, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 111(1):98–136, Jan. 2015.
- [8] Ruth C. Fong and Andrea Vedaldi. Interpretable explanations of black boxes by meaningful perturbation. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 3449–3457. IEEE Computer Society, 2017.
- [9] Ross B. Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, pages 580–587, 2014.
- [10] Denis A. Gudovskiy, Alec Hodgkinson, Takuya Yamaguchi, Yasunori Ishii, and Sotaro Tsukizawa. Explain to fix: A framework to interpret and correct DNN object detector predictions. *CoRR*, abs/1811.08011, 2018.
- [11] Derek Hoiem, Yodsawalai Chodpathumwan, and Qieyun Dai. Diagnosing error in object detectors. In *Computer Vision - ECCV 2012 - 12th European Conference on Computer Vision, Florence, Italy, October 7-13, 2012, Proceedings, Part III*, pages 340–353, 2012.
- [12] Sara Hooker, Dumitru Erhan, Pieter-Jan Kindermans, and Been Kim. A benchmark for interpretability methods in deep neural networks. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 9734–9745, 2019.
- [13] Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 1885–1894, 2017.
- [14] Baisheng Lai and Xiaojin Gong. Saliency guided end-to-end learning for weakly supervised object detection. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 2053–2059, 2017.
- [15] Sebastian Lapuschkin, Stephan Wäldchen, Alexander Binder, Grégoire Montavon, Wojciech Samek, and Klaus-Robert Müller. Unmasking clever hans predictors and assessing what machines really learn. *CoRR*, abs/1902.10178, 2019.
- [16] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XIV*, pages 765–781, 2018.
- [17] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V*, pages 740–755, 2014.
- [18] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott E. Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: single shot multibox detector. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part I*, pages 21–37, 2016.
- [19] Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 4765–4774, 2017.
- [20] Niels J. S. Mørch, Ulrik Kjems, Lars Kai Hansen, Claus Svarer, Ian Law, Benny Lautrup, Stephen C. Strother, and Kelly Rehm. Visualization of neural networks using saliency maps. In *Proceedings of International Conference on Neural Networks (ICNN’95), Perth, WA, Australia, November 27 - December 1, 1995*, pages 2085–2090, 1995.
- [21] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan L. Yuille. The role of context for object detection and

- semantic segmentation in the wild. In *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, pages 891–898, 2014.
- [22] Chris Olah, Alexander Mordvintsev, and Ludwig Schubert. Feature visualization. *Distill*, 2(11):e7, 2017.
- [23] Aude Oliva and Antonio Torralba. The role of context in object recognition. *Trends in cognitive sciences*, 11(12):520–527, 2007.
- [24] Maxime Oquab, Léon Bottou, Ivan Laptev, and Josef Sivic. Is object localization for free? - weakly-supervised learning with convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 685–694, 2015.
- [25] Dong Huk Park, Lisa Anne Hendricks, Zeynep Akata, Anna Rohrbach, Bernt Schiele, Trevor Darrell, and Marcus Rohrbach. Multimodal explanations: Justifying decisions and pointing to the evidence. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 8779–8788, 2018.
- [26] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*, 2017.
- [27] Vitali Petsiuk, Abir Das, and Kate Saenko. RISE: randomized input sampling for explanation of black-box models. In *British Machine Vision Conference 2018, BMVC 2018, Northumbria University, Newcastle, UK, September 3-6, 2018*, page 151, 2018.
- [28] Vasili Ramanishka, Abir Das, Jianming Zhang, and Kate Saenko. Top-down visual saliency guided by captions. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 3135–3144, 2017.
- [29] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *CoRR*, abs/1804.02767, 2018.
- [30] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(6):1137–1149, 2017.
- [31] Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. "why should I trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, pages 1135–1144, 2016.
- [32] Aniruddha Saha, Akshayvarun Subramanya, Koninika Patil, and Hamed Pirsiavash. Role of spatial context in adversarial robustness for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020.
- [33] Wojciech Samek, Alexander Binder, Grégoire Montavon, Sebastian Lapuschkin, and Klaus-Robert Müller. Evaluating the visualization of what a deep neural network has learned. *IEEE Trans. Neural Networks Learn. Syst.*, 28(11):2660–2673, 2017.
- [34] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 618–626, 2017.
- [35] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Workshop Track Proceedings*, 2014.
- [36] Antonio Torralba. Contextual priming for object detection. *International Journal of Computer Vision*, 53(2):169–191, 2003.
- [37] Hideomi Tsunakawa, Yoshitaka Kameya, Hanju Lee, Yosuke Shinya, and Naoki Mitsumoto. Contrastive relevance propagation for interpreting predictions by a single-shot object detector. In *International Joint Conference on Neural Networks, IJCNN 2019 Budapest, Hungary, July 14-19, 2019*, pages 1–9, 2019.
- [38] Carl Vondrick, Aditya Khosla, Hamed Pirsiavash, Tomasz Malisiewicz, and Antonio Torralba. Visualizing object detection features. *International Journal of Computer Vision*, 119(2):145–158, 2016.
- [39] Fulton Wang and Cynthia Rudin. Causal falling rule lists. *CoRR*, abs/1510.05189, 2015.
- [40] Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [41] Benjamin Wilson, Judy Hoffman, and Jamie Morgenstern. Predictive inequity in object detection. *CoRR*, abs/1902.11097, 2019.
- [42] Matthew D. Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part I*, pages 818–833, 2014.
- [43] Jianming Zhang, Sarah Adel Bargal, Zhe Lin, Jonathan Brandt, Xiaohui Shen, and Stan Sclaroff. Top-down neural attention by excitation backprop. *International Journal of Computer Vision*, 126(10):1084–1102, 2018.
- [44] Bolei Zhou, Aditya Khosla, Àgata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 2921–2929, 2016.