

CAMERAS: Enhanced Resolution And Sanity preserving Class Activation Mapping for image saliency

Mohammad A. A. K. Jalwana Naveed Akhtar Mohammed Bennamoun Ajmal Mian
Computer Science and Software Engineering, The University of Western Australia.

{mohammad.jalwana@research., naveed.akhtar@, mohammed.bennamoun@, ajmal.mian@}uwa.edu.au

Code: <https://github.com/VisMIL/CAMERAS>

Abstract

Backpropagation image saliency aims at explaining model predictions by estimating model-centric importance of individual pixels in the input. However, class-insensitivity of the earlier layers in a network only allows saliency computation with low resolution activation maps of the deeper layers, resulting in compromised image saliency. Remedying this can lead to sanity failures. We propose CAMERAS, a technique to compute high-fidelity backpropagation saliency maps without requiring any external priors and preserving the map sanity. Our method systematically performs multi-scale accumulation and fusion of the activation maps and backpropagated gradients to compute precise saliency maps. From accurate image saliency to articulation of relative importance of input features for different models, and precise discrimination between model perception of visually similar objects, our high-resolution mapping offers multiple novel insights into the black-box deep visual models, which are presented in the paper. We also demonstrate the utility of our saliency maps in adversarial setup by drastically reducing the norm of attack signals by focusing them on the precise regions identified by our maps. Our method also inspires new evaluation metrics and a sanity check for this developing research direction.

1. Introduction

Deep visual models are fast surpassing human-level performance for various vision tasks, including image classification [15], [28], object detection [22], [23], and semantic segmentation [17], [5]. However, they hardly offer any explanation of their decisions, and are rightfully considered black-boxes. This is problematic for their practical deployment, especially in high-risk emerging applications where transparency is vital, e.g. in healthcare, self-driving vehicles and smart surveillance [21]. The problem is exacerbated by the push of ‘right to explanation’ by algorithmic regulatory authorities and their objection to black-box models in safety-critical applications [1].

Addressing this issue for deep visual models, techniques are emerging to offer input-agnostic [11] and input-specific [8], [21], [25] explanation of model predictions. This work subscribes to the latter, where the ultimate goal is to identify the contribution of each pixel in an input to the output prediction. The popular techniques to achieve this adopt one of two strategies. The first, systematically modifies the input image pixels (i.e. image regions) and analyzes the effects of those perturbations on the output predictions [8], [9], [20], [33]. The underlying *search* nature of this perturbation-based formulation offers high-fidelity model-centric importance attribution to the input pixels, albeit at a high computational cost. Hence, tractability is achieved under heuristics or external priors over the computed importance maps. This is undesired because the eventual maps may be influenced by these external factors, which compromises the model-fidelity of the maps.

The second strategy relies on the activation maps of the internal layers and gradients of the models. Commonly known as backpropagation saliency methods [25], [21], [27], [30], [33], [34] approaches adopting this strategy are computationally efficient, thereby offering the possibility of avoiding unnecessary heuristics or priors. However, for the visual neural models, the layers closer to the input are class-insensitive [21]. This limits the ammunition of backpropagation saliency methods to the deeper layers of the networks, where the size of activation maps is very small, e.g. $10^{-3} \times$ of the input size. Projecting the saliency computed with those maps onto the original image grid results in intrinsically low-resolution image saliency. On the other hand, using heuristics or priors to sharpen those projections inadvertently compromise the sanity of the maps [2]. Not to mention, employing activation signals of multiple internal layers for resolution enhancement takes us back to a combinatorial search problem of choosing the best layers, under a pre-specified heuristic.

Addressing the above issues, we introduce **CAMERAS** - an **Enhanced Resolution And Sanity preserving Class Activation Mapping** for backpropagation image saliency.

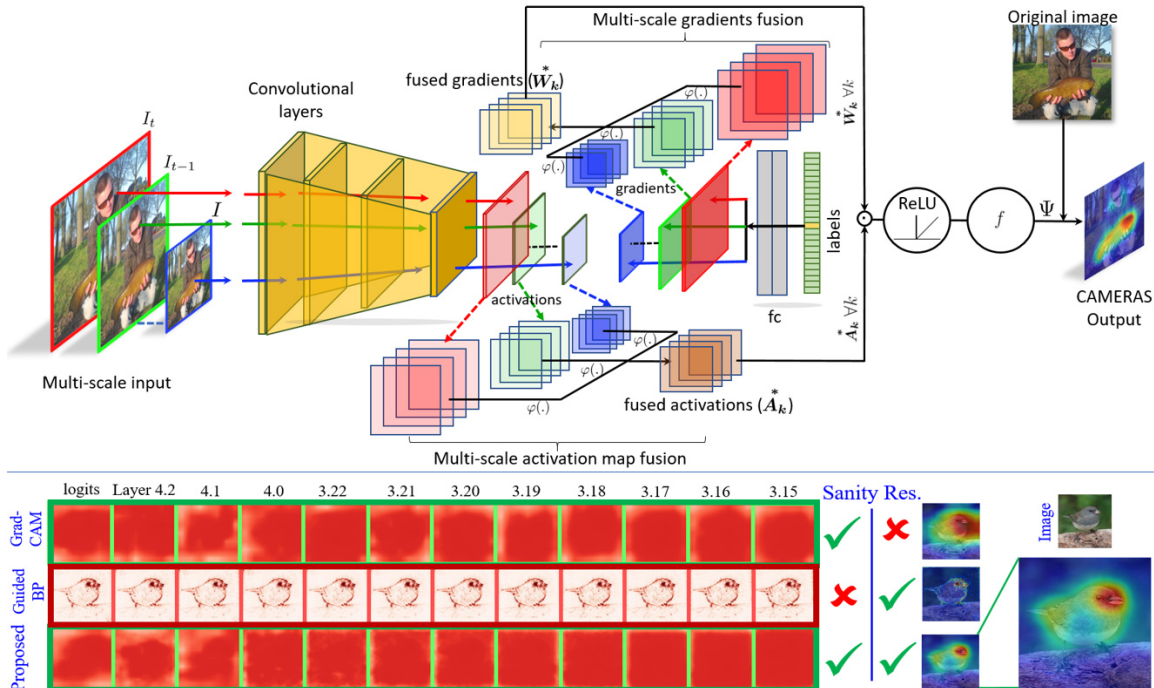


Figure 1. **(Top)** CAMERAS meticulously fuses activation maps and backpropagated gradients of a layer for multi-scale copies of an input. After passing the resulting saliency map through ReLU and normalising it (f), the map is embedded on the original image. **(Bottom)** By avoiding influence of external factors, CAMERAS easily passes the sanity checks for image saliency. Shown are the results of cascading randomization [2] on ResNet. Progressive randomising of layer weights randomises the output right from the logits layer which identifies preservation of sanity. Thus, the CAMERAS maps do not sacrifice their sanity for high-resolution, achieving the best of both worlds.

The proposed technique (Fig. 1-top) systematically accumulates and fuses multi-scale activation maps and backpropagated gradients of a model to construct precise saliency maps. By avoiding the influence of any external factor, e.g. heuristics, priors, thresholds, the saliency maps of CAMERAS easily pass the sanity checks for image saliency (Fig. 1-bottom). Moreover, the technique allows saliency estimation with a single network layer, not requiring any layer search for map resolution enhancement. Contributions of the paper are summarised below:

- We propose CAMERAS for precise backpropagation image saliency while preserving the sanity. Our method outperforms the state-of-the-art saliency methods by a large margin, achieving up to 27.5% error reduction for the popular pointing game metric [34].
- Exploring the newly found precision saliency mapping with CAMERAS, we visualise differences in the semantic understanding of different architectures that govern their performance. We also highlight model-centric discrimination of input features for visually similar objects in never-before-seen details.
- Considering the equivalent treatment of deep models as differentiable programs by the fast-developing parallel field of adversarial learning, we enhance the widely considered strongest adversarial attack

PGD [18] with our saliency technique - drastically improving the efficacy of the attack.

- The ability of precise saliency computation allowed by CAMERAS calls for new quantitative metrics and sanity checks. We contribute two new evaluation metrics and a sanity check to advance this research direction.

2. Related work

The literature has seen multiple techniques that perturb input pixels and measure its effects on the model outputs to identify salient regions in input images [9], [8], [20]. Perturbing every possible combination of pixels has an exponential complexity. Therefore, such techniques often rely on fixed subsets of pixel combinations for tractability. Moreover, high non-linearity of deep models further restricts the saliency map to be reliable under fixed devised perturbation subsets. RISE [20] and Occlusion [33] generate attribution maps by weighing perturbation masks corresponding to the changes in the output scores. Other techniques, such as Meaningful perturbations [9], Extremal perturbations [8], Real-time saliency [6] and [24] cast the problem into an optimization objective. Though effective, these methods face a common issue of allowing a channel for external influence on the resulting maps in the form of e.g. heuristics, external constraints, priors or threshold etc.

Backpropagation saliency methods [27], [25], [33] aim at extracting information from within the model to identify model-centric salient regions in an input image. Relying on layer activations and model gradients, these methods are also known to be computationally more effective [36], [12]. Simoyan et al. [27] first used model gradients as a possible explanation of output predictions. Different adaptations have since been proposed to mitigate the inherent noise sensitivity of model gradients. Guided back prop [30] and DeConvNet [33] alter the backpropagation rules of model ReLU layers, while SmoothGrad [29] computes the average gradients over samples in the close vicinity of the original one. Similarly, DeepLIFT [26], LRP [4] and Excitation Backprop [34] recast the backpropagation rules such that the sum of attribution signal becomes unity. Sundararajan et al. [31] interpolated multiple attribution maps to reduce the signal noise. There are also instances of exploring saliency computation using various layers of deep models by merging the layer activation maps with the gradient information. Such methods include CAM [35], its generalized adaption GradCAM [25], linear approximation [14] and NormGrad [21]. For such methods, it is found that the layers closer to output generate better saliency maps because those layers are more sensitive to the high level class features.

Beyond the visual quality of saliency maps, a few works have also critically explored the reliability of these maps for different methods [19], [13]. Adebayo et al. [2] first introduced sanity checks for image saliency methods, highlighting that visual appeal of the maps alone can be misleading. They evaluated sensitivity of the results of popular techniques to model parameters. Surprisingly, ‘model-centric saliency maps’ computed by multiple methods were found insensitive to the model - failing the sanity check. The techniques avoiding external influences on the map, e.g. Grad-CAM [25] easily passed the test. This finding also resonated with other subsequent sanity checks [21].

Evaluation of image saliency methods is a challenging problem because deep model representation is not always aligned with the human visual system [32]. Hence, indirect evaluation of image saliency is often done by analysing its weak localization performance. For instance, Pointing game score [34] is a commonly used metric for quantitative evaluation of image saliency results [8], [21]. It measures the correlation between the maximal point in a saliency map with the semantic labels of the pixel. Its later adaptation [9] measures the overlap in the bounding boxes from saliency maps and the ground truth. Petsiuk et al. [20] proposed insertion-deletion metrics to measure the impact of perturbations over image patches in the order of importance to quantify saliency accuracy. Nevertheless, these metrics are meant to be weak indicators of saliency maps due to the imprecise nature of the maps computed by the earlier methods, which is no longer the case for CAMERAS.

3. Proposed Approach

Before discussing the details of our technique, we first provide a closer look at the broader paradigm of backpropagation saliency computation. We use Grad-CAM [25] - a popular technique - as a test case to motivate the proposed method. The text below highlights only the relevant aspects of the test case for intuition.

3.1. Saliency computation with backpropagation

Let $\mathbf{I} \in \mathbb{R}^{c \times h \times w}$ be an input image with ‘ c ’ channels. A deep visual classifier $\mathcal{K}(\mathbf{I})$ maps \mathbf{I} to a prediction vector $\mathbf{y}^\ell \in \mathbb{R}^L$, where ‘ L ’ is the total number of classes. Here, ‘ ℓ ’ indicates the predicted label of \mathbf{I} under the premise that the ℓ^{th} coefficient of \mathbf{y} has the largest value. It is well-known that a neural network is a hierarchical composition of representation layers. Rebuffi et al. [21] demonstrated that among these layers, those closer to the input learn class-insensitive features. Thus, the deeper layers hold more promise for computing image saliency for a model. Grad-CAM [25] takes a pragmatic approach to single out the last convolutional layer to estimate the saliency map.

Let us denote the k^{th} activation map of the last convolutional layer of a network as $\mathbf{A}_k(\mathbf{I}) \in \mathbb{R}^{m \times n}$. Focusing on Grad-CAM, the technique first computes an intermediate representation $\mathbf{S}(\mathbf{I}) \in \mathbb{R}^{m \times n}$, such that $\mathbf{S}^{(i,j)}(\mathbf{I}) = \sum_k w_k \mathbf{A}_k^{(i,j)}(\mathbf{I})$, where w_k is given by Eq. (1). Henceforth, we ignore the argument (\mathbf{I}) for clarity, unless required.

$$w_k = \frac{1}{(m+n)} \sum_{i=1}^m \sum_{j=1}^n \left(\frac{\partial \mathbf{y}^\ell}{\partial \mathbf{A}_k^{(i,j)}} \right). \quad (1)$$

In the above expressions, $\mathbf{X}^{(i,j)}$ is the (i,j) coefficient of \mathbf{X} . The computed \mathbf{S} is later extended to the final saliency map $\Psi \in \mathbb{R}^{h \times w}$, as $f(\mathbf{S}) : \mathbf{S} \rightarrow \Psi$, where the function $f(\cdot)$ must account for interpolating an $m \times n$ matrix for a $h \times w$ grid (along other complementary transformations).

3.2. Sub-optimality of backpropagation methods

Observing Grad-CAM (and similar methods) from the above perspective reveals two performance drain-holes in backpropagation image saliency computations.

(a) *Over-simplification of the weights w_k* : Since \mathbf{A}_k is an activation map, individual coefficients of this matrix should have different importance for the final prediction. Indeed, this is also reflected in the values of individual backpropagated gradients for these coefficients - computed with the expression in the parenthesis in Eq. (1). Since the ultimate objective of image saliency is to compute importance of *individual* pixels, losing information with over-simplification of w_k is not conducive. Grad-CAM takes an extreme approach of representing w_k with a scalar value. The main reason for that is, it is actually detrimental to plainly replace w_k with an encoding $\mathbf{W}_k \in \mathbb{R}^{m \times n}$ such that

$S = \sum_k \mathbf{W}_k \odot \mathbf{A}_k$. Here, \odot is the point-wise product and \mathbf{W}_k encodes individual backpropagated gradients. Gradients are extremely sensitive to signal variations. Hence, even a small activation change can result in (a misleading) exaggerated weight alteration for the activation map, resulting in incorrect image saliency. Grad-CAM is able to mitigate this problem by averaging the gradients in Eq. (1). However, this remedy comes at the cost of losing the fine-grained information about the gradients.

Though centered around Grad-CAM, the above discussion points to a simple, yet powerful generic notion for effective backpropagation saliency computation. That is, to better leverage the backpropagated gradients, the differential information of the gradients (in \mathbf{W}_k) is still at our disposal to exploit. Fusing activation maps with this information promises more precise image saliency.

(b) *Large interpolated segments*: Typically, the activation maps of the deeper convolutional layer in visual classifiers are (spatially) much smaller than the input images. For instance, in ResNet [10], the 7×7 maps of the last convolutional layer are 1024 times smaller than the 224×224 inputs. Thus, for $m \ll h$ and $n \ll w$, a saliency map Ψ computed with the class sensitive deeper layers must be mainly composed of interpolated segments. To contextualize, in the above ResNet example, 99.9% of the values are generated by the function $f(\mathbf{S}) : \mathbf{S} \rightarrow \Psi$ in Grad-CAM. This automatically renders Ψ a low-resolution map, leaving alone the issue of correctness of the importance assigned to the individual pixels in the eventual saliency map. The low resolution of saliency maps has also spawned methods to improve $f(\cdot)$ [2], [25]. However, those techniques inevitably rely on external information (including heuristics) for the transform $\mathbf{S} \rightarrow \Psi$ due to unavailability of further useful information from the model itself. This leads to sanity check failures because the operands are no longer purely grounded in the original model.

3.3. The room for improvement

From the above discussion, it is clear that whereas useful techniques exist for backpropagation image saliency, the paradigm is yet to fully harness the backpropagated gradients and resolution enhancement of the activation maps for precise image saliency. Both limitations are rooted in the very nature of the underlying signals. Leveraging these signals from multiple layers can potentially help in partially overcoming the issues. However, this possibility is also restricted by the class-insensitivity of the earlier network layers and combinatorial nature of the problem. Moreover, there is evidence that multi-layer fusion can often adversely affect image saliency [21]. Hence, a technique specifically targeting the class-sensitive last layer, while allowing minimal loss of differential information across backpropagated gradients and improving activation map upsampling, holds significant promises for better saliency computation.

3.4. CAMERAS

Building directly on the insights in § 3.3, we devise CAMERAS - an enhanced resolution and sanity preserving class activation mapping scheme. The approach is illustrated in Fig. (1-top) and explained below in a top-down manner, keeping the flow of the above discussion.

Our method eventually computes a saliency map as:

$$\Psi = f \left(\text{ReLU} \left(\sum_k \mathbf{W}_k^* \odot \mathbf{A}_k^* \right) \right), \forall k, \quad (2)$$

where $\mathbf{W}_k^* \in \mathbb{R}^{h \times w}$ encodes the differential information of the backpropagated gradients for the k^{th} activation map in a network layer, $\mathbf{A}_k^* \in \mathbb{R}^{h \times w}$ is an enhanced resolution encoding for the activation map itself, and $f(\cdot)$ performs an element-wise normalisation in the range $[0, 1]$. The \mathbf{W}_k^* and \mathbf{A}_k^* are defined as follows

$$\begin{aligned} \mathbf{W}_k^* &= \mathbb{E}_t \left[\varphi_t \left(\mathbf{W}_k(\varphi_t(\mathbf{I}, \zeta_t)), \zeta_o \right) \right], \\ \mathbf{A}_k^* &= \mathbb{E}_t \left[\varphi_t \left(\mathbf{A}_k(\varphi_t(\mathbf{I}, \zeta_t)), \zeta_o \right) \right], \end{aligned} \quad (3)$$

where $\varphi_t(\mathbf{X}, \zeta_t)$ is the t^{th} up-sampling applied to resize \mathbf{X} to the dimensions ζ_t - provided as a tuple. We fix $\zeta_o = (h, w)$ for $\mathbf{I} \in \mathbb{R}^{c \times h \times w}$. This will be explained shortly. We compute the (i, j) coefficient of \mathbf{W}_k as $\left(\frac{\partial y^t}{\partial \mathbf{A}_k^{(i,j)}} \right)$. The overall process of generating an image saliency map with CAMERAS is summarized as Algorithm 1.

The algorithm computes the desired saliency map Ψ by an iterative multi-scale accumulation of activation maps and gradients for the κ^{th} layer of the model. In the t^{th} iteration, the input image \mathbf{I} gets up-sampled to ζ_t based on the maximum desired size ζ_m and the number of steps N allowed to reach that size (lines 4,5). Provided that the input up-scaling does not alter the model prediction, the activation maps and backpropagated gradients to the κ^{th} layer are also up-sampled and stored. We show this on lines 6-10 of the algorithm. Notice that, we use calligraphic symbols to distinguish 3D tensors from matrices (e.g. \mathcal{A} instead of \mathbf{A} for activations) in the algorithm for clarity. The newly introduced symbol $\nabla \mathcal{J}(\kappa, \ell)$ on line 9 denotes the collective backpropagated gradients to the κ^{th} layer *w.r.t.* the predicted label ℓ . Also notice, on lines 8, 9, up-sampling of the activation maps and gradients are performed to match the original image size ζ_o . This is because, the same accumulated signals are eventually transformed into the saliency map of the original image. We iteratively accumulate the up-sampled activation maps and gradients, and finally compute their averages on lines 12 and 13. On line 14, we compute the saliency map by solving Eq. (2). Here, matrix notation is intentionally used to match the original equation.

Algorithm 1 CAMERAS algorithm

Input: Classifier \mathcal{K} , image $\mathbf{I} \in \mathbb{R}^{c \times h \times w}$, maximum size ζ_m , steps N , interpolation function $\varphi(\cdot)$, layer κ
Output: Image saliency map $\Psi \in \mathbb{R}^{h \times w}$

- 1: Initialize $\mathcal{A}_o, \mathcal{W}_o$ to $\mathbf{0}$ tensors, $\zeta_o = (h, w)$ and $t = t_m = 0, \ell = \mathcal{K}(\mathbf{I})$
- 2: **while** $t \leq N$ **do**
- 3: $t \leftarrow t + 1$
- 4: $\zeta_t \leftarrow \zeta_{t-1} + \lfloor \frac{\zeta_m}{N} \rfloor (t - 1)$
- 5: $\mathbf{I}_t \leftarrow \varphi_t(\mathbf{I}, \zeta_t)$
- 6: **if** $\mathcal{K}(\mathbf{I}_t) \rightarrow \ell$ **then**
- 7: $t_m \leftarrow t_m + 1$
- 8: $\mathcal{A}_t \leftarrow \mathcal{A}_{t-1} + \varphi_t(\mathcal{A}(\mathbf{I}_t, \kappa), \zeta_o)$
- 9: $\mathcal{W}_t \leftarrow \mathcal{W}_{t-1} + \varphi_t(\nabla \mathcal{J}(\kappa, \ell), \zeta_o)$
- 10: **end if**
- 11: **end while**
- 12: $\mathcal{A} \leftarrow \mathcal{A}_t / t_m$
- 13: $\mathcal{W} \leftarrow \mathcal{W}_t / t_m$
- 14: $\Psi = f\left(\text{ReLU}\left(\sum_k \mathbf{W}_k^* \odot \mathbf{A}_k^*\right)\right), \forall k$
- 15: **return**

In Algorithm 1, CAMERAS is shown to expect four input parameters, along with the classifier and the image. We discuss the choice of $\varphi(\cdot)$ in § 3.4.1 where we eventually propose to keep this function fixed. The algorithm optionally allows κ for computing saliency maps using layers other than the last convolutional layer of the model. For all the experiments presented in the main paper, we keep κ fixed to the last convolutional layer. This is due to the well-known class-sensitivity of the deeper layers of CNNs [21]. Essentially, the only choice to be made is for the values of parameters ζ_m and N , which are related as $\zeta_m = cN\zeta_o$, where c is the step size. Trading-off performance with efficiency, the choice of these parameters is mainly governed by the available computational resources. For larger ζ_m and N values, performance of CAMERAS roughly improves monotonically - generally saturating at $\zeta_m \approx (1K, 1K)$ for $\zeta_o = (224, 224)$ for the popular ImageNet models. For this ζ_m range, the performance is largely insensitive for $N \in [5, 10]$. We give further analysis of parameter values in the supplementary material. In the presented experiments, we empirically choose $\zeta_m = (1K, 1K)$ and $N = 7$.

3.4.1 Sanity preservation and strength

In CAMERAS, we do not impose any prior over the saliency map, nor we use any heuristic to guide its computation process. The technique also preserves model fidelity by requiring no structural (or any other) alteration to the original model. It mainly relies on primitive arithmetic operations over the model signals. These attributes also characterize those other methods that pass the popular sanity checks for backpropagation saliency [2], [21], albeit

resulting in low-resolution saliency maps. In CAMERAS, the only source of any ‘potential’ external influence on the resulting map is through the interpolation function $\varphi(\cdot)$. We conjecture that as long as $\varphi(\cdot)$ is a first-order function defined over the signals originating in the model itself, CAMERAS maps will always preserve their sanity because the maps would fully originate in the model. To preclude any unintentional prior over the maps, our formulation dictates the use of simpler functions as $\varphi(\cdot)$. Hence, we choose to fix bi-linear interpolation as $\varphi(\cdot)$.

To analyse the reasons of extraordinary performance of CAMERAS (see § 4), we provide a brief theoretical perspective on the accumulation of multi-scale interpolated signals exploited by our method, using the results below.

Lemma 3.1: For $\mathbf{X} \in \mathbb{R}^{h \times w}$ and its interpolated approximation $\widetilde{\mathbf{X}} = \varphi(\mathbf{X})$ s.t. $\mathbf{X} \in \mathbb{R}^{m \times n}$ and $\mathbf{X} \neq \widetilde{\mathbf{X}}$, $\|\mathbf{X}^* - \widetilde{\mathbf{X}}\| = f((h-m), (w-n))$ for $m < h$ and $n < w$, where $\varphi(\cdot)$ denotes bi-linear interpolation and $f(\cdot)$ is a monotonic function over its arguments.

Lemma 3.2: For $\widetilde{\mathbf{X}}_z = \varphi(\mathbf{X}_z)$, where $\mathbf{X}_z \in \mathbb{R}^{m \times n}$ and $\widetilde{\mathbf{X}}_p = \varphi(\mathbf{X}_p)$, where $\mathbf{X}_p \in \mathbb{R}^{p \times q}$ s.t. $p < m$ and $n < q$, $\|\mathbf{X}^* - \widetilde{\mathbf{X}}_z\| - \|\mathbf{X}^* - \widetilde{\mathbf{X}}_p\| \leq 0$.

Corollary: $\mathbb{E}_z[\|\mathbf{X}^* - \widetilde{\mathbf{X}}_z\|] \leq \|\mathbf{X}^* - \widetilde{\mathbf{X}}_p\|, \forall \widetilde{\mathbf{X}}_z \cup \widetilde{\mathbf{X}}_p$.

The lemma 3.1 states that bi-linear interpolation tends to be more accurate when the difference in the dimensions of the source signal and the target grid is smaller. The lemma 3.2 can also be easily verified as $\widetilde{\mathbf{X}}_z$ is at least as accurate a projection of \mathbf{X}^* as $\widetilde{\mathbf{X}}_p$, according to lemma 3.1. This necessarily makes its error equal to or smaller than the error of $\widetilde{\mathbf{X}}_p$. In the light of lemma 3.2, the corollary affirms that the expected error of a set of interpolated signals is upper-bounded by the error of the least accurate signal in the set.

The above analysis highlights an important aspect. The CAMERAS results will necessarily be as accurate as operating our algorithm on the original input size, and will improve monotonically thereof with the up-sampled inputs. This is because the activations and gradients of the up-sampled inputs would map more accurately on the original image grid, as per the above results. This is significant because it allows CAMERAS to use the differential information in the backpropagated gradient maps while accounting for the noise-sensitivity of the gradients by averaging out these signals across multiple scales. This is the key strength of the proposed technique.

4. Evaluation

We perform a thorough qualitative and quantitative evaluation of CAMERAS on large-scale models and compare the performance with the state-of-the-art methods.

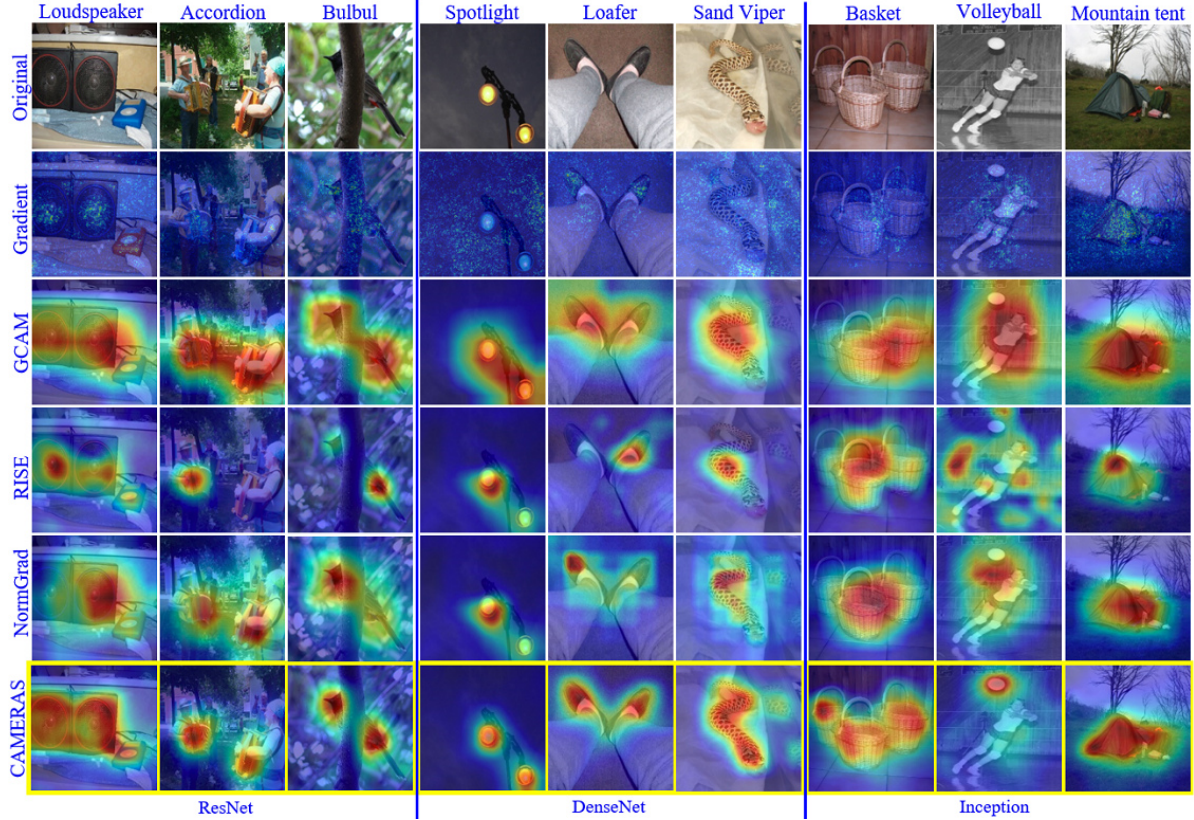


Figure 2. Visual comparison of saliency maps with the state-of-the-art methods. Representative maps are shown for random class images for ResNet-50, DenseNet-121 and Inception-v3. Class labels are provided at the top. Results of proposed CAMERAS are in yellow box.

4.1. Qualitative results

In Fig. 2, we qualitatively compare the results of our technique with the existing saliency methods for ImageNet models of ResNet, DenseNet and Inception. Representative maps of randomly chosen images are provided. See the supplementary material for more visualisations. The high quality of CAMERAS maps is apparent in the figure. A quick inspection reveals that our technique maintains its performance across a variety of scenarios, including clear objects (Loudspeaker, Mountain tent), occluded objects (Bulbul), and multiples instances of objects (Accordion, Basket). Observing carefully, CAMERAS maps provide precise maps even for small and relatively complex geometric shapes, e.g. Volleyball, Spotlight, Loafer, Hognose snake. Interestingly, our method is able to attach appropriate importance even to the reflection of the Hognose snake. Adoption of saliency maps to complex geometric shapes is a direct consequence of enabling precise saliency mapping while sealing-off any external influence on the maps. CAMERAS is able to maintain its characteristic precision across different models and images. These are highly promising results for explainability of modern deep visual classifiers.

4.2. Quantitative results

A quantum leap in performance with CAMERAS is also observed in our quantitative results. Saliency maps are typ-

ically evaluated by measuring their correlation with the semantic annotations of the image. Pointing game [34] is a popular metric for that purpose, which considers the computed saliency for every object class in the image. If the maximal point in the saliency map is contained within the object, it is considered a hit; otherwise, a miss. The performance is measured as the percentage of successful hits. We refer to [34] for more details on the metric. Table 1 benchmarks the performance of CAMERAS for pointing game on 4,952 images of PASCAL VOC test set [7], and $\sim 50K$ images of COCO 2014 validation set [16]. Our technique consistently shows superior performance, achieving up to 27.5% error reduction. The gain is higher for ResNet as compared to VGG due the better performance of the original ResNet that permits better saliency.

The pointing game generally disregards precision of the saliency maps by focusing only on the maximal points. Arguably, crudeness of the saliency maps computed by the earlier methods influenced this evaluation metric. The possibility of precise saliency computation (by CAMERAS) calls for new metrics that account for the finer details of saliency maps. We propose ‘positive map density’ and ‘negative map density’ as two suitable metrics, respectively defined as: $\rho_{\text{map}}^+ = P(\mathcal{K}(\mathbf{I} \odot \Psi)) / \sum_i \sum_j \Psi^{(i,j)} \times (h \times w)$, and $\rho_{\text{map}}^- = P(\mathcal{K}(\mathbf{I} \odot \mathbf{1} - \Psi)) / \sum_i \sum_j (\mathbf{1} - \Psi^{(i,j)}) \times (h \times w)$.

Method	VOC07 Test (All/Diff)		COCO14 Val (All/Diff)	
	VGG16	ResNet50	VGG16	ResNet50
Center [8]	69.6/42.4	69.6/42.4	27.8/19.5	27.8/19.5
Gradient [27]	76.3/56.9	72.3/56.8	37.7/31.4	35.0/29.4
DeConv. [33]	67.5/44.2	68.6/44.7	30.7/23.0	30.0/21.9
Guid [30]	75.9/53.0	77.2/59.4	39.1/31.4	42.1/35.3
MWP [34]	77.1/56.6	84.4/70.8	39.8/32.8	49.6/43.9
cMWP [34]	79.9/66.5	90.7/82.1	49.7/44.3	58.5/53.6
RISE* [20]	86.9/75.1	86.4/78.8	50.8/45.3	54.7/50.0
GradCAM [25]	86.6/74.0	90.4/82.3	54.2/49.0	57.3/52.3
Extremal* [8]	88.0 /76.1	88.9/78.7	51.5/45.9	56.5/51.5
NormGrad [21]	81.9/64.8	84.6/72.2	-	-
CAMERAS	86.2/76.2	94.2/88.8	55.4/50.7	69.9/66.4

Table 1. Mean accuracy on pointing game over the full data (*All*) splits and subset of difficult images (*Diff*), as specified in [34]. The results of other schemes are generated with TorchRay package [8], and *’ indicates an average over 3 runs for improved performance.

Model	Positive map density ($\rho_{map}^+ \uparrow$)			Negative map density ($\rho_{map}^- \downarrow$)		
	NGrad	GCAM	Ours	NGrad	GCAM	Ours
ResNet	1.67	2.33	3.20	0.96	0.86	0.81
DenseNet	1.76	2.35	3.23	1.02	0.94	0.83
Inception	2.19	2.18	3.15	0.95	1.04	0.93

Table 2. The proposed metric scores on ImageNet validation set for the saliency maps of Norm-Grad (NGrad) [21], Grad-CAM (GCAM) [25], and our method.

Here, $P(\cdot)$ is the predicted probability of the actual label of the object. Other notations follow the conventions from above. For an estimated saliency map, the value of ρ_{map}^+ improves if higher importance is attached to a smaller number of pixels that retain higher confidence of the model on the original label. In the extreme case of all the pixels deemed maximally important (saliency value 1), the score depicts the model’s confidence on the object label. On the other hand, the value of ρ_{map}^- decreases if lesser importance is attached by the saliency method to more pixels that do not influence the prediction confidence on the original label. Lower value of this metric is more desired.

Combined ρ_{map}^+ and ρ_{map}^- provide a comprehensive quantification of the quality of the saliency map. We provide results of our technique, Grad-CAM [25] and the recent Norm-Grad [21] on our metrics in Table 2. Due to page limits, we provide further discussion on the proposed metrics in the supplementary material.

5. CAMERAS for analysis

The precise CAMERAS results allow model analysis with backpropagation saliency in unprecedented details. Below, we present a few interesting examples.

The label attribution problem: It is known that deep visual classifiers sometimes learn incorrect association of labels with the objects in input images [9], [8]. For instance, for the image of Chocolate Sauce in Fig. 3, Inception is found to associate the said label to spoon instead of the sauce in the cup [9], [8]. This revelation was possible only through the input perturbation-based methods for attribution due to their precise nature, however, only after fine-tuning a list of parameters for the specific images.

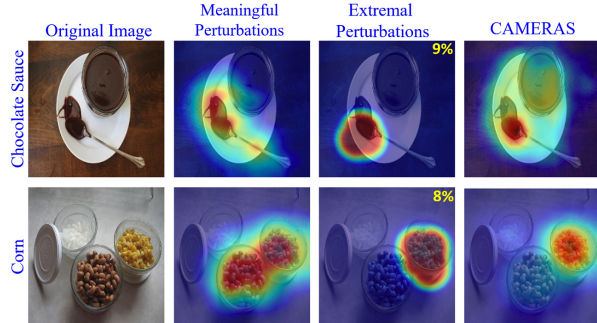


Figure 3. Proposed CAMERAS allows verification of label attribution even more precisely than the input perturbation-based methods. Meaningful [9] and Extremal [8] perturbation methods require image-specific fine tuning of parameters. The latter is confined to 8% and 9% of the image area to achieve the results.

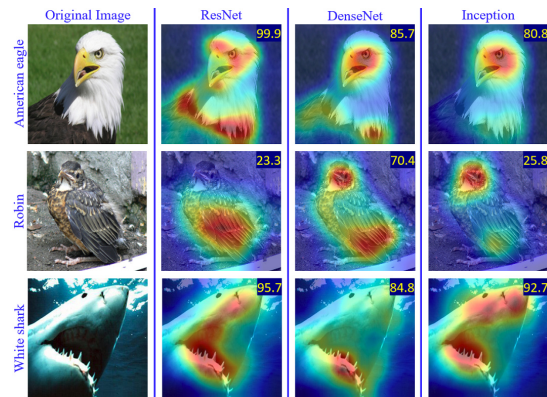


Figure 4. Precise saliency of CAMERAS reveals similarity in the level of attention on fine-grained features causes similarity in the prediction confidence (given as percentages) of different models.

CAMERAS is the first backpropagation saliency method to achieve this result, without requiring any image-specific fine tuning. Our method verifies the original results of the perturbation-based methods with an even better precision.

Prediction confidence: In Fig. 4, CAMERAS results reveal that prediction confidence on individual images is often strongly influenced by a model’s attention on fine-grained features. Different visual models may pay similar attention to the same features to achieve similar confidence scores.

Discrimination of similar objects: Precise saliency mapping of CAMERAS also reveals clear differences of the features learned by the models for visually similar objects. In Fig. 5, we show the saliency maps for multiple examples of ‘Chain-link Fence’ and ‘Swing’ for two high performing models. Notice how the models pay high attention to the individual chain knots (left) as compared to the larger chain structures (right) to distinguish the two classes. These results also reinforce the importance of ‘not enforcing’ any priors on the map (e.g. smoothness [8]). The shown results provide the first instance of clear saliency differences between similar object features under backpropagation saliency mapping without external priors.

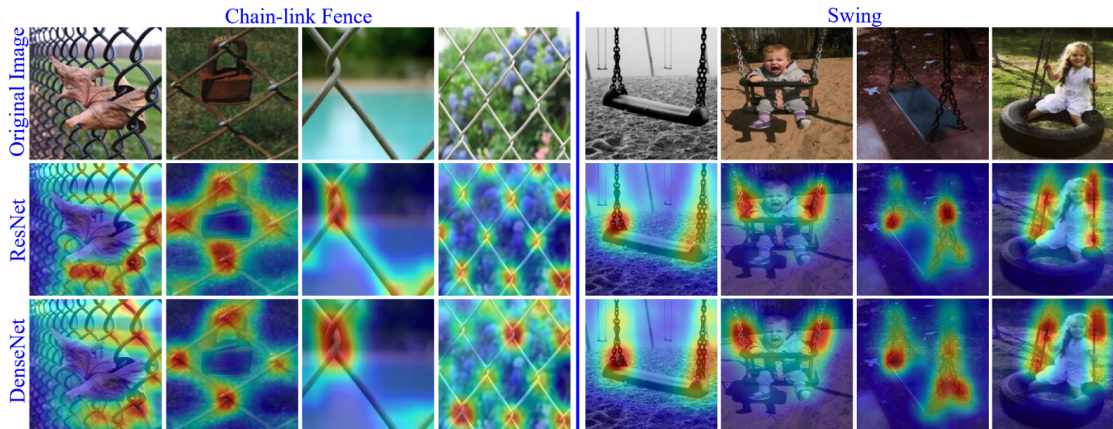


Figure 5. CAMERAS saliency maps reveal the differences of features learned by models to distinguish visually similar objects.

6. Adversarial Attack Enhancement

Similar to the backpropagation saliency methods, most of the adversarial attacks on deep visual classifiers [3] treat the models as differentiable programs. Using model gradients, they engineer additive noise (i.e. perturbations) that alter model predictions on an input. To avoid attack suspicion, the perturbations must be kept norm-bounded. Projected Gradient Descent (PGD) [18] is considered one of the strongest attacks [3] that computes holistic perturbations to fool the models. Using PGD as an example, we show that precision saliency of CAMERAS can significantly enhance these attacks by confining the perturbations to the regions considered more salient by our method, see Fig. 6.

We iteratively solve for the following using the PGD

$$\min_{\mathbf{p}} (\mathcal{J}(\mathbf{I}_{\mathbf{p}}, \ell_{ll}) + \beta \|\mathbf{p} \odot (\mathbf{1} - \Psi)\|_2), \quad (4)$$

where $\mathbf{I}_{\mathbf{p}}$ is the perturbed image, $\mathcal{J}(\cdot)$ is the cross entropy loss, ℓ_{ll} is the least likely label of the clean image, \mathbf{p} is the perturbation, and $\beta = 50$ is an empirically chosen scaling factor. In (4), we allow the perturbation signal to grow freely for our salient regions while restricting it in the other regions. By focusing only on the most important regions, we are able to drastically reduce the required perturbation norm. Maintaining 99.99% fooling confidence for ResNet-50 on all images of ImageNet validation set, we successfully reduced the PGD perturbation norm by 56.5% on average with our CAMERAS enhancement. See the supplementary material for more details. Other adversarial attacks can also be enhanced similarly with CAMERAS.

6.1. Sanity test with adversarial perturbation

Gradient-based adversarial attacks algorithmically compute minimal perturbations to image pixels to maximally change the model prediction. This objective coincides with the objective of image saliency computation, thereby providing a natural sanity check for the saliency methods. That is, the effects of corruption with an adversarial perturbation

to image pixels should correspond to the importance of the pixels identified by the saliency method. Leveraging this fact, we develop a sanity check for saliency methods that operates on pixel-by-pixel basis, which is more suited for precise image saliency. We provide details of the test in the supplementary material due to page limits.




Figure 6. Enhancement of PGD with CAMERAS saliency. Confining the perturbation to high importance regions discovered by CAMERAS drastically reduces visual perceptibility of the attack while maintaining the fooling ratio and confidence on the wrong labels. For clarity, ℓ_{∞} -norm of 12/255 is chosen for perturbation.

7. Conclusion

We introduced CAMERAS to compute precise saliency maps using the gradient backpropagation strategy. Our technique is shown to preserve the sanity of the computed saliency maps by avoiding external influence and priors over the maps. High precision of our saliency maps allow better explanation of deep visual model predictions. We also demonstrated application of CAMERAS to enhance adversarial attacks, and used this to introduce a new sanity check for high-fidelity saliency methods.

Acknowledgment This research was supported by ARC Discovery Grant DP190102443 and partially by DP150100294 and DP150104251. The Titan V used in our experiments was donated by NVIDIA corporation.

References

- [1] Amina Adadi and Mohammed Berrada. Peeking inside the black-box: A survey on explainable artificial intelligence (xai). *IEEE Access*, 6:52138–52160, 2018. [1](#)
- [2] Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, and Been Kim. Sanity checks for saliency maps. In *Advances in Neural Information Processing Systems*, pages 9505–9515, 2018. [1](#), [2](#), [3](#), [4](#), [5](#)
- [3] Naveed Akhtar and Ajmal Mian. Threat of adversarial attacks on deep learning in computer vision: A survey. *IEEE Access*, 6:14410–14430, 2018. [8](#)
- [4] 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(7):e0130140, 2015. [3](#)
- [5] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. [1](#)
- [6] Piotr Dabkowski and Yarín Gal. Real time image saliency for black box classifiers. In *Advances in Neural Information Processing Systems*, pages 6967–6976, 2017. [2](#)
- [7] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results. <http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html>. [6](#)
- [8] Ruth Fong, Mandela Patrick, and Andrea Vedaldi. Understanding deep networks via extremal perturbations and smooth masks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2950–2958, 2019. [1](#), [2](#), [3](#), [7](#)
- [9] Ruth C Fong and Andrea Vedaldi. Interpretable explanations of black boxes by meaningful perturbation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3429–3437, 2017. [1](#), [2](#), [3](#), [7](#)
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [4](#)
- [11] Mohammad AAK Jalwana, Naveed Akhtar, Mohammed Bennamoun, and Ajmal Mian. Attack to explain deep representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9543–9552, 2020. [1](#)
- [12] Andrei Kapishnikov, Tolga Bolukbasi, Fernanda Viégas, and Michael Terry. Xrai: Better attributions through regions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4948–4957, 2019. [3](#)
- [13] Pieter-Jan Kindermans, Sara Hooker, Julius Adebayo, Maximilian Alber, Kristof T Schütt, Sven Dähne, Dumitru Erhan, and Been Kim. The (un) reliability of saliency methods. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, pages 267–280. Springer, 2019. [3](#)
- [14] Pieter-Jan Kindermans, Kristof Schütt, Klaus-Robert Müller, and Sven Dähne. Investigating the influence of noise and distractors on the interpretation of neural networks. *arXiv preprint arXiv:1611.07270*, 2016. [3](#)
- [15] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. [1](#)
- [16] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. [6](#)
- [17] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. [1](#)
- [18] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017. [2](#), [8](#)
- [19] Aravindh Mahendran and Andrea Vedaldi. Salient deconvolutional networks. In *European Conference on Computer Vision*, pages 120–135. Springer, 2016. [3](#)
- [20] Vitali Petsiuk, Abir Das, and Kate Saenko. Rise: Randomized input sampling for explanation of black-box models. *arXiv preprint arXiv:1806.07421*, 2018. [1](#), [2](#), [3](#), [7](#)
- [21] Sylvestre-Alvise Rebuffi, Ruth Fong, Xu Ji, and Andrea Vedaldi. There and back again: Revisiting backpropagation saliency methods. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8839–8848, 2020. [1](#), [3](#), [4](#), [5](#), [7](#)
- [22] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7263–7271, 2017. [1](#)
- [23] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. [1](#)
- [24] Marco Tulio 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*, pages 1135–1144, 2016. [2](#)
- [25] 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 *Proceedings of the IEEE International Conference on Computer Vision*, pages 618–626, 2017. [1](#), [3](#), [4](#), [7](#)
- [26] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. *arXiv preprint arXiv:1704.02685*, 2017. [3](#)
- [27] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. 2014. [1](#), [3](#), [7](#)

- [28] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [1](#)
- [29] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*, 2017. [3](#)
- [30] Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*, 2014. [1](#), [3](#), [7](#)
- [31] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. *arXiv preprint arXiv:1703.01365*, 2017. [3](#)
- [32] Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. *arXiv preprint arXiv:1805.12152*, 2018. [3](#)
- [33] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014. [1](#), [2](#), [3](#), [7](#)
- [34] 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. [1](#), [2](#), [3](#), [6](#), [7](#)
- [35] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016. [3](#)
- [36] Luisa M Zintgraf, Taco S Cohen, Tameem Adel, and Max Welling. Visualizing deep neural network decisions: Prediction difference analysis. *arXiv preprint arXiv:1702.04595*, 2017. [3](#)