

SESS: Saliency Enhancing with Scaling and Sliding

Osman Tursun , Simon Denman , Sridha Sridharan , and Clinton Fookes 

SAIVT Lab, Queensland University of Technology, Australia
{osman.tursun,s.denman,s.sridharan,c.fookes}@qut.edu.au

Abstract. High-quality saliency maps are essential in several machine learning application areas including explainable AI and weakly supervised object detection and segmentation. Many techniques have been developed to generate better saliency using neural networks. However, they are often limited to specific saliency visualisation methods or saliency issues. We propose a novel saliency enhancing approach called **SESS** (Saliency Enhancing with Scaling and Sliding). It is a method and model agnostic extension to existing saliency map generation methods. With SESS, existing saliency approaches become robust to scale variance, multiple occurrences of target objects, presence of distractors and generate less noisy and more discriminative saliency maps. SESS improves saliency by fusing saliency maps extracted from multiple patches at different scales from different areas, and combines these individual maps using a novel fusion scheme that incorporates channel-wise weights and spatial weighted average. To improve efficiency, we introduce a pre-filtering step that can exclude uninformative saliency maps to improve efficiency while still enhancing overall results. We evaluate SESS on object recognition and detection benchmarks where it achieves significant improvement. The code is released publicly to enable researchers to verify performance and further development. Code is available at <https://github.com/neouyghur/SESS>.

1 Introduction

Approaches that generate saliency or importance maps based on the decision of deep neural networks (DNNs) are critical in several machine learning application areas including explainable AI and weakly supervised object detection and semantic segmentation. High-quality saliency maps increase the understanding and interpretability of a DNN’s decision-making process, and can increase the accuracy of segmentation and detection results.

Since the development of DNNs, numerous approaches have been proposed to efficiently produce high-quality saliency maps. However, most methods have limited transferability and versatility. Existing methods are designed for DNN models with specific structures (i.e. a global average pooling layer), for certain types of visualisation (for details refer to Sec. 2), or to address a specific limitation. For instance, CAM [24] requires a network with global average pooling. Guided backpropagation (Guided-BP) [18] is restricted to gradient-based

approaches. Score-CAM [20] seeks to reduce the method’s running-time, while SmoothGrad [17] aims to generate saliency maps with lower noise.

In this work, we propose Saliency Enhancing with Scaling and Sliding (SESS), a model and method agnostic black-box extension to existing saliency visualisation approaches. SESS is only applied to the input and output spaces, and thus does not need to access the internal structure and features of DNNs, and is not sensitive to the design of the base saliency method. It also addresses multiple limitations that plague existing saliency methods. For example, in Fig. 1, SESS shows improvements when applied to three different saliency methods. The saliency map extracted with the gradient-based approach (Guided-BP) is discriminative but noisy. Saliency maps generated by the activation-based method Grad-CAM [14] and perturbation-based method RISE [12] generate smooth saliency maps, but lack detail around the target object and fail to precisely separate the target object from the scene. With SESS, the results of all three methods become less noisy and a more discriminative boundary around the target is obtained.

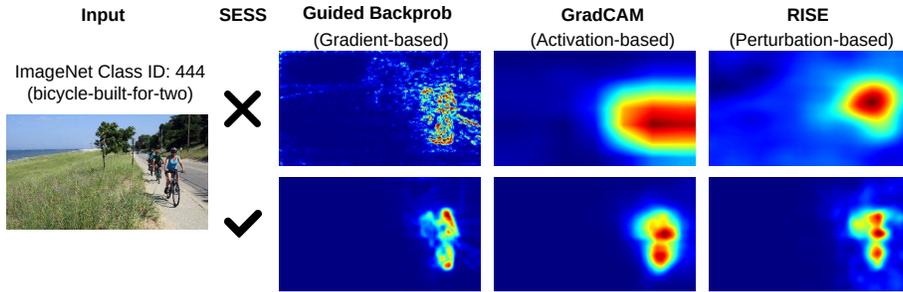


Fig. 1: Example results of three well-known deep neural network visualisation methods with and without SESS. Each of these methods represents one type of saliency map extraction technique. With SESS, all methods generate less noisy and more discriminative saliency maps. The results are extracted with ResNet50, and layer4 is used for Grad-CAM. Target ImageNet class ID is 444 (bicycle-built-for-two).

SESS addresses the following limitations of existing approaches:

- **Weak scale invariance:** Several studies claim that generated saliency maps are inconsistent when there are scale differences [7,21], and we also observe that generated saliency maps are less discriminative when the target objects are comparatively small (See Fig. 1 and Fig. 8).
- **Inability to detect multiple occurrences:** Some deep visualisation methods (i.e., Grad-CAM) fail to capture multiple occurrences of the same object in a scene [1,11]. (See Fig. 8).
- **Impacted by distractors:** Extracted saliency maps frequently incorrectly highlight regions when distractors exist. This is especially true when the

class of the distractor returns a high confidence score, or is correlated with the target class.

- **Noisy results:** Saliency maps extracted with gradient based visualisation approaches [15,19] appear visually noisy as shown in Fig 1.
- **Less discriminative results:** Activation based approaches (e.g., Grad-CAM) tend to be less discriminative, often highlighting large regions around the target such that background regions are often incorrectly captured as being salient.
- **Fixed input size requirements:** Neural networks with fully-connected layers like VGG-16 [16] require a fixed input size. Moreover, models perform better when the input size at inference is the same as the input size during training. As such, most visualisation methods resize the input to a fixed size. This impacts the resolution and aspect ratio, and may cause poor visualisation results [21].

SESS is a remedy for all of the limitations mentioned above. SESS extracts multiple equally sized (i.e., 224×224) patches from different regions of multiple scaled versions of an input image through resizing and sliding window operations. This step ensures that it is robust to scale variance and multiple occurrences. Moreover, since each extracted patch is equal in size to the default input size of the model, SESS takes advantage of high-resolution inputs and respects the aspect ratio of the input image. Each extracted patch will contribute to the final saliency map, and the final saliency map is the fusion of the saliency maps extracted from patches. In the fusion step, SESS considers the confidence score of each patch, which serves to reduce noise and the impact of distractors while increasing SESS’s discriminative power.

The increased performance of SESS is achieved countered a reduction in efficiency due to the use of multiple patches. Quantitative ablation studies show using more scales and denser sliding windows are beneficial, but increase computational costs. To reduce this cost, SESS uses a pre-filtering step that filters out background regions with low target class activation scores. Compared to saliency extraction, the inference step is efficient as it only requires a single forward pass and can exploit parallel computation and batch processing. As such, SESS obtains improved saliency masks with a small increase in run-time requirements. Ablation studies show that the proposed method outperforms its base saliency methods when using pre-filtering with a high pre-filter ratio. In a Pointing Game experiment [22] all methods with SESS achieved significant improvements, despite of a pre-filter ratio of 99% that excludes the majority of extracted patches from saliency generation.

We quantitatively and qualitatively evaluate SESS and conduct ablation studies regarding multiple scales, pre-filtering and fusion. All experimental results show that SESS is a useful and versatile extension to existing saliency methods.

To summarize, the main contributions of this work are as follows:

- We propose, SESS, a model and method agnostic black-box extension to existing saliency methods which is simple and efficient.

- We demonstrate that SESS increases the visual quality of saliency maps, and improves their performance on object recognition and localisation tasks.

2 Related Work

Deep Saliency Methods: Numerous deep neural network-based visualisation methods have been developed in recent years. Based on how the saliency map is extracted, they can be broadly categorised into three groups: gradient-based [15,17,19], class activation-based [24,14,23,20], and perturbation-based [5,12,2] methods.

Gradient-based methods interpret the gradient with respect to the input image as a saliency map. They are efficient as they only require a single forward and backward propagation operation. However, saliency maps generated from raw gradients are visually noisy. Activation-based methods aggregate target class activations of a selected network layer to generate saliency maps. Compared with gradient-based methods, activation-based methods are less noisy, but are also less discriminative and will often incorrectly show strong activations in nearby background regions. Perturbation-based methods generate saliency maps by measuring the changes in the output when the input is perturbed. Perturbation-based methods are slow when compared to most gradient- and activation-based approaches, as they require multiple queries.

Methods can also be split into black-box and white-box according to whether they access the model architecture and parameters. Except for some perturbation-based methods [12,5], saliency methods are all white-box in nature [15,17,19,24,14]. White-box methods are usually more computationally efficient than black-box methods, and require a single forward and backward pass through the network. However, black-box methods are model agnostic, while white-box methods may only work on models with specific architectural features.

Approaches can also be one-shot or multi-shot in nature. One-shot approaches require a single forward and backward pass. Most gradient- and activation-based methods are single-shot. However, multi-shot variants are developed to obtain further improvements. For example, SmoothGrad [17] generates a sharper visualisations through multiple passes across noisy samples of the input image. Integrated Gradients (IG) [19] addresses the “gradient saturation” problem by taking the average of the gradients across multiple interpolated images. Augmented Grad-CAM [10] generates high-resolution saliency maps through multiple low-resolution saliency maps extracted from augmented variants of the input. Smooth Grad-CAM++ [11] utilises the same idea proposed in SmoothGrad to generate sharper saliency maps. To the best of our knowledge, all perturbation methods are multi-shot in nature, as they require multiple queries to the model, each of which has a different perturbation.

Attempts have been made to make multi-shot approaches more efficient. Most such approaches seek to create perturbation masks in an efficient way. Dabkowski *et al.* [2] generates a perturbation mask with a second neural network. Score-CAM [20] uses class activation maps (CAM) as masks; and Group-CAM [23]

follows a similar idea, but further reduces the number of masks through the merging of adjacent maps.

The proposed SESS is a method and model agnostic saliency extension. It can be a “plug-and-play” extension for any saliency methods. However, like perturbation methods, it requires multiple queries. As such, for the sake of efficiency, single-shot and efficient multi-shot approaches are most appropriate for use with SESS.

Enhancing Deep Saliency Visualisations: Many attempts have been made to generate discriminative and low-noise saliency maps. Early Gradient-based methods are visually noisy, and several methods have been proposed to address this. Guided-BP [18] ignores zero gradients during backpropagation by using a RELU as the activation unit. SmoothGrad [18] takes the average gradient of noisy samples [17] to generate cleaner results.

The first of the activation-based methods, CAM, is model sensitive. It requires the model apply a global average pooling over convolutional feature map channels immediately prior to the classification layer [12]. Later variants such as Grad-CAM relax this restriction by using average channel gradients as weights. However, Grad-CAM [14] is also less discriminative, and is unable to locate multiple occurrences of target objects. Grad-CAM++ [1] uses positive partial derivatives of features maps as weights. Smooth Grad-CAM++ [11] combines techniques from both Grad-CAM++ and SmoothGrad to generate sharper visualisations.

Perturbation methods are inefficient, as they send multiple queries to the model. For example, RISE [12] sends 8000 queries to the model to evaluate the importance of regions covered by 8000 randomly selected masks. Recent works reduce the number of masks by using channels in CAMs as masks. For instance, Score-CAM uses all channels in CAMs, while Group-CAM further minimises the number of masks by grouping the channels of CAMs.

All the aforementioned methods have successfully improved certain issues relating to saliency methods, but have limited transferability and versatility. In comparison, SESS is a model and method agnostic extension, which can be applied to any existing saliency approach (though we note that single-pass or efficient multi-pass methods are most suitable). Moreover, SESS is robust to scale-variance, noise, multiple occurrences and distractors. SESS can generate clean and focused saliency maps, and significantly improves the performance of saliency methods for image recognition and detection tasks.

3 Saliency Enhancing with Scaling and Sliding (SESS)

In this section, we introduce SESS. A system diagram is shown in Figure 2, and the main steps are described in Algorithm 1. The implementation of SESS is simple and includes six steps: *multi-scaling*, *sliding window*, *pre-filtering*, *saliency extraction*, *saliency fusion*, and *smoothing*. The first four steps are applied to the input space, and the last two steps are applied at the output space. SESS

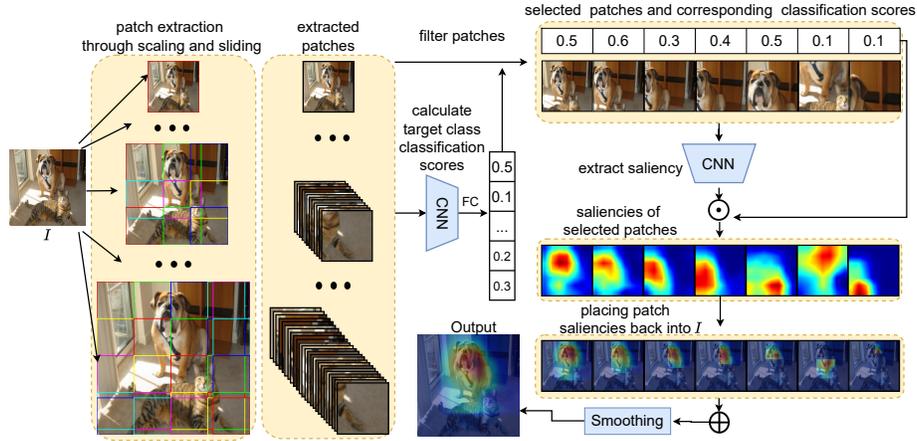


Fig. 2: The SESS Process: SESS includes six major steps: multi-scaling, sliding window, pre-filtering, saliency extraction, saliency fusion and smoothing.

is therefore a black-box extension and a model and method agnostic approach. Each of these steps will be discussed in detail in this section.

Algorithm 1 SESS

Input: Image I , Model f , Target class c , Scale n , Window size (w, h) , Pre-filtering ratio r

Output: Saliency map L_{sess}^c

- 1: $M, P \leftarrow []$
 - 2: **for** $i \in [1, \dots, n]$ **do** ▷ Scaling
 - 3: $M.append(\text{resize}(I, 224 + 64 \times (i - 1)))$
 - 4: **end for**
 - 5: **for** $m \in M$ **do** ▷ Extracting patches
 - 6: $P.append(\text{sliding-window}(m, w, h))$
 - 7: **end for**
 - 8: $B \leftarrow \text{batchify}(P)$
 - 9: $S^c \leftarrow f(B, c)$ ▷ S^c as activation scores of class c
 - 10: $S_{fil}^c, P_{fil} \leftarrow \text{pre-filtering}(S^c, P, r)$ ▷ filter out patches whose class c activation score is lower than top $(100 - r)\%$
 - 11: $A \leftarrow \text{saliency_extraction}(P_{fil}, f, c)$ ▷ get saliency maps of patches after pre-filtering
 - 12: $L \leftarrow \text{calibration}(P_{fil}, A)$ ▷ L is a tensor with shape $n \times w \times h$
 - 13: $L' \leftarrow L \otimes S_{fil}^c$ ▷ Apply channel-wise weight
 - 14: $L_{sess}^c = \text{weighted_average}(L')$ ▷ Apply binary weights to obtain the average of the non-zero values
 - 15: **return** L_{sess}^c
-

Multi-scaling: Generating multiple scaled versions of the input image I is the first step of SESS. In this study, the number of scales, n , ranges from 1 to 12. The set of sizes of all scales is equal to $\{224 + 64 * (i - 1) | i \in \{1, 2, \dots, n\}\}$. The smallest size is equal to the default size of pre-trained models, and the largest size is approximately four times the smallest size. The smaller side of I is resized to the given scale, while respecting the original aspect ratio. M represents the set of all I s at different scales. Benefits of multi-scaling include:

- Most saliency extraction methods are scale-variant. Thus saliency maps generated at different scales are inconsistent. By using multiple scales and combining the saliency results from these, scale-invariance is achievable.
- Small objects will be distinct and visible in salience maps after scaling.

Sliding window: For efficiency the sliding window step occurs after multi-scaling, which calls n resizing operations. A sliding window is applied to each image in M to extract patches. The width w and height h of the sliding window is set to 224. Thus patch sizes are equal to the default input size of pre-trained models in PyTorch¹. The sliding operation starts from the top-left corner of the given image, and slides through from top to bottom and left to right. By default, for efficiency, the step-size of the sliding window is set to 224, in other words there is no overlap between neighbouring windows. However, patches at image boundaries are allowed to overlap with their neighbours to ensure that the entire image is sampled. The minimum number of generated patches is $\sum_{i=1}^n [0.25i + 0.75]^2$. When I has equal width and height and $n = 1$, only one patch of size 224×224 will be extracted, and SESS will return the same results as it’s base saliency visualisation method. Thus, SESS can be viewed as a generalisation of existing saliency extraction methods.

Pre-filtering: To increase the efficiency of SESS, a pre-filtering step is introduced. Generating saliency maps for each extracted patch is computationally expensive. Generally, only a few patches are extracted that contain objects which belong to the target class, and they have comparatively large target class activation scores. Calculating target class activation scores requires only a forward pass, and can be sped-up by exploiting batch operations. After sorting the patches based on activation scores, only patches that have a score in the top $(100 - r)\%$ of patches are selected to generate saliency maps. Here, we denote r the pre-filter ratio. When $r = 0$, no pre-filter is applied. As shown in Fig. 3, when r increases only the region which covers the target object remains, and the number of patches is greatly reduced. For instance, only four patches from an initial set of 303 patches are retained after applying a pre-filter with $r = 99$, and these patches are exclusively focussed on the target object. Of course, a large pre-filter ratio i.e., $r > 50$ will decrease the quality of the generated saliency maps as shown in Fig. 3. Note we use notation S_{fil}^c to represent the class “c” activation scores of the remaining patches after filtering.

Saliency extraction: The saliency maps for the patches retained after pre-filtering are extracted with a base saliency extraction method. Any saliency

¹ <https://pytorch.org>

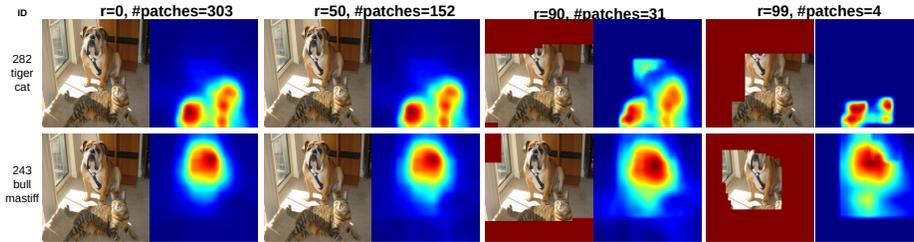


Fig. 3: Visualisation of regions and saliency maps after pre-filtering, when computing saliency maps for the target classes “tiger cat” (top row) and “bull mastiff” (bottom row). All patches that overlap with the red region are removed after pre-filtering.

extraction method is suitable; however real-time saliency extraction methods including Grad-CAM, Guided-BP and Group-CAM are recommended for efficiency. Each extracted saliency map is normalised with Min-Max normalisation. **Saliency fusion:** Since each patch is extracted from a different position or a scaled version of I , a calibration step is applied before fusion. Each saliency map is overlaid on a zero mask image which has the same size as the scaled I from which it was extracted. Then all masks are resized to the same size as I . Here, notation L represents the channel-wise concatenation of all masks. L has n channels of size $w \times h$. Before fusion, a channel-wise weight is applied. S_{fil}^c , the activation scores of patches after filtering, is used as the weight. The weighted L' is then obtained using,

$$L' = L \otimes S_{fil}^c. \quad (1)$$

Finally, a weighted average that excludes non-zero values is applied at each spatial position for fusion. The modified weighted average is used over uniform average to ignore the zero saliency values introduced during the calibration step. Thus, the saliency value at (i, j) of the final saliency map becomes,

$$L_{sess}(i, j) = \frac{\sum_{i=1}^n L'(n, i, j) * \sigma(L'(n, i, j))}{\sum_{i=1}^n \sigma(L'(n, i, j))}, \quad (2)$$

where $\sigma(x) = 1$ if $x > \theta$, else $\sigma(x) = 0$, and $\theta = 0$. A Min-Max normalisation is applied after fusion.

Smoothing: Visual artefacts typically remain between patches after fusion, as shown in Fig. 4. Gaussian filtering is applied to eliminate these artefacts. This paper sets the kernel size to 11 and $\sigma = 5$.

4 Experiments

In this section, we first conduct a series of ablation studies to find the optimal hyper-parameters and show the significance of the steps in our approach. Then,

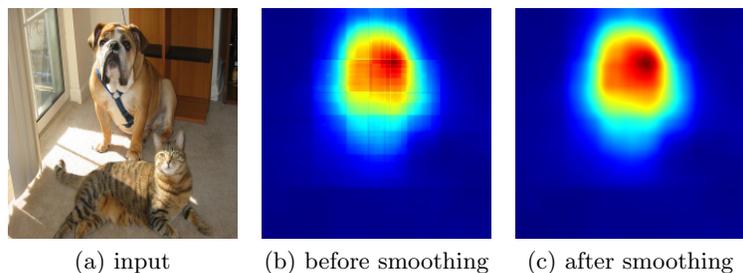


Fig. 4: An example of the effect of the smoothing step. After the smoothing step, edge artefacts are removed and the generated saliency is more visually pleasing.

we qualitatively and quantitatively evaluate the efficiency and effectiveness of SESS compared to other widely used saliency methods.

4.1 Experimental Setup

All experiments are conducted on the validation split of the three publicly available datasets: ImageNet-1k [13], PASCAL VOC07 [3] and MSCOCO2014 [9]. Pre-trained VGG-16 (layer: Feature.29) [16] and ResNet-50 (layer: layer4) [6] networks are used as backbones in our experiments. We used Grad-CAM [14], Guided-BP [18] and Group-CAM [23] as base saliency extraction methods. Grad-CAM and Guided-BP are selected as widely used representations of activation-based and gradient-based approaches. We selected Group-CAM as a representative perturbation-based method given it’s efficiency.

In qualitative experiments, the number of scales and the pre-filter ratio are set to 12 and 0, and smoothing is applied. In quantitative experiments, to reduce computation time and fair comparison, we employ fewer scales and higher pre-filtering ratios, and omit the smoothing step.

4.2 Ablation Studies

We conduct ablation studies on 2000 random images selected from the validation split of ImageNet-1k [13]. ImageNet pre-trained VGG-16 and ResNet-50 networks are used during the ablation study, and Grad-CAM is used as the base saliency method. Insertion and deletion scores [12] are used as evaluation metrics. The intuition behind this metric is that deletion/insertion of pixels with high saliency will cause a sharp drop/increase in the classification score of the target class. The area under the classification score curve (AUC) is used as the quantitative indicator of the insertion/deletion score. A lower deletion score and a higher insertion score indicates a high-quality saliency map. We also reported the overall score as in [23], where the overall score is defined as $AUC(insertion) - AUC(deletion)$. The implementation is the same as [23]. 3.6%

of pixels are gradually deleted from the original image in the deletion test, while 3.6% of pixels are recovered from a highly blurred version of the original image in the insertion test.

Scale: To study the role of multi-scale inputs, we tested different numbers of scales with insertion and deletion tests. As Figure 5 shows, for both VGG-16 and ResNet-50, when the number of scales increases, the insertion scores increase and the deletion scores decrease. Improvements begin to plateau once five scales are used, and converge once ten scales are used. Overall, the improvement is clear even when using images from the ImageNet dataset, where the main object typically covers the majority of the image, and the role of scaling is less apparent.

Pre-filtering ratio: To find a high pre-filtering ratio which increases efficiency whilst retaining high performance, we test 10 different global filters from 0 to 0.9. The insertion score decreases as the pre-filter ratio increases, while the deletion scores fluctuate just slightly until the pre-filter ratio reaches 0.6, after which they increase sharply. This shows pre-filter ratio can be set to 0.5 for both high quality and efficiency. However, we used a pre-filter ratio larger than 0.9 in quantitative experiments.

Channel-wise weights: In the fusion step of SESS, channel-wise weights are applied. Figure 7 qualitatively shows the role of the channel-wise weights. With the channel-wise weights, the extracted saliency maps are more discriminative, better highlighting relevant image regions. Without the channel-wise weights, background regions are more likely to be detected as salient.

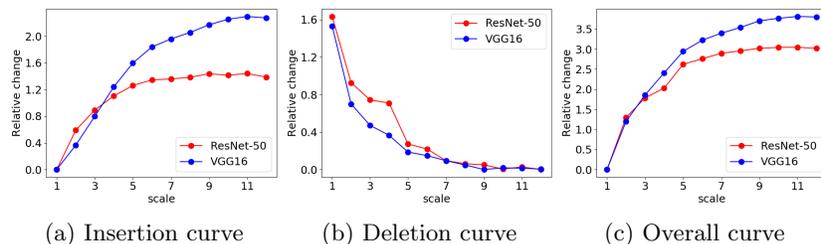


Fig. 5: Ablation study considering scale factor in terms of deletion (lower AUC is better), insertion (higher AUC is better), and overall score (higher AUC is better) on the ImageNet-1k validation split (on a randomly selected set of 2k images).

4.3 Image Recognition Evaluation

Here, we also use the insertion and deletion metrics to evaluate the performance of the proposed SESS. We evaluate SESS with three base saliency extraction approaches (Grad-CAM, Guided-BP and Group-CAM) and two backbones (VGG-16 and ResNet-50) on 5000 randomly selected images from the ImageNet-1k

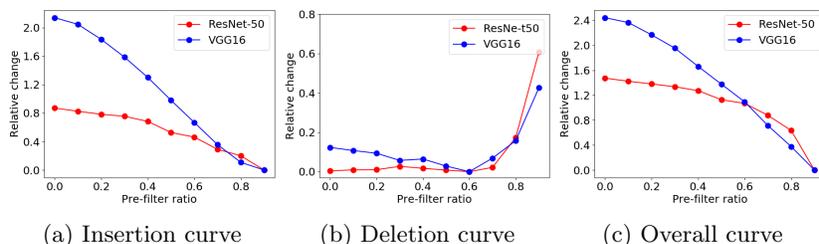


Fig. 6: Ablation study considering the pre-filtering operation in terms of deletion (lower is better), insertion (higher is better), and overall (higher is better) scores on the ImageNet-1k validation split (on a randomly selected set of 2k images).

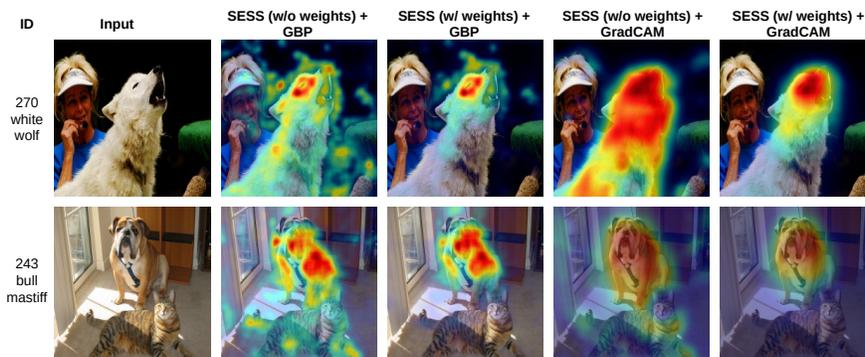


Fig. 7: Impact of the channel-wise weights: The use of the channel-wise weights suppresses activations in background regions, and results in a more focussed saliency map.

validation split. Considering the efficiency, the number of scales and pre-filter ratio are set to 10 and 0.9. With SESS, all three methods with two different backbones achieve improvements, especially Guided-BP whose overall score increases by nearly 5%.

4.4 Qualitative Results

This section qualitatively illustrates how much visual improvement SESS brings over the base visualisation approaches such as Grad-CAM [14] and Guided-BP [18]. As a baseline, we selected five visualisation approaches: Guided-BP, SmoothGrad [17], RISE [12], Score-CAM [20] and Grad-CAM. ResNet-50 is selected as the backbone for all methods. We selected challenging cases for queries, including instances with multiple occurrences of the target classes, the presence of distractors, small targets, and curved shapes.

Method	Model/layer	SESS	Insertion (\uparrow)	Deletion (\downarrow)	Over-all (\uparrow)
Grad-CAM [14]	ResNet-50		68.1	12.1	56.0
		✓	68.6	11.3	57.3
	VGG-16		60.6	9.1	51.5
		✓	60.3	8.1	52.2
Guided-BP [18]	ResNet-50		47.8	11.0	36.8
		✓	53.0	12.0	41.0
	VGG-16		38.8	6.8	32.0
		✓	44.3	6.9	37.4
Group-CAM [23]	ResNet-50		68.2	12.1	56.2
		✓	68.8	11.3	57.4
	VGG-16		61.1	8.8	52.3
		✓	61.1	8.1	53.1

Table 1: Comparison in terms of deletion (lower is better), insertion (higher is better), and the overall (higher is better) scores on a randomly selected set of 5000 images from the ImageNet-1k validation split.

As shown in Fig. 8, visualisation results with SESS are more discriminative and contain less noise. SESS reduces noise and suppresses distractors from the saliency maps, while making the Grad-CAM maps more discriminative and robust to small scales and multiple occurrences.

4.5 Running Time

We calculated the average running time of Grad-CAM, Guided-BP and Group-CAM with/without SESS on a randomly selected set of 5000 images from the ImageNet-1k validation split. Since SESS’s running time is decided by the pre-filter ratio and the number of scales, we calculated SESS’s running time with two scales (6 and 12) and three pre-filtering ratios (0, 50%, 99%). For comparison, we also calculated the average running time of RISE [12], Score-CAM [20] and XRAI [8]. These experiments are conducted with an NVIDIA T4 Tensor Core GPU and four Intel Xeon 6140 CPUs. Results are given in Table 2. With SESS, the average computation time increased, though it is substantially reduced by using a higher pre-filter ratio and a lower number of scales. Compared to perturbation-based methods, activation/gradient-based methods with SESS are efficient. For instance, in the worst case, SESS requires 16.66 seconds which is still over twice as fast as RISE and XRAI, which require more than 38 seconds for a single enquiry.

4.6 Localisation Evaluation

In this section, we evaluate SESS using the Pointing Game introduced in [22]. This allows us to evaluate the performance of the generated saliency maps on weakly-supervised object localisation tasks. The localisation accuracy of each class is calculated using $Acc = \#Hits / (\#Hits + \#Misses)$. $\#Hit$ is increased

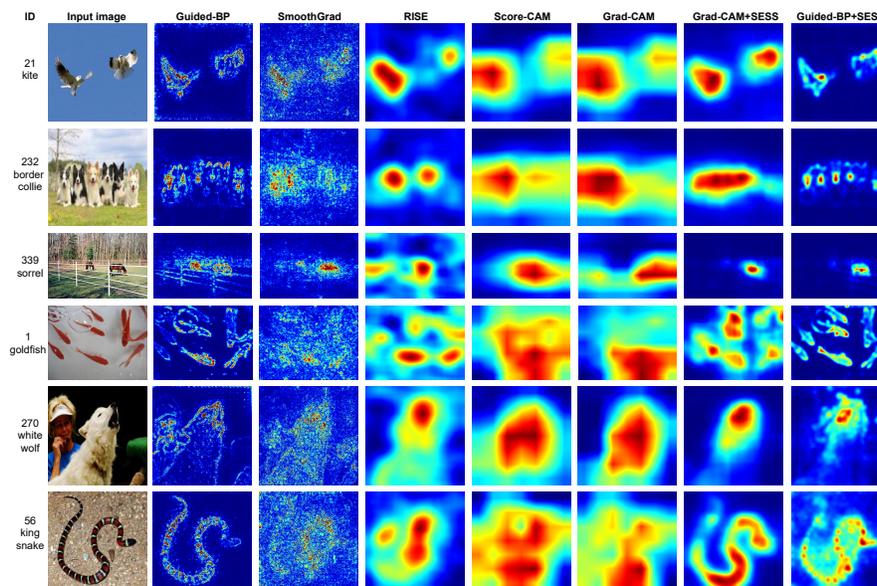


Fig 8: Qualitative comparison with SOTA saliency methods. Target class IDs and inputs images are shown in the first two columns. Later columns show saliency maps for (from left to right) Guided-BP, SmoothGrad, RISE, Score-CAM, Grad-CAM, Grad-CAM with SESS, and Guided-BP with SESS.

Method	Without SESS	With SESS		
		Pre-filter=0% Scale=6/12	Pre-filter=50% Scale=6/12	Pre-filter=99% Scale=6/12
Grad-CAM [14]	0.03	1.23/3.96	0.84/2.36	0.42/0.70
Guided-BP [18]	0.04	1.29/4.09	0.86/2.37	0.41/0.73
Group-CAM [23]	0.13	4.42/16.66	2.54/8.54	0.51/0.92
RISE [12]	38.25	-	-	-
Score-CAM [20]	2.47	-	-	-
XRAI [8]	42.17	-	-	-

Table 2: Comparison of average run-time (in seconds) of a single query from a set of 5000 randomly selected images from the ImageNet-1K validation split.

by one if the highest saliency point is within the ground-truth bounding box of the target class, otherwise $\#Misses$ is incremented. The overall performance is measured by computing the mean accuracy across different categories. Higher accuracy indicates better localisation performance.

We conduct the Pointing Game on the test split of PASCAL VOC07 and the validation split of MSCOCO2014. VGG16 and ResNet50 networks are used as backbones, and are initialised with the pre-trained weights provided by [22].

For implementation, we adopt the TorchRay² library. Grad-CAM, Guided-BP and Group-CAM are chosen as base saliency methods. For efficiency and fair comparison, we set the pre-filter ratio to 99% and the number of scales to 10. As per [22], the results of both the “all” and “difficult” sets are reported. The “difficult” set includes images with small objects (covering less than 1/4 of the image) and distractors.

Results are shown in Table 3. With SESS, all three methods achieved significant improvements, especially on the “difficult” set. The average improvement for all cases is 11.2%, and the average improvement for difficult cases is 19.8%. Grad-CAM with SESS achieved SOTA results. The results further demonstrate that the multi-scaling and sliding window steps of SESS are beneficial when scale variance and distractors exist.

Method	SESS	VOC07 Test (All/Diff)		COCO Val. (All/Diff)	
		VGG16	ResNet50	VGG16	ResNet50
Grad-CAM [14]		86.6/74.0	90.4/82.3	54.2/49.0	57.3/52.3
	✓	90.4/80.8	93.0/86.1	62.0/57.8	67.0/63.2
Guided-BP [18]		75.9/53.0	77.2/59.4	39.1/31.4	42.1/35.3
	✓	79.4/64.2	86.0/75.7	39.5/34.5	44.0/39.4
Group-CAM [23]		80.2/64.9	84.2/71.0	47.4/41.1	48.6/42.4
	✓	89.5/79.8	92.4/85.3	61.2/56.9	66.2/62.3
RISE [12]		86.9/75.1	86.4/78.8	50.8/45.3	54.7/50.0
EBP [22]		77.1/56.6	84.5/70.8	39.8/32.8	49.6/43.9
EP [4]		88.0/76.1	88.9/78.7	51.5/45.9	56.5/51.5

Table 3: Comparative evaluation on the Pointing Game [22].

5 Conclusion

In this work, we proposed SESS, a novel model and method agnostic extension for saliency visualisation methods. As qualitative results show, with SESS, the generated saliency maps are more visually pleasing and discriminative. Improved quantitative experimental results on object recognition and detection tasks demonstrate that SESS is beneficial for weakly supervised object detection and recognition tasks.

References

1. Chattopadhyay, A., Sarkar, A., Howlader, P., Balasubramanian, V.N.: Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In: 2018 IEEE winter conference on applications of computer vision (WACV). pp. 839–847. IEEE (2018)

² <https://facebookresearch.github.io/TorchRay>

2. Dabkowski, P., Gal, Y.: Real time image saliency for black box classifiers. *Advances in neural information processing systems* **30** (2017)
3. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. *International journal of computer vision* **88**(2), 303–338 (2010)
4. Fong, R., Patrick, M., Vedaldi, A.: Understanding deep networks via extremal perturbations and smooth masks. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 2950–2958 (2019)
5. Fong, R.C., Vedaldi, A.: Interpretable explanations of black boxes by meaningful perturbation. In: *Proceedings of the IEEE international conference on computer vision*. pp. 3429–3437 (2017)
6. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
7. Jo, S., Yu, I.J.: Puzzle-cam: Improved localization via matching partial and full features. In: *2021 IEEE International Conference on Image Processing (ICIP)*. pp. 639–643. IEEE (2021)
8. Kapishnikov, A., Bolukbasi, T., Viégas, F., Terry, M.: Xrai: Better attributions through regions. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4948–4957 (2019)
9. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *European conference on computer vision*. pp. 740–755. Springer (2014)
10. Morbidelli, P., Carrera, D., Rossi, B., Fragneto, P., Boracchi, G.: Augmented grad-cam: Heat-maps super resolution through augmentation. In: *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 4067–4071. IEEE (2020)
11. Omeiza, D., Speakman, S., Cintas, C., Weldermariam, K.: Smooth grad-cam++: An enhanced inference level visualization technique for deep convolutional neural network models. *arXiv preprint arXiv:1908.01224* (2019)
12. Petsiuk, V., Das, A., Saenko, K.: Rise: Randomized input sampling for explanation of black-box models. *arXiv preprint arXiv:1806.07421* (2018)
13. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International journal of computer vision* **115**(3), 211–252 (2015)
14. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: *Proceedings of the IEEE international conference on computer vision*. pp. 618–626 (2017)
15. Simonyan, K., Vedaldi, A., Zisserman, A.: Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034* (2013)
16. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014)
17. Smilkov, D., Thorat, N., Kim, B., Viégas, F., Wattenberg, M.: Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825* (2017)
18. Springenberg, J., Dosovitskiy, A., Brox, T., Riedmiller, M.: Striving for simplicity: The all convolutional net. In: *ICLR (workshop track)* (2015)
19. Sundararajan, M., Taly, A., Yan, Q.: Axiomatic attribution for deep networks. In: *International conference on machine learning*. pp. 3319–3328. PMLR (2017)

20. Wang, H., Wang, Z., Du, M., Yang, F., Zhang, Z., Ding, S., Mardziel, P., Hu, X.: Score-cam: Score-weighted visual explanations for convolutional neural networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops. pp. 24–25 (2020)
21. Wang, Y., Zhang, J., Kan, M., Shan, S., Chen, X.: Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12275–12284 (2020)
22. Zhang, J., Bargal, S.A., Lin, Z., Brandt, J., Shen, X., Sclaroff, S.: Top-down neural attention by excitation backprop. *International Journal of Computer Vision* **126**(10), 1084–1102 (2018)
23. Zhang, Q., Rao, L., Yang, Y.: Group-cam: Group score-weighted visual explanations for deep convolutional networks. arXiv preprint arXiv:2103.13859 (2021)
24. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2921–2929 (2016)