

Hallucinating Saliency Maps for Fine-grained Image Classification for Limited Data Domains

Carola Figueroa-Flores^{1,2}, Bogdan Raducanu¹, David Berga¹ and Joost van de Weijer¹

¹*Computer Vision Center, Edifici "O" - Campus UAB, 8193 Bellaterra, Barcelona, Spain*

²*Department of Computer Science and Information Technology, Universidad del Bío Bío, Chile*

Keywords: Fine-grained Image Classification, Saliency Detection, Convolutional Neural Networks.

Abstract: It has been shown that saliency maps can be used to improve the performance of object recognition systems, especially on datasets that have only limited training data. However, a drawback of such an approach is that it requires a pre-trained saliency network. In the current paper, we propose an approach which does not require explicit saliency maps to improve image classification, but they are learned implicitly, during the training of an end-to-end image classification task. We show that our approach obtains similar results as the case when the saliency maps are provided explicitly. We validate our method on several datasets for fine-grained classification tasks (Flowers, Birds and Cars), and show that especially for domains with limited data the proposed method significantly improves the results.

1 INTRODUCTION

Fine-grained image recognition has as objective to recognize many subcategories of a super-category. Examples of well-known fine-grained datasets are Flowers (Nilsback and Zisserman, 2008), Cars (Krause et al., 2013) and Birds (Welinder et al., 2010). The challenge of fine-grained image recognition is that the differences between classes are often very subtle, and only the detection of small highly localized features will correctly lead to the recognition of the specific bird or flower species. An additional challenge of fine-grained image recognition is the difficulty of data collection. The labelling of these datasets requires experts and subcategories can be very rare which further complicates the collection of data. Therefore, the ability to train high-quality image classification systems from few data is an important research topic in fine-grained object recognition.

Most of the state-of-the-art general object classification approaches (Wang et al., 2017; Krizhevsky et al., 2012) have difficulties in the fine-grained recognition task, which is more challenging due to the fact that basic-level categories (e.g. different bird species or flowers) share similar shape and visual appearance. Early works have focused on localization and classification of semantic parts using either explicit annotation (Zhang et al., 2014; Lin et al., 2015; Zhang et al., 2016a; Ding et al., 2019; Du et al., 2020) or weak la-

beling (Zheng et al., 2017; Fu et al., 2017; Wang et al., 2020). The main disadvantage of these approaches was that they required two different 'pipelines', for detection and classification, which makes the joint optimization of the two subsystems more complicated. Therefore, more recent approaches are proposing end-to-end strategies with the focus on improving the feature representation from intermediate layers in a CNN through higher order statistics modeling (Cai et al., 2017; Wang et al., 2018a).

One recent approach which obtained good fine-grained recognition results, especially with only few labeled data is proposed in (Flores et al., 2019). The main idea is that a saliency image can be used to modulate the recognition branch of a fine-grained recognition network. We will refer to this technique as saliency-modulated image classification (SMIC). This is especially beneficial when only few labeled data is available. The gradients which are backpropagated are concentrated on the regions which have high saliency. This prevents backpropagation of gradients of uninformative background parts of the image which could lead to overfitting to irrelevant details. A major drawback of this approach is that it requires an explicit saliency algorithm which needs to be trained on a saliency dataset.

In order to overcome the lack of sufficient data for a given modality, a common strategy is to introduce a 'hallucination' mechanism which emulates the ef-

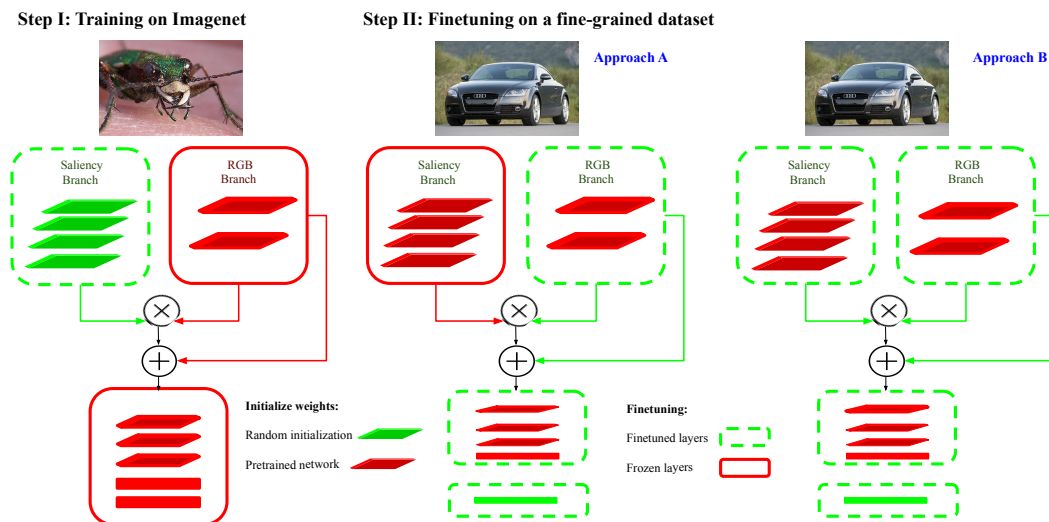


Figure 1: Overview of our method. We process an RGB input image through two branches: one branch extracts the RGB features and the other one is used to learn saliency maps. The resulting features are merged via a modulation layer, which continues with a few more convolutional layers and a classification layer. The network is trained in two steps.

fect of genuine data. For instance, in (Hoffman et al., 2016), they use this ‘hallucination’ strategy for RGB-D object detection. A hallucination network is trained to learn complementary RGB image representation which is taught to mimic convolutional mid-level features from a depth network. At test time, images are processed jointly through the RGB and hallucination networks, demonstrating an improvement in detection performance. This strategy has been adopted also for the case of few-shot learning (Hariharan and Girshick, 2017; Wang et al., 2018b; Zhang et al., 2019). In this case, the hallucination network has been used to produce additional training sample used to train jointly with the original network (also called a meta-learner).

In this paper, we address the major drawback of SMIC, by implementing a hallucination mechanism in order to remove the requirement for providing saliency images for training obtained using one of the existing algorithms (Bylinskii et al.,). In other words, we show that the explicit saliency branch which requires training on a saliency image dataset, can be replaced with a branch which is trained end-to-end for the task of image classification (for which no saliency dataset is required). We replace the saliency image with the input RGB image (see Figure 1). We then pre-train this network for the task of image classification using a subset from ImageNet validation dataset. During this process, the saliency branch will learn to identify which regions are more discriminative. In a second phase, we initialize the weights of the saliency branch with these pre-trained weights. We then train the system end-to-end on the fine-grained dataset using only the RGB images. Results show that

the saliency branch improves fine-grained recognition significantly, especially for domains with few training images.

We briefly summarize below our main contributions:

- we propose an approach which hallucinates saliency maps that are fused together with the RGB modality via a modulation process,
- our method does not require any saliency maps for training (like in these works (Murabito et al., 2018; Flores et al., 2019)) but instead is trained indirectly in an end-to-end fashion by training the network for image classification,
- our method improves classification accuracy on three fine-grained datasets, especially for domains with limited data.

The paper is organized as follows. Section II is devoted to review the related work in fine-grained image classification and saliency estimation. Section III presents our approach. We report our experimental results in Section IV. Finally, Section V contains our conclusions.

2 RELATED WORK

2.1 Fine-grained Image Classification

A first group of approaches on fine-grained recognition operate on a two-stage pipeline: first detecting some object parts and then categorizing the objects using this information.

The work of (Huang et al., 2016) first localizes a set of part keypoints, and then simultaneously processes part and object information to obtain highly descriptive representations. Mask-CNN (Wei et al., 2018) also aggregates descriptors for parts and objects simultaneously, but using pixel-level masks instead of keypoints. The main drawback of these models is the need of human annotation for the semantic parts in terms of keypoints or bounding boxes. To partially alleviate this tedious task of annotation, (Xiao et al., 2015) propose a weakly-supervised approach based on the combination of three types of attention in order to guide the search for object parts in terms of 'what' and 'where'. A further improvement has been reported in (Zhang et al., 2016b), where the authors propose an approach free of any object / part annotation. Their method explores a unified framework based on two steps of deep filter response picking. On the other hand, (Wang et al., 2020) propose an end-to-end discriminative feature-oriented Gaussian Mixture Model (DF-GMM) to learn low-rank feature maps which alleviates the discriminative region diffusion problem in high-level feature maps and thus find better fine-grained details.

A second group of approaches merges these two stages into an end-to-end learning framework which optimizes simultaneously both part localization and fine-grained classification. This is achieved by first finding the corresponding parts and then comparing their appearance (Wang et al., 2018a). In (Xie et al., 2017), their framework first performs unsupervised part candidates discovery and global object discovery which are subsequently fed into a two-stream CNN in order to model jointly both the local and global features. In (Chen et al., 2019), they propose an approach based on 'Destruction and Construction Learning' whose purpose is to force the network to understand ζ_i the semantics of each region. For destruction, a region confusion mechanism (RCM) forces the classification network to learn from discriminative regions. For construction, the region alignment network restores the original region layout by modeling the semantic correlation among regions. A similar idea has been pursued in (Du et al., 2020), where they propose a progressive training strategy to encourage the network to learn features at different granularities (using a random jigsaw patch generator) and afterwards fuse them together. Some other works introduce an attention mechanism. For instance, (Zheng et al., 2017) propose a novel part learning approach by a multi-attention convolutional neural network (MA-CNN) without bounding box/part annotations. MA-CNN jointly learns part proposals (defined as multiple attention areas with strong discrimination ability)

and the feature representations on each part. Similar approaches have been reported in (Sun et al., 2018; Luo et al., 2019). In (Ding et al., 2019), they propose a network which learns sparse attention from class peak responses (which usually corresponds to informative object parts) and implements spatial and semantic sampling. Finally, in (Ji et al., 2020), the authors present an attention convolutional binary neural tree in a weakly-supervised approach. Different root-to-leaf paths in the tree network focus on different discriminative regions using the attention transformer inserted into the convolutional operations along edges of the tree. The final decision is produced as the summation of the predictions from the leaf nodes.

In another direction, some end-to-end frameworks aim to enhance the intermediate representation learning capability of a CNN by encoding higher-order statistics. For instance in (Gao et al., 2016) they capture the second-order information by taking the outer-product over the network output and itself. Other approaches focus on reducing the high feature dimensionality (Kong and Fowlkes, 2017) or extracting higher order information with kernelized modules (Cai et al., 2017). In (Wang et al., 2018a), they learn a bank of convolutional filters that capture class-specific discriminative patches without extra part or bounding box annotations. The advantage of this approach is that the network focuses on classification only and avoids the trade-off between recognition and localization.

Regardless, most fine-grained approaches use the object ground-truth bounding box at test time, achieving a significantly lower performance when this information is not available. Moreover, automatically discovering discriminative parts might require large amounts of training images. Our approach is more general, as it only requires image level annotations at training time and could easily generalize to other recognition tasks.

2.2 Saliency Estimation

Initial efforts in modelling saliency involved multi-scale representations of color, orientation and intensity contrast. These were often biologically inspired such as the well-known work by Itti et al. (Itti et al., 1998). From that model, a myriad of models were based on handcrafting these features in order to obtain an accurate saliency map (Borji and Itti, 2013; Bylinskii et al., 2015), either maximizing (Bruce and Tsotsos, 2005) or learning statistics of natural images (Torralla et al., 2006; Harel et al., 2007). Saliency research was propelled further by the availability of large data sets which enabled the use machine learn-

ing algorithms (Borji, 2018), mainly pretrained on existing human fixation data.

The question of whether saliency is important for object recognition and object tracking has been raised in (Han and Vasconcelos, 2010). Latest methods (Borji, 2018) take advantage of end-to-end convolutional architectures by finetuning over fixation prediction (Kümmerer et al., 2016; ?; ?). But the main goal of these works was to estimate a saliency map, not how saliency could contribute to object recognition. In this paper instead, we propose an approach which does not require explicit saliency maps to improve image classification, but they are learned implicitly, during the training of an end-to-end image classification task. We show that our approach obtains similar results as the case when the saliency maps are provided explicitly.

3 PROPOSED METHOD

Several works have shown that having the saliency map of an image can be helpful for object recognition and fine-grained recognition in particular (Murabito et al., 2018; Flores et al., 2019). The idea is twofold: the saliency map can help focus the attention on the relevant parts of the image to improve the recognition, and it can help guide the training by focusing the backpropagation to the relevant image regions. In (Flores et al., 2019), the authors show that saliency-modulated image classification (SMIC) is especially efficient for training on datasets with few labeled data. The main drawback of these methods is that they require a trained saliency method. Here we show that this restriction can be removed and that we can hallucinate the saliency image from the RGB image. By training the network for image classification on the imageNet dataset we can obtain the saliency branch without human groundtruth images.

3.1 Overview of the Method

The overview of our proposed network architecture is illustrated in Figure 1. Our network consists of two branches: one to extract the features from an RGB image, and the other one (saliency branch) to generate the saliency map from the same RGB image. Both branches are combined using a *modulation layer* (represented by the \times symbol) and are then processed by several shared layers of the joint branch which finally ends up with a classification layer. The RGB branch followed by the joint branch resembles a standard image classification network. The novelty of our architecture is the introduction of the saliency

branch, which transforms the generated saliency image into a *modulation image*. This modulation image is used to modulate the characteristics of the RGB branch, putting more emphasis on those characteristics that are considered important for the fine-grained recognition task. In the following sections we provide the details of the network architecture, the operation of the modulation layer, and finally, how our saliency map is generated. We explain our model using AlexNet (Krizhevsky et al., 2012) as the base classification network, but the theory could be extended to other convolutional neural network architectures. For instance, in the experimental results section, we also consider the ResNet-152 architecture (He et al., 2016).

3.2 Hallucination of Saliency Maps from RGB Images

The function of the visual attention maps is to focus on the location of the characteristics necessary to identify the target classes, ignoring anything else that may be irrelevant to the classification task. Therefore, given an input RGB image, our saliency branch should be able to produce a map of the most salient image locations useful for classification purposes.

To achieve that, we apply a CNN-based saliency detector consisting of four convolutional layers (based on the AlexNet architecture)¹. The output from the last convolutional layer, i.e. one with 384 dimensional feature maps with a spatial resolution of 13×13 (for a 227×227 RGB input image), is further processed using a 1×1 convolution and then a function of activation ReLU. This is to calculate the saliency score for each "pixel" in the feature maps of the previous layer, and to produce a single channel map. Finally, to generate the input for the subsequent classification network, the 13×13 saliency maps are upsampled to 27×27 (which is the default input size of the next classification module) through bilinear interpolation. We justify the size of the output maps by claiming that saliency is a primitive mechanism, used by humans to direct attention to objects of interest, which is evoked by coarse visual stimuli. Therefore, our experiments (see section IV) show that 13×13 feature maps can encode the information needed to detect salient areas and drive a classifier with them.

¹We vary the number of convolutional layers in the experimental section and found four to be optimal.

3.3 Fusion of RGB and Saliency Branches

Consider an input image $I(x, y, z)$, where $z = \{1, 2, 3\}$ indicate the three color channels of the image. Also consider a saliency map $s(x, y)$. In Flores et al. (Flores et al., 2019), a network $h(I, s)$ was trained which performed image classification based on the input image I and the saliency map s . Here, we replace the saliency map (which was generated by a saliency algorithm) by a hallucinated saliency map $h(I, \hat{s}(I))$. The hallucinated saliency map \hat{s} is trained end-to-end and estimated from the same input image I without the need of any ground truth saliency data.

The combination of the hallucinated saliency map \hat{s} , which is the output of the saliency branch, and the RGB branch is done with modulation. Consider the output of the i^{th} layer of the network, l^i , with dimension $w_i \times h_i \times z_i$. Then we define the modulation as

$$\hat{l}^i(x, y, z) = l^i(x, y, z) \cdot \hat{s}(x, y), \quad (1)$$

resulting in the saliency-modulated layer \hat{l}^i . Note that a single hallucinated saliency map is used to modulate all i feature maps of \hat{l} .

In addition to the formula in Eq. (1) we also introduce a skip connection from the RGB branch to the beginning of the joint branch, defined as

$$\hat{l}^i(x, y, z) = l^i(x, y, z) \cdot (\hat{s}(x, y) + 1). \quad (2)$$

This skip connection is depicted in Figure 1 (+ symbol). It prevents the modulation layer from completely ignoring the features from the RGB branch. This is inspired by this work (Flores et al., 2019) that found this approach beneficial when using attention for network compression.

We train our architecture in an end-to-end manner. The backpropagated gradient for the modulation layer into the image classification branch is equal defined as:

$$\frac{\partial L}{\partial l^i} = \frac{\partial L}{\partial \hat{l}^i} \cdot (\hat{s}(x, y) + 1), \quad (3)$$

where L is the loss function of the network. We can see that the saliency map modulates both the forward pass (see Eq. (2)) as well as the backward pass in the same manner; in both cases putting more weight on the features that are on locations with high saliency, and putting less weight on the irrelevant features. We show in the experiments that this helps the network train more efficiently, also on datasets with only few labeled samples. The modulation prevents the network from overfitting to the background.

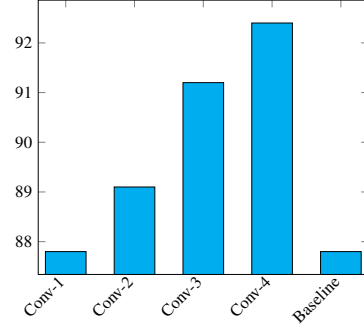


Figure 2: Graph shows the classification accuracy on Flowers for various number of layers in the saliency branch. Best results are obtained with four convolutional layers. Baseline refers to the method without saliency branch.




Figure 3: Graph shows the classification accuracy on Flowers. Various points for fusing the saliency and RGB branch are evaluated. Best results are obtained when fusion is placed before the pool-2 layer. Baseline refers to the method without saliency branch.

3.4 Training on Imagenet and Fine-tuning on a Target Dataset

As can be seen in Figure 1, the training of our approach is divided into two steps: first, training on Imagenet and second, fine-tuning on a target dataset.

Step 1: Training of Saliency Branch on Imagenet.

As explained above, the aim of the saliency branch is to hallucinate (generate) a saliency map directly from an RGB input image. This network is constructed by initializing the RGB branch with pretrained weights from Imagenet. The weights of the saliency branch are initialized randomly using the Xavier method (see Figure 1, left image). The network is then trained selectively, using the ImageNet validation set: we allow to train only the layers corresponding to the saliency branch (depicted by the surrounding dotted line) and freeze all the remaining layers (depicted through the continuous line boxes).

Step 2: Fine-tuning on a Target Dataset. In this step, we initialize the RGB branch with the weights

Table 1: Classification accuracy for Flowers, Cars, and Birds dataset (results are the average over three runs), using **AlexNet** as base network. Results are provided for varying number of training images, from 1 until 30; K refers to using the number of training images used in the official dataset split. The rightmost column shows the average. The * indicates that the method requires an explicit saliency method. Our method (Approach B) obtains similar results as SMIC but without the need of a pretrained saliency network trained on a saliency dataset.

	#train images	1	2	3	5	10	15	20	25	30	K	AVG
Flowers	Baseline-RGB	31.8	45.8	53.1	63.6	72.4	76.9	81.2	85.1	87.2	87.8	68.3
	Baseline-RGB + scratch SAL	34.3	48.9	54.3	65.9	73.1	77.4	82.3	85.9	88.9	89.1	70.0
	SMIC (Flores et al., 2019)*	37.6	51.9	57.1	68.5	75.2	79.7	84.9	88.2	91.2	92.3	72.7
	Approach A	36.9	51.3	56.9	67.8	74.9	78.4	82.9	88.1	90.9	92.0	72.0
	Approach B	37.3	51.7	57.2	68.7	75.6	78.7	83.8	88.4	91.7	92.5	72.6
Cars	Baseline-RGB	4.1	7.8	11.7	17.3	25.5	31.1	38.5	42.2	47.2	60.0	28.5
	Baseline-RGB + scratch SAL	5.9	10.7	14.4	19.1	27.4	32.9	38.5	44.0	48.7	61.5	30.3
	SMIC (Flores et al., 2019)*	9.3	14.0	18.0	22.8	30.0	34.7	40.4	46.0	50.0	61.4	32.7
	Approach A	9.3	14.3	17.4	22.3	28.4	35.3	39.7	45.7	50.1	61.9	32.4
	Approach B	9.8	15.1	18.4	22.9	28.8	35.1	39.9	45.8	49.7	62.9	32.8
Birds	Baseline-RGB	9.1	13.6	19.4	27.7	37.8	44.3	48.0	50.0	54.2	57.0	34.8
	Baseline-RGB + scratch SAL	10.4	14.9	20.3	28.3	38.6	43.9	46.9	48.4	50.7	55.7	35.8
	SMIC (Flores et al., 2019)*	13.1	18.9	22.2	30.2	38.7	44.3	48.0	50.0	54.2	57.0	37.7
	Approach A	11.8	18.3	22.1	29.3	39.1	44.4	47.8	49.7	53.1	56.5	37.2
	Approach B	12.9	18.7	22.7	29.7	39.4	44.1	48.2	49.9	53.9	57.7	37.7

pre-trained from Imagenet and the saliency branch with the corresponding pre-trained weights from Step 1. The weights of the top classification layer are initialized randomly, using the Xavier method. Then, this network is then further fine-tuned on a target dataset, selectively. We distinguish two cases:

- Approach A. We freeze the layers of the saliency branch and we allow all the other layers in the network to be trained. This process is depicted by the continuous line surrounding the saliency branch and the dotted line for the rest (see the Figure 1, middle image).
- Approach B. We allow all layers to be trained. Since we consider training on datasets with only few labels this could result in overfitting, since it requires all the weights of the saliency branch to be learned (see the Figure 1, right image).

In the experiments we evaluate both approaches to training the network.

4 EXPERIMENTS

4.1 Experimental Setup

Datasets. To evaluate our approach, we used three standard datasets used for fine-grained image classification:

- *Flowers*: Oxford Flower 102 dataset (Nilsback and Zisserman, 2008) has 8,189 images divided in 102 classes.

- *Birds*: CUB200 has 11,788 images of 200 different bird species (Welinder et al., 2010).

- *Cars*: the CARS-196 dataset in (Krause et al., 2013) contains 16,185 images of 196 car classes.

Network Architectures. We evaluate our approach using two network architectures: Alexnet (Krizhevsky et al., 2012) and Resnet-152 (He et al., 2016). In both cases, the weights were pretrained on Imagenet and then finetuned on each of the datasets mentioned above. The networks were trained for 70 epochs with a learning rate of 0.0001 and a weight decay of 0.005. The top classification layer was initialized from scratch using Xavier method (Glorot and Bengio, 2010).

Evaluation Protocol. To validate our approach, we follow the same protocol as in (Flores et al., 2019). For the image classification task, we train each model with subsets of k training images for $k \in \{1, 2, 3, 5, 10, 15, 20, 25, 30, K\}$, where k is the total number of training images for the class. We keep 5 images per class for validation and 5 images per class for test. We report the performance in terms of accuracy, i.e. percentage of correctly classified samples. We show the results as an average over three runs.

4.2 Fine-grained Image Classification Results

Optimal Depth and Fusion Saliency Branch. First we evaluate the saliency branch with a varying number of convolutional layers. The results are presented in Figure 2. We found that four convolutional layers

Table 2: Classification accuracy for Flowers, Cars, and Birds dataset (results are the average over three runs), using **ResNet152** as base network. Results are provided for varying number of training images, from 1 until 30; K refers to using the number of training images used in the official dataset split. The rightmost column shows the average. The * indicates that the method requires an explicit saliency method. Our method (Approach B) obtains similar results as SMIC but without the need of a pretrained saliency network trained on a saliency dataset.

	#train images	1	2	3	5	10	15	20	25	30	K	AVG
Flowers	Baseline-RGB	39.0	60.1	68.0	82.5	89.0	92.0	92.1	93.3	94.2	95.4	80.3
	Baseline-RGB + scratch SAL	40.1	63.8	69.7	83.9	89.7	91.9	92.9	93.8	95.1	97.1	81.8
	SMIC (Flores et al., 2019)*	42.6	64.2	70.9	85.5	90.9	92.7	94.0	95.0	97.0	97.8	83.1
	Approach A	42.4	64.5	70.7	85.2	90.3	92.4	93.3	94.3	96.5	97.9	82.8
	Approach B	42.7	64.5	71.0	85.1	90.4	92.5	93.1	94.7	96.8	98.1	82.9
Cars	Baseline-RGB	30.9	45.8	53.1	62.7	70.9	73.9	79.9	88.7	89.2	90.7	68.6
	Baseline-RGB + scratch SAL	33.8	46.1	54.8	63.8	71.7	74.9	80.9	88.1	89.1	91.0	69.4
	SMIC (Flores et al., 2019)*	34.7	47.9	55.2	64.9	72.1	75.8	82.1	90.0	91.1	92.4	70.6
	Approach A	34.1	47.0	56.3	64.7	71.9	75.3	81.7	89.0	90.8	91.7	70.2
	Approach B	34.0	47.5	55.4	64.7	71.8	75.5	81.9	89.3	91.0	92.1	70.3
Birds	Baseline-RGB	24.9	35.3	44.1	53.3	63.8	71.8	75.7	79.3	82.9	83.7	61.5
	Baseline-RGB + scratch SAL	26.3	36.1	45.2	53.9	64.3	72.1	76.3	79.9	83.1	83.4	62.1
	SMIC (Flores et al., 2019)*	28.1	37.9	46.8	55.2	65.3	73.1	77.0	82.9	84.4	86.1	63.7
	Approach A	26.9	36.9	46.1	54.2	64.9	72.8	77.1	81.4	83.4	84.8	62.9
	Approach B	27.1	37.0	46.2	54.9	65.4	72.8	77.1	81.3	83.8	85.1	63.1

led to a significant increase in performance. In addition, we look at the best RGB branch layer to perform the fusion of the saliency branch and the RGB branch. The results are presented in Figure 3. It is found to be optimal to fuse the two branches before the Pool-2 layer for AlexNet². Based on these experiments, we use four convolutional layers in the saliency branch and fuse before the second pool layer for the remainder of the experiments and for all datasets.

Evaluation on Scarce Data Domain. As described in section III, we consider two alternative ways to train the saliency branch on the target dataset: keeping the saliency branch fixed (Approach A) or allowing it to finetune (Approach B). In this section, we compare these two approaches with respect to the Baseline-RGB and Baseline-RGB + scratch SAL (where Saliency branch is initialized from scratch without pretraining on Imagenet). In addition, we compare to the SMIC method of Flores et al. (Flores et al., 2019) who also reports results for small training datasets. We do not compare to other fine-grained methods here, because they do not report results when only considering few labeled images. The experiments are performed on *Flowers*, *Cars* and *Birds* datasets and can be seen in Table 1. The average improvement of accuracy of our *Approach A* and *B* with respect the *Baseline-RGB* is 3.7% and 4.3%, respectively for the *Flowers* dataset; 3.9% and 4.3%, respectively for the *Cars* dataset; and 2.4% and 2.9%, respectively for the *Birds* dataset. Our *Approach B* is especially advantageous when we compare it with

²In a similar study, we found that for Resnet-152 the optimal fusion is after the forth residual block.

the SMIC approach, where an additional algorithm is needed to generate the saliency map. It is therefore advantageous to also finetune the saliency branch on the target data even when we only have a few labeled images per class.

In Table 2, we show the same results but now for ResNet152. One can see that the results improve significantly, especially for *Cars* results improve a lot. The same general conclusions can be drawn : *Approach B* obtains better results than *Approach A* and the method obtains similar results as SMIC but without the need of a pretrained saliency network.

Qualitative Results. Table 3 shows some qualitative results for the case when the pretrained version of our approach predicts the correct label, meanwhile the Baseline-RGB fails. Alternatively, in Table 4 depicts the opposite case: the Baseline-RGB predicts the correct label of the test images, meanwhile the pretrained version of our approach fails. In both cases, the saliency images have been generated using our Approach B. A possible explanation for the failures in this latter case could be that the saliency images are not able to capture the relevant region of the image for fine-grained discrimination. Thus, the saliency-modulated layer focuses on the wrong features for the task.

5 CONCLUSIONS

In this work, we proposed a method to improve fine-grained image classification by means of saliency maps. Our method does not require explicit saliency

Table 3: Some success examples on Flowers: when the prediction done by Baseline-RGB fails to infer the right label for some test images, but the prediction by our approach is correct. Example image contains image of the wrongly predicted class.

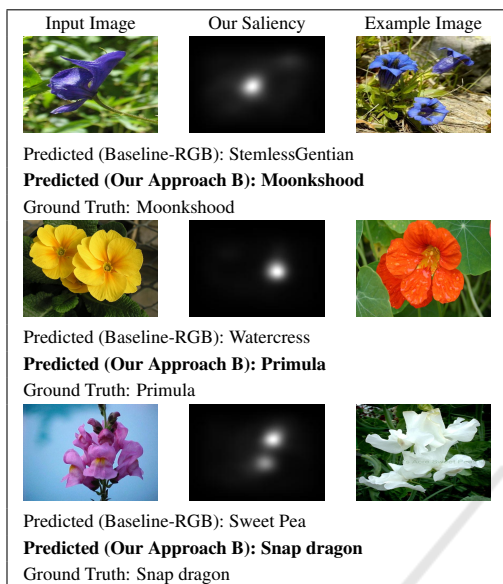
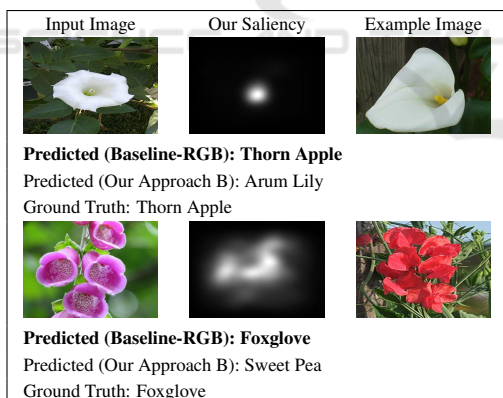


Table 4: Some failure examples on Flowers: when the prediction done by our method fails to infer the right label for some test images, but the prediction by Baseline-RGB is correct. Example image contains image of the wrongly predicted class.



maps, but they are learned implicitly during the training of an end-to-end deep convolutional network. We validated our method on several datasets for fine-grained classification tasks (Flowers, Birds and Cars). We showed that our approach obtains similar results as the SMIC method (Flores et al., 2019) which required explicit saliency maps. We showed that combining RGB data with saliency maps represents a significant advantage for object recognition, especially for the case when training data is limited.

ACKNOWLEDGEMENTS

The authors acknowledge the Spanish project PID2019-104174GB-I00 (MINECO). and the CERCA Programme of Generalitat de Catalunya. Carola Figueroa is supported by a Ph.D. scholarship from CONICYT (now ANID), Chile.

REFERENCES

- Borji, A. (2018). Saliency prediction in the deep learning era: Successes, limitations, and future challenges. *arXiv*, arXiv:1810.03716.
- Borji, A. and Itti, L. (2013). State-of-the-art in visual attention modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(1):185–207.
- Bruce, N. D. B. and Tsotsos, J. K. (2005). Saliency based on information maximization. In *Proceedings of the 18th International Conference on Neural Information Processing Systems, NIPS'05*, pages 155–162, Cambridge, MA, USA. MIT Press.
- Bylinskii, Z., DeGennaro, E., Rajalingham, R., Ruda, H., Zhang, J., and Tsotsos, J. (2015). Towards the quantitative evaluation of visual attention models. *Vision Research*, 116:258–268.
- Bylinskii, Z., Judd, T., Borji, A., Itti, L., Durand, F., Oliva, A., and Torralba, A. Mit saliency benchmark. <http://saliency.mit.edu/>.
- Cai, S., Zuo, W., and Zhang, L. (2017). Higher-order integration of hierarchical convolutional activations for fine-grained visual categorization. In *Proc. of ICCV*, pages 511–520.
- Chen, Y., Bai, Y., Zhang, W., and Mei, T. (2019). Destruction and construction learning for fine-grained image recognition. In *Proc. of CVPR*, pages 5157–5166.
- Ding, Y., Zhou, Y., Zhu, Y., Ye, Q., and Jiao, J. (2019). Selective sparse sampling for fine-grained image recognition. In *Proc. of ICCV*, pages 6599–6608.
- Du, R., Chang, D., Bhunia, A., Xie, J., Ma, Z., Song, Y.-Z., and Guo, J. (2020). In *Proc. of ECCV*, pages 1–16.
- Flores, C. F., Gonzalez-Garcia, A., van de Weijer, J., and Raducanu, B. (2019). Saliency for fine-grained object recognition in domains with scarce training data. *Pattern Recognition*, 94:62–73.
- Fu, J., Zheng, H., and Mei, T. (2017). Look closer to see better: Recurrent attention convolutional neural network for fine-grained image recognition. In *Proc. of CVPR*, pages 4438–4446.
- Gao, Y., Beijbom, O., Zhang, N., and Darrell, T. (2016). Compact bilinear pooling. In *Proc. of CVPR*, pages 317–326.
- Glorot, X. and Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. In *International Conference on Artificial Intelligence and Statistics*, pages 249–256.
- Han, S. and Vasconcelos, N. (2010). Biologically plausible saliency mechanisms improve feedforward object recognition. *Vision Research*, 50:2295—2307.

- Harel, J., Koch, C., and Perona, P. (2007). Graph-based visual saliency. In Schölkopf, B., Platt, J. C., and Hoffman, T., editors, *Advances in Neural Information Processing Systems 19*, pages 545–552. MIT Press.
- Hariharan, B. and Girshick, R. (2017). Low-shot visual recognition by shrinking and hallucinating features. In *Proc. of ICCV*, pages 3018–3027.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778.
- Hoffman, J., Gupta, S., and Darrell, T. (2016). Learning with side information through modality hallucination. In *Proc. of CVPR*, pages 826–834.
- Huang, S., Xu, Z., Tao, D., and Zhang, Y. (2016). Part-stacked cnn for fine-grained visual categorization. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1173–1182.
- Itti, L., Koch, C., and Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11):1254–1259.
- Ji, R., Wen, L., Zhang, L., Du, D., Wu, Y., Zhao, C., Liu, X., and Huang, F. (2020). Attention convolutional binary neural tree for fine-grained visual categorization. In *Proc. of CVPR*, pages 10468–10477.
- Kümmerer, M., Wallis, T. S. A., and Bethge, M. (2016). Deepgaze ii: Reading fixations from deep features trained on object recognition. *arXiv preprint arXiv:1610.01563*.
- Kong, S. and Fowlkes, C. (2017). Low-rank bilinear pooling for fine-grained classification. In *Proc. of CVPR*, pages 365–374.
- Krause, J., Stark, M., Deng, J., and Fei-Fei, L. (2013). 3d object representations for fine-grained categorization. In *4th IEEE Workshop on 3D Representation and Recognition, at ICCV*, pages 1–8.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 1097–1105.
- Lin, D., Shen, X., Lu, C., and Jia, J. (2015). Deep lac: deep localization, alignment and classification for fine-grained recognition. In *Proc. of CVPR*, pages 1666–1774.
- Luo, W., Yang, X., Mo, X., Lu, Y., Davis, L., Li, J., Yang, J., and Lim, S.-N. (2019). Cross-x learning for fine-grained visual categorization. In *Proc. of ICCV*.
- Murabito, F., Spampinato, C., Palazzo, S., Pogorelov, K., and Riegler, M. (2018). Top-down saliency detection driven by visual classification. *Computer Vision and Image Understanding*, 172:67–76.
- Nilsback, M.-E. and Zisserman, A. (2008). Automated flower classification over a large number of classes. In *Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729.
- Sun, M., Yuan, Y., Zhou, F., and Ding, E. (2018). Multi-attention multi-class constraint for fine-grained image recognition. In *Proc. of ECCV*, pages 834–850.
- Torralba, A., Oliva, A., Castelhano, M. S., and Henderson, J. M. (2006). Contextual guidance of eye movements and attention in real-world scenes: The role of global features in object search. *Psychological Review*, 113(4):766–786.
- Wang, F., Jiang, M., Qian, C., Yang, S., Li, C., Zhang, H., Wang, X., and Tang, X. (2017). Residual attention network for image classification. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3156–3164.
- Wang, Y., Morariu, V. I., and Davis, L. S. (2018a). Learning a discriminative filter bank within a cnn for fine-grained recognition. In *Proc. of CVPR*, pages 4148–4157.
- Wang, Y.-X., Girshick, R., Hebert, M., and Hariharan, B. (2018b). Low-shot learning from imaginary data. In *Proc. of CVPR*, pages 7278–7286.
- Wang, Z., Wang, S., Yang, S., Li, H., Li, J., and Li, Z. (2020). Weakly supervised fine-grained image classification via gaussian mixture model oriented discriminative learning. In *Proc. of CVPR*, pages 9749–9758.
- Wei, X.-S., Xie, C.-W., Wu, J., and Shen, C. (2018). Mask-cnn: Localizing parts and selecting descriptors for fine-grained bird species categorization. *Pattern Recognition*, 76:704 – 714.
- Welinder, P., Branson, S., Mita, T., Wah, C., Schroff, F., Belongie, S., and Perona, P. (2010). Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology.
- Xiao, T., Xu, Y., Yang, K., Zhang, J., Peng, Y., and Zhang, Z. (2015). The application of two-level attention models in deep convolutional neural network for fine-grained image classification. In *Proc. of CVPR*, pages 842–850.
- Xie, G.-S., Zhang, X.-Y., Yang, W., Xu, M., Yan, S., and Liu, C.-L. (2017). Lg-cnn: From local parts to global discrimination for fine-grained recognition. *Pattern Recognition*, 71:118–131.
- Zhang, H., Xu, T., Elhoseiny, M., Huang, X., Zhang, S., Elgammal, A., and Metaxas, D. (2016a). Spda-cnn: Unifying semantic part detection and abstraction for fine-grained recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1143–1152.
- Zhang, H., Zhang, J., and Koniusz, P. (2019). Few-shot learning via saliency-guided hallucination of samples. In *Proc. of CVPR*, pages 2770–2779.
- Zhang, N., Donahue, J., Girshick, R., and Darrell, T. (2014). Part-based r-cnns for fine-grained category detection. In *European Conference on Computer Vision*, pages 834–849.
- Zhang, X., Xiong, H., Zhou, W., Lin, W., and Tian, Q. (2016b). Picking deep filter responses for fine-grained image recognition. In *Proc. of CVPR*, pages 1134–1142.
- Zheng, H., Fu, J., Mei, T., and Luo, J. (2017). Learning multi-attention convolutional neural network for fine-grained image recognition. In *Proc. of ICCV*, pages 5209–5217.