

Attend and Rectify: a Gated Attention Mechanism for Fine-Grained Recovery.

Pau Rodríguez¹, Josep M. Gonfaus², Guillem Cucurull¹, F. Xavier Roca¹, and Jordi Gonzàlez¹

¹ Computer Vision Center and Universitat Autònoma de Barcelona (UAB),
Campus UAB, 08193 Bellaterra, Catalonia Spain.

pau.rodriguez@cvc.uab.cat

² Visual Tagging Services, Parc de Recerca, Campus UAB

Abstract. We propose a novel attention mechanism to enhance Convolutional Neural Networks for fine-grained recognition. It learns to attend to lower-level feature activations without requiring part annotations and uses these activations to update and rectify the output likelihood distribution. In contrast to other approaches, the proposed mechanism is modular, architecture-independent and efficient both in terms of parameters and computation required. Experiments show that networks augmented with our approach systematically improve their classification accuracy and become more robust to clutter. As a result, Wide Residual Networks augmented with our proposal surpasses the state of the art classification accuracies in CIFAR-10, the Adience gender recognition task, Stanford dogs, and UEC Food-100.

Keywords: Deep Learning · Convolutional Neural Networks · Attention

1 Introduction

Humans and animals process vast amounts of information with limited computational resources thanks to attention mechanisms which allow them to focus resources on the most informative chunks of information [1,3,29]

This work is inspired by the advantages of visual and biological attention mechanisms, for tackling fine-grained visual recognition with Convolutional Neural Networks (CNN) [17]. This is a particularly difficult task since it involves looking for details in large amounts of data (images) while remaining robust to deformation and clutter. In this sense, different attention mechanisms for fine-grained recognition exist in the literature: (i) iterative methods that process images using "glimpses" with recurrent neural networks (RNN) or long short-term memory (LSTM) [26,38], (ii) feed-forward attention mechanisms that augment vanilla CNNs, such as the Spatial Transformer Networks (STN) [11], or top-down feed-forward attention mechanisms (FAM) [23]. Although it is not applied to fine-grained recognition, the Residual Attention introduced by [32] is another example of feed-forward attention mechanism that takes advantage of residual

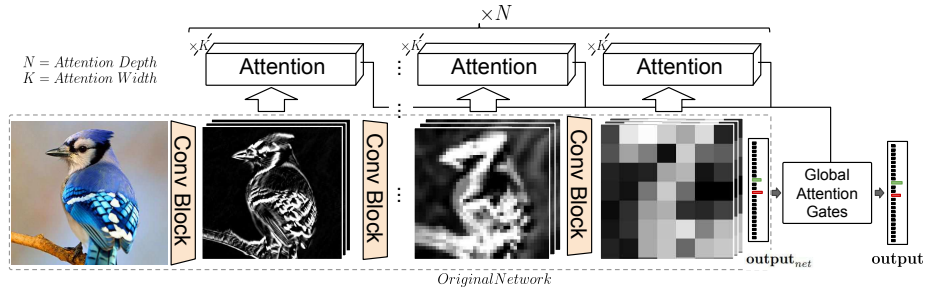


Fig. 1: The proposed mechanism. The original CNN is augmented with N attention modules at N different depths. Each attention module applies K attention heads to the network feature maps to make a class prediction based on local information. The original network output_{net} is then corrected based on the local features by means of the global attention gates, resulting in the final output .

connections [8] to enhance or dampen certain regions of the feature maps in an incremental manner.

Thus, most of the existing attention mechanisms are either limited by having to perform multiple passes through the data [26], by carefully designed architectures that should be trained from scratch [11], or by considerably increasing the needed amount of memory and computation, thus introducing computational bottlenecks [12]. Hence, there is still the need of models with the following learning properties: (i) Detect and process in detail the most informative parts of an image for learning models more robust to deformation and clutter [21]; (ii) feed-forward trainable with SGD for achieving faster inference than iterative models [26,38], together with faster convergence rate than Reinforcement Learning-based (RL) methods [26,19]; (iii) preserve low-level detail for a direct access to local low-level features before they are modified by residual identity mappings. This is important for fine-grained recognition, where low-level patterns such as textures can help to distinguish two similar classes. This is not fulfilled by Residual Attention, where low-level features are subject to noise after traversing multiple residual connections [32].

In addition, desirable properties for attention mechanisms applied to CNNs would be: (i) **Modular and incremental**, since the same structure can be applied at each layer on any convolutional architecture, and it is easy to adapt to the task at hand; (ii) **Architecture independent**, that is, being able to adapt any pre-trained architecture such as VGG [27] or ResNet [8]; (iii) **Low computational impact** implying that it does not result in a significant increase in memory and computation; and (iv) **Simple** in the sense that it can be implemented in few lines of code, making it appealing to be used in future work.

Based on all these properties, we propose a novel attention mechanism that learns to attend low-level features from a standard CNN architecture through a set of replicable Attention Modules and gating mechanisms (see Section 3). Concretely, as it can be seen in Figure 1, any existing architecture can be aug-

mented by applying the proposed model at different depths, and replacing the original loss by the proposed one. It is remarkable that the modules are independent of the original path of the network, so in practice, it can be computed in parallel to the rest of the network. The proposed attention mechanism has been included in a strong baseline like Wide Residual Networks (WRN) [35], and applied on CIFAR-10, CIFAR-100 [15], and five challenging fine-grained recognition datasets. The resulting network, called Wide Attentional Residual Network (WARN) systematically enhances the performance of WRNs and surpasses the state of the art in various classification benchmarks.

2 Related Work

There are different approaches to fine-grained recognition [37]: (i) vanilla deep CNNs, (ii) CNNs as feature extractors for localizing parts and do alignment, (iii) ensembles, (iv) attention mechanisms. In this work, we focus on (iv), the attention mechanisms, which aim to discover the most discriminative parts of an image to be processed in greater detail, thus ignoring clutter and focusing on the most distinctive traits. These parts are central for fine-grained recognition, where the inter-class variance is small and the intra-class variance is high.

Different fine-grained attention mechanisms can be found in the literature. [33] proposed a *two-level attention* mechanism for fine-grained classification on different subsets of the ILSVRC [25] dataset, and the CUB200_2011. In this model, images are first processed by a bottom-up object proposal network based on R-CNN [36] and selective search [28]. Then, the softmax scores of another ILSVRC2012 pre-trained CNN, which they call *FilterNet*, are thresholded to prune the patches with the lowest parent class score. These patches are then classified to fine-grained categories with a *DomainNet*. Spectral clustering is also used on the *DomainNet* filters in order to extract parts (head, neck, body, etc.), which are classified with an SVM. Finally, the part- and object-based classifier scores are merged to get the final prediction. The *two-level attention* obtained state of the art results on CUB200-2011 with only class-level supervision. However, the pipeline must be carefully fine-tuned since many stages are involved with many hyper-parameters.

Differently from *two-level attention*, which consists of independent processing and it is not end-to-end, Sermanet *et al.* proposed to use a deep CNN and a Recurrent Neural Network (RNN) to accumulate high multi-resolution “glimpses” of an image to make a final prediction [26]. However, reinforcement learning slows down convergence and the RNN adds extra computation steps and parameters.

A more efficient approach was presented by Liu *et al.* [19], where a fully-convolutional network is trained with reinforcement learning to generate confidence maps on the image and use them to extract the parts for the final classifiers whose scores are averaged. Compared to previous approaches, in the work done by [19], multiple image regions are proposed in a single timestep thus, speeding up the computation. A greedy reward strategy is also proposed in order to increase the training speed. The recent approach presented by [5] uses a clas-

sification network and a recurrent attention proposal network that iteratively refines the center and scale of the input (RA-CNN). A ranking loss is used to enforce incremental performance at each iteration.

Zhao *et al.* proposed to enforce multiple non-overlapped attention regions [38]. The overall architecture consists of an attention canvas generator, which extracts patches of different regions and scales from the original image; a VGG-16 [27] CNN is then used to extract features from the patches, which are aggregated with a long short-term memory [9] that attends to non-overlapping regions of the patches. Classification is performed with the average prediction in each region. Similarly, in [39], they proposed the Multi-Attention CNN (MA-CNN) to learn to localize informative patches from the output of a VGG-19 and use them to train an ensemble of part classifiers.

In [12], they propose to extract global features from the last layers of a CNN, just before the classifier and use them to attend relevant regions in lower level feature activations. The attended activations from each level are then spatially averaged, channel-wise concatenated, and fed to the final classifier. The main differences with [12] are: (i) attention maps are computed in parallel to the base model, while the model in [12] requires output features for computing attention maps; (ii) WARN uses fewer parameters, so dropout is not needed to obtain competitive performance (these two factors clearly reflect in gain of speed); and (iii) gates allow our model to ignore/attend different information to improve the performance of the original model, while in [12] the full output function is replaced. As a result, WARN obtains 3.44% error on CIFAR10, outperforming [12] while being 7 times faster w/o parallelization.

All the previously described methods involve multi-stage pipelines and most of them are trained using reinforcement learning (which requires sampling and makes them slow to train). In contrast, STNs, FAM, the model in [12], and our approach jointly propose the attention regions and classify them in a single pass. Moreover, different from STNs and FAM our approach only uses one CNN stream, it can be used on pre-trained models, and it is far more computationally efficient than STNs, FAM, and [12] as described next.

3 Our approach

Our approach consists of a universal attention module that can be added after each convolutional layer without altering pre-defined information pathways of any architecture (see Figure 1). This is helpful since it seamlessly augments any architecture such as VGG and ResNet with no extra supervision, *i.e.* no part labels are necessary. Furthermore, it also allows being plugged into any existing trained network to quickly perform transfer learning approaches.

The attention module consists of three main submodules depicted in Figure 2 (a): (i) the attention heads \mathcal{H} , which define the most relevant regions of a feature map, (ii) the output heads \mathcal{O} , generate an hypothesis given the attended information, and (iii) the confidence gates \mathcal{G} , which output a confidence score for

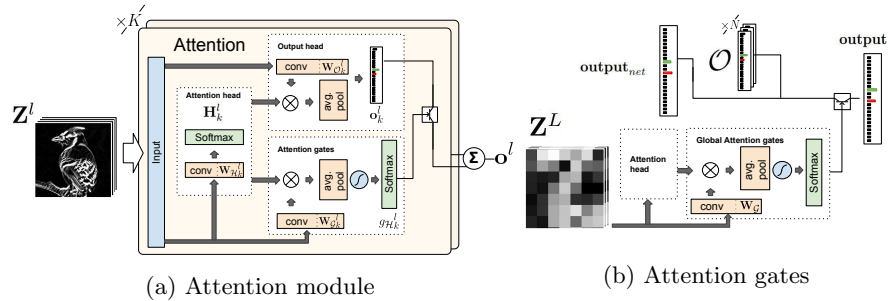


Fig. 2: (a) Attention Module: K attention heads \mathbf{H}_k^l are applied to a feature map \mathbf{Z}^l , and information is aggregated with the layer attention gates. (b) Global attention: global information from the last feature map \mathbf{Z}^L is used to compute the gating scores that produce the final **output** as the weighted average of the outputs of the attention modules and the original network \mathbf{output}_{net}

each attention head. Each of these modules is described in detail in the following subsections.

3.1 Overview

As it can be seen in Figure 1, a convolution layer is applied to the output of the augmented layer, producing K attentional heatmaps. These attentional maps are then used to spatially average the local class probability scores for each of the feature maps, and produce the final class probability vector. This process is applied to an arbitrary number N of layers, producing N class probability vectors. Then, the model learns to correct the initial prediction by attending the lower-level class predictions. This is the final combined prediction of the network. In terms of probability, the network corrects the initial likelihood by updating the prior with local information.

3.2 Attention head

Inspired by [38] and the *transformer* architecture presented by [30], and following the notation established by [35], we have identified two main dimensions to define attentional mechanisms: (i) the number of layers using the attention mechanism, which we call *attention depth* (AD), and (ii) the number of attention heads in each attention module, which we call *attention width* (AW). Thus, a desirable property for any universal attention mechanism is to be able to be deployed at any arbitrary *depth* and *width*.

This property is fulfilled by including K attention heads \mathcal{H}_k (width), depicted in Figure 1, into each attention module (depth)³. Then, the attention heads at

³ Notation: $\mathcal{H}, \mathcal{O}, \mathcal{G}$ are the set of attention heads, output heads, and attention gates respectively. Uppercase letters refer to functions or constants, and lowercase ones to indices. Bold uppercase letters represent matrices and bold lowercase ones vectors.

layer $l \in [1..L]$, receive the feature activations $\mathbf{Z}^l \in \mathbb{R}^{c \times h \times w}$ of that layer as input, and output K attention masks:

$$\mathbf{H}^l = \text{spatial_softmax}(\mathbf{W}_{\mathcal{H}^l} * \mathbf{Z}^l), \quad (1)$$

where $\mathbf{H}^l \in \mathbb{R}^{K \times h \times w}$ is the output matrix of the l^{th} attention module, $\mathbf{W}_{\mathcal{H}^l} : \mathbb{R}^{c \times h \times w} \rightarrow \mathbb{R}^{K \times h \times w}$ is a convolution kernel with output dimensionality K used to compute the attention masks corresponding to the attention heads \mathbf{H}_k , and $*$ denotes the convolution operator. The *spatial_softmax*, which performs the *softmax* operation channel-wise on the spatial dimensions of the input, is used to enforce the model to learn the most relevant region of the image. Sigmoid units could also be used at the risk of degeneration to all-zeros or all-ones. To prevent the attention heads at the same depth to collapse into the same region, we apply the regularizer proposed in [38].

3.3 Output head

To obtain the class probability scores, the input feature map \mathbf{Z}_k^l is convolved with a kernel:

$$\mathbf{W}_{\mathcal{O}_k^l} \in \mathbb{R}^{\text{channels} \times h \times w} \rightarrow \mathbb{R}^{\#\text{labels} \times h \times w},$$

h, w represent the spatial dimensions, and *channels* is the number of input channels to the module. This results on a spatial map of class probability scores:

$$\mathbf{O}_k^l = \mathbf{W}_{\mathcal{O}_k^l} * \mathbf{Z}_k^l. \quad (2)$$

Note that this operation can be done in a single pass for all the K heads by setting the number of output channels to $\#\text{labels} \cdot K$. Then, class probability vectors \mathbf{O}_k^l are weighted by the spatial attention scores and spatially averaged:

$$\mathbf{o}_k^l = \sum_{x,y} \mathbf{H}_k^l \odot \mathbf{O}_k^l, \quad (3)$$

where \odot is the element-wise product, and $x \in \{1..width\}, y \in \{1..height\}$. The attention scores \mathbf{H}_k^l are a 2d flat mask and the product with each of the input channels of \mathbf{Z}^l is done by broadcasting, *i.e.* repeating \mathbf{H}_k^l for each of the channels of \mathbf{Z}^l .

3.4 Layered attention gates

The final output \mathbf{o}^l of an attention module is obtained by a weighted average of the K output probability vectors, through the use of head attention gates $\mathbf{g}_{\mathcal{H}^l} \in \mathbb{R}^{|\mathcal{H}^l|}$, $\sum_k g_{\mathcal{H}^l} = 1$.

$$\mathbf{o}^l = \sum_k g_{\mathcal{H}^l} \mathbf{o}_k^l. \quad (4)$$

Where $g_{\mathcal{H}}$ is obtained by first convolving \mathbf{Z}^l with

$$\mathbf{W}_{\mathbf{g}}^l \in \mathbb{R}^{channels \times h \times w} \rightarrow \mathbb{R}^{|\mathcal{H}| \times h \times w},$$

and then performing a spatial weighted average:

$$\mathbf{g}_{\mathcal{H}}^l = softmax(\tanh(\sum_{x,y} (\mathbf{W}_{\mathbf{g}}^l * \mathbf{Z}^l) \odot \mathbf{H}_l)). \quad (5)$$

This way, the model learns to choose the attention head that provides the most meaningful output for a given attention module.

3.5 Global attention gates

In order to let the model learn to choose the most discriminative features at each depth to disambiguate the output prediction, a set of relevance scores \mathbf{c} are predicted at the model output, one for each attention module, and one for the final prediction. This way, through a series of gates, the model can learn to query information from each level of the network conditioned to the global context. Note that, unlike in [12], the final predictions do not act as a bottleneck to compute the output of the attention modules.

The relevance scores are obtained with an inner product between the last feature activation of the network \mathbf{Z}^L and the gate weight matrix $\mathbf{W}_{\mathcal{G}}$:

$$\mathbf{c} = \tanh(\mathbf{W}_{\mathcal{G}} \mathbf{Z}^L). \quad (6)$$

The gate values $\mathbf{g}_{\mathcal{O}}$ are then obtained by normalizing the set of scores by means of a *softmax* function:

$$g_{\mathcal{O}}^l = \frac{e^{c_k}}{\sum_{i=1}^{|\mathcal{G}|} e^{c_i}}, \quad (7)$$

where $|\mathcal{G}|$ is the total number of gates, and c_i is the i^{th} confidence score from the set of all confidence scores. The final output of the network is the weighted sum of the attention modules:

$$\mathbf{output} = g_{net} \cdot \mathbf{output}_{net} + \sum_{l \in \{1..|\mathcal{O}|\}} g_{\mathcal{O}}^l \cdot \mathbf{o}^l, \quad (8)$$

where g_{net} is the gate value for the original network output (\mathbf{output}_{net}), and \mathbf{output} is the final output taking the attentional predictions \mathbf{o}^l into consideration. Note that setting the output of \mathcal{G} to $\frac{1}{|\mathcal{O}|}$, corresponds to averaging all the outputs. Likewise, setting $\{\mathcal{G} \setminus G_{output}\} = 0, G_{output} = 1$, *i.e.* the set of attention gates is set to zero and the output gate to one, corresponds to the original pre-trained model without attention.

It is worth noting that all the operations that use \mathbf{Z}^l can be aggregated into a single convolution operation. Likewise, the fact that the attention mask is generated by just one convolution operation, and that most masking operations



Fig. 3: Samples from the five fine-grained datasets. (a) Adience, (b) CUB200 Birds, (c) Stanford Cars, (d) Stanford Dogs, (e) UEC-Food100

are directly performed in the label space, or can be projected into a smaller dimensionality space, makes the implementation highly efficient. Additionally, the direct access to the output gradients makes the module fast to learn, thus being able to generate foreground masks from the beginning of the training and refining them during the following epochs.

4 Experiments

We empirically demonstrate the impact on the accuracy and robustness of the different modules in our model on Cluttered Translated MNIST and then compare it with state-of-the-art models such as DenseNets and ResNeXt. Finally, we demonstrate the universality of our method for fine-grained recognition through a set of experiments on five fine-grained recognition datasets, as detailed next.

4.1 Datasets

Cluttered Translated MNIST⁴ Consists of 40×40 images containing a randomly placed MNIST [16] digit and a set of D randomly placed distractors, see Figure 5b. The distractors are random 8×8 patches from other MNIST digits.

CIFAR⁵ The CIFAR dataset consists of 60K 32×32 images in 10 classes for CIFAR-10, and 100 for CIFAR-100. There are 50K training and 10K test images.

Stanford Dogs [13]. The Stanford Dogs dataset consists of 20.5K images of 120 breeds of dogs, see Figure 3d. The dataset splits are fixed and they consist of 12k training images and 8.5K validation images.

UEC Food 100 [20]. A Japanese food dataset with 14K images of 100 different dishes, see Figure 3e. In order to follow the standard procedure (*e.g.* [2,6]), bounding boxes are used to crop the images before training.

Adience dataset [4]. The adience dataset consists of 26.5 K images distributed in eight age categories (02, 46, 813, 1520, 2532, 3843, 4853, 60+), and gender

⁴ <https://github.com/deepmind/mnist-cluttered>

⁵ <https://www.cs.toronto.edu/~kriz/cifar.html>

labels. A sample is shown in Figure 3a. The performance on this dataset is measured using 5-fold cross-validation.

Stanford Cars [14]. The Cars dataset contains 16K images of 196 classes of cars, see Figure 3c. The data is split into 8K training and 8K testing images.

Caltech-UCSD Birds 200 [31]. The CUB200-2011 birds dataset (see Figure 3b) consists of 6K train and 5.8K test bird images distributed in 200 categories. Although bounding box, segmentation, and attributes are provided, we perform raw classification as done by [11].

4.2 Ablation study

We evaluate the submodules of our method on the Cluttered Translated MNIST following the same procedure as in [21]. The proposed attention mechanism is used to augment a CNN with five 3×3 convolutional layers and two fully-connected layers in the end. The three first convolution layers are followed by Batch-normalization and a spatial pooling. Attention modules are placed starting from the fifth convolution (or pooling instead) backward until AD is reached. Training is performed with SGD for 200 epochs, and a learning rate of 0.1, which is divided by 10 after epoch 60. Models are trained on a $200k$ images train set, validated on a $100k$ images validation set, and tested on $100k$ test images. Weights are initialized using He *et al.* [7]. Figure 4 shows the effects of the different hyperparameters of the proposed model. The performance without attention is labeled as **baseline**. Attention models are trained with softmax attention gates and regularized with [38], unless explicitly specified.

First, we test the importance of AD for our model by increasingly adding attention layers with $AW = 1$ after each pooling layer. As it can be seen in Figure 4b, greater AD results in better accuracy, reaching saturation at $AD = 4$, note that for this value the receptive field of the attention module is 5×5 px , and thus the performance improvement from such small regions is limited. Figure 4c shows training curves for different values of AW , and $AD = 4$. As it can be seen, small performance increments are obtained by increasing the number of attention heads even with a single object present in the image.

Then, we use the best AD and AW , *i.e.* $AD, AW = 4$, to verify the importance of using softmax on the attention masks instead of sigmoid (1), the effect of using gates (Eq. 7), and the benefits of regularization [38]. Figure 4d confirms that ordered by importance: gates, softmax, and regularization result in accuracy improvement, reaching 97.8%. In particular, gates play an important role in discarding the distractors, especially for high AW and high AD .

Finally, in order to verify that attention masks are not overfitting on the data, and thus generalize to any amount of clutter, we run our best model so far (Figure 4d) on the test set with an increasing number of distractors (from 4 to 64). For the comparison, we included the baseline model before applying our approach and the same baseline augmented with an STN [11] that reached comparable performance as our best model in the validation set. All three models were trained with the same dataset with eight distractors. Remarkably, as it

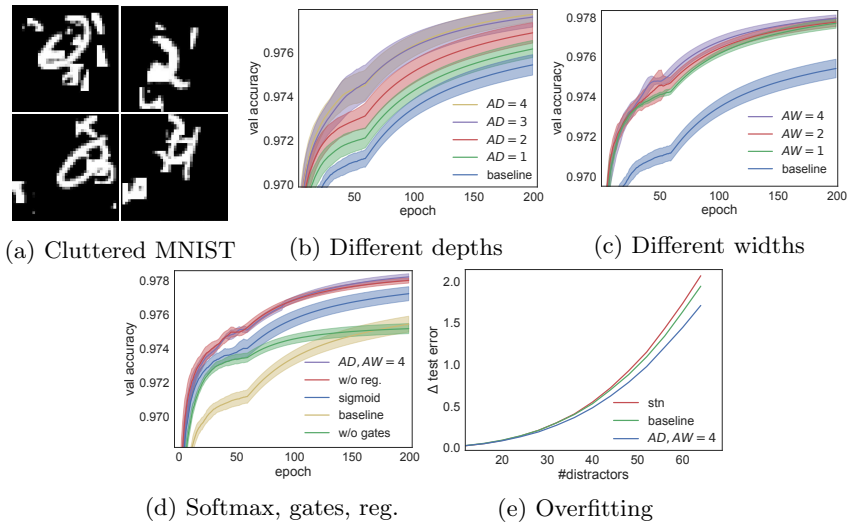


Fig. 4: Ablation experiments on Cluttered Translated MNIST. **baseline** indicates the original model before being augmented with attention. (a) shows a sample of the cluttered MNIST dataset. (b) the effect of increasing the attention depth (AD), for attention width $AW = 1$. (c) effect of increasing AW, for $AD=4$. (d) best performing model ($AD, AW = 4$, softmax attention gates, and regularization [38]) vs unregularized, sigmoid attention, and without gates. (e) test error of the baseline, attention ($AD, AW = 4$), and spatial transformer networks (stn), when trained with different amounts of distractors.

can be seen in Figure 4e, the attention augmented model demonstrates better generalization than the baseline and the STN.

4.3 Training from scratch

We benchmark the proposed attention mechanism on CIFAR-10 and CIFAR-100, and compare it with the state of the art. As a base model, we choose Wide Residual Networks, a strong baseline with a large number of parameters so that the additional parameters introduced by our model (WARN) could be considered negligible. The same WRN baseline is used to train an **att2** model [12], we refer to this model as WRN-att2. Models are initialized and optimized following the same procedure as in [35]. Attention Modules are systematically placed after each of the three convolutional groups starting by the last one until the attention depth has been reached in order to capture information at different levels of abstraction and fine-grained resolution, this same procedure is followed in [12]. The model is implemented with pytorch [22] and run on a single workstation with two NVIDIA 1080Ti.⁶

⁶ <https://github.com/prlz77/attend-and-rectify>

Table 1: Error rate on CIFAR-10 and CIFAR-100 (%). Results that surpass all other methods are in blue, results that surpass the baseline are in black bold font. Total network depth, attention depth, attention width, the usage of dropout, and the amount of floating point operations (Flop) are provided in columns 1-5 for fair comparison

	Net	Depth	AD	AW	Dropout	GFlop	CIFAR-10	CIFAR-100
Resnext [34]	29	-	-			10.7	3.58	17.31
Densenet [10]	250	-	-			5.4	3.62	17.60
	190	-	-			9.3	3.46	17.18
WRN [35]	28	-	-			5.2	4	19.25
	28	-	-	✓		5.2	3.89	18.85
	40	-	-	✓		8.1	3.8	18.3
WRN-att2 [12]	28	2	-			5.7	4.10	21.20
	28	2	-	✓		5.7	3.60	20.00
	40	2	-	✓		8.6	3.90	19.20
WARN	28	2	4			5.2	3.60	18.72
	28	3	4			5.3	3.45	18.61
	28	3	4	✓		5.3	3.44	18.26
	40	3	4	✓		8.2	3.46	17.82

First, the same ablation study performed in Section 4.2 is repeated on CIFAR100. We consistently reached the same conclusions as in Cluttered-MNIST: accuracy improves 1.5% by increasing attention depth from 1 to #residual.blocks, and width from 1 to 4. Gating performs 4% better than a simpler linear projection, and 3% with respect to simply averaging the output vectors. A 0.6% improvement is also observed when regularization is activated. Interestingly, we found sigmoid attention to perform similarly to softmax. With this setting, WARN reaches 17.82% error on CIFAR100. In addition, we perform an experiment blocking the gradients from the proposed attention modules to the original network to analyze whether the observed improvement is due to the attention mechanism or an optimization effect due to introducing shortcut paths to the loss function [18]. Interestingly, we observed a 0.2% drop on CIFAR10, and 0.4% on CIFAR100, which are still better than the baseline. Note that a performance drop should be expected, even without taking optimization into account, since backpropagation makes intermediate layers learn to gather more discriminative features for the attention layers. It is also worth noting that fine-grained accuracy improves even when fine-tuning (gradients are multiplied by 0.1 in the base model), see Section 4.4. In contrast, the approach in [12] does not converge when gradients are not sent to the base model since classification is directly performed on intermediate feature maps (which continuously shift during training).

As seen in Table 1, the proposed Wide Attentional Residual Network (WARN) improves the baseline model for CIFAR-10 and CIFAR-100 even without the use

Table 2: Number of parameters, floating point operations (Flop), time (s) per validation epoch, and error rates (%) on CIFAR-10 and CIFAR-100. The "Time" column shows the amount of seconds to forward the validation dataset with batch size 256 on a single GPU

	Depth	Params	GFlop	Time	CIFAR-10	CIFAR-100
ResNext	29	68M	10.7	5.02s	3.58	17.31
Densenet	190	26M	9.3	6.41s	3.46	17.18
WRN	40	56M	8.1	0.18s	3.80	18.30
WRN-att2	40	64M	8.6	0.24s	3.90	19.20
WARN	28	37M	5.3	0.17s	3.44	18.26
WARN	40	56M	8.2	0.18s	3.46	17.82




Fig. 5: Comparison of the best performing Resnext, Densenet, WRN, WRN-att2, and WARN on the CIFAR-10 and CIFAR-100. Validation accuracy is reported as a function of the number of GFLOPs.

of Dropout and outperforms the rest of the state of the art in CIFAR-10 while being remarkably faster, as it can be seen in Table 2. Remarkably, the performance on CIFAR-100 makes WARN competitive when compared with Densenet and Resnext, while being up to 36 times faster. We hypothesize that the increase in accuracy of the augmented model is limited by the base network and even better results could be obtained when applied on the best performing baseline. Interestingly, WARN shows superior performance even without the use of dropout; this was not possible with [12], which requires dropout to achieve competitive performances, since they introduce more parameters to the augmented network. The computing efficiency of the top performing models is shown in Figure 5. WARN provides the highest accuracy per GFlop on CIFAR-10, and is more competitive than WRN, and WRN-att2 on CIFAR-100.

Table 3: Results on six fine-grained recognition tasks. *DSP* means that the cited model uses Domain Specific Pre-training. *HR* means the cited model uses high-resolution images. Accuracies that improve the baseline model are in black bold font, and highest accuracies are in blue

	Dogs	Food	Cars	Gender	Age	Birds
SotA	RA-CNN [5]	Inception [6]	MA-CNN [39]	FAM [23]	DEX [24]	MA-CNN [39]
DSP				✓	✓	
HR	✓		✓			✓
Accuracy	87.3	81.5	92.8	93.0	64.0	86.5
WRN	89.6	84.3	88.5	93.9	57.4	84.3
WARN	92.9	85.5	90.0	94.6	59.7	85.6

4.4 Transfer Learning

We fine-tune an augmented WRN-50-4 pre-trained on Imagenet [25] and report higher accuracy on five different fine-grained datasets: Stanford Dogs, UEC Food-100, Adience, Stanford Cars, CUB200-2001 compared to the WRN baseline. All the experiments are trained for 100 epochs, with a batch size of 64. The learning rate is first set to 10^{-3} to all layers except the attention modules and the classifier, for which it ten times higher. The learning rate is reduced by a factor of 0.1 every 30 iterations and the experiment is automatically stopped if a plateau is reached. The network is trained with standard data augmentation, *i.e.* random 224×224 patches are extracted from 256×256 images with random horizontal flips. Since the aim of this work is to demonstrate that the proposed mechanism universally improves the baseline CNNs for fine-grained recognition, we follow the same training procedure in all datasets. Thus, we do not use 512×512 images, which are central for state-of-the-art methods such as RA-CNNs, MA-CNNs, or color jitter [6] for food recognition. The proposed method is able to obtain state of the art results in Adience Gender, Stanford dogs and UEC Food-100 even when trained with lower resolution.

As seen in table 3, WRN substantially increase their accuracy on all benchmarks by just fine-tuning them with the proposed attention mechanism. Moreover, we report the highest accuracy scores on Stanford Dogs, UEC Food, and Gender recognition, and obtain competitive scores when compared with models that use high resolution images, or domain-specific pre-training. For instance, in [24] a domain-specific model pre-trained on millions of faces is used for age recognition, while our baseline is a general-purpose WRN pre-trained on the Imagenet. It is also worth noting that the performance increase on CUB200-2011 (+1.3%) is higher than the one obtained in STNs with 224×224 images (+0.8%) even though we are augmenting a stronger baseline. This points out that the proposed mechanism might be extracting complementary information that is not extracted by the main convolutional stream. As seen in table 4, WARN not only increases the absolute accuracy, but it provides a high efficiency per introduced parameter. A sample of the attention masks for each dataset is shown on Figure

Table 4: Increment of accuracy (%) per Million of parameters

	Dogs	Food	Cars	Gender	Age	Birds	Average
WRN	1.3	1.2	1.3	1.4	0.8	1.2	1.2
WARN	6.9	2.5	3.1	1.5	4.0	2.5	3.4




Fig. 6: Attention masks for each dataset: (a) Stanford dogs, (b) Stanford cars, (c) Adience gender, (d) CUB birds, (e) Adience age, (f) UEC food. As it can be seen, the masks help to focus on the foreground object. In (c), the attention mask focuses on ears for gender recognition, possibly looking for earrings

6. As it can be seen, the attention heads learn to ignore the background and to attend the most discriminative parts of the objects. This matches the conclusions of Section 4.2.

5 Conclusion

We have presented a novel attention mechanism to improve CNNs. The proposed model learns to attend the most informative parts of the CNN feature maps at different depth levels and combines them with a gating function to update the output distribution.

We suggest that attention helps to discard noisy uninformative regions, avoiding the network to memorize them. Unlike previous work, the proposed mechanism is modular, architecture independent, fast, simple, and yet WRN augmented with it obtain state-of-the-art results on highly competitive datasets while being 37 times faster than DenseNet, 30 times faster than ResNeXt, and making the augmented model more parameter-efficient. When fine-tuning on a transfer learning task, the attention augmented model showed superior performance in each recognition dataset. Moreover, state of the art performance is obtained on dogs, gender, and food. Results indicate that the model learns to extract local discriminative information that is otherwise lost when traversing the layers of the baseline architecture.

Acknowledgments Authors acknowledge the support of the Spanish project TIN2015-65464-R (MINECO/FEDER), the 2016FI B 01163 grant of Generalitat de Catalunya, and the COST Action IC1307 iV&L Net. We also gratefully acknowledge the support of NVIDIA Corporation with the donation of a Tesla K40 GPU and a GTX TITAN GPU, used for this research.

References

1. Anderson, J.R.: Cognitive psychology and its implications. New York, NY, US: WH Freeman/Times Books/Henry Holt and Co (1985)
2. Chen, J., Ngo, C.W.: Deep-based ingredient recognition for cooking recipe retrieval. In: ACM MM. pp. 32–41. ACM (2016)
3. Desimone, R., Duncan, J.: Neural mechanisms of selective visual attention. Annual review of neuroscience **18**(1), 193–222 (1995)
4. Eidinger, E., Enbar, R., Hassner, T.: Age and gender estimation of unfiltered faces. TIFS **9**(12), 2170–2179 (2014)
5. Fu, J., Zheng, H., Mei, T.: Look closer to see better: recurrent attention convolutional neural network for fine-grained image recognition. In: CVPR (2017)
6. Hassannejad, H., Matrella, G., Ciampolini, P., De Munari, I., Mordonini, M., Cagnoni, S.: Food image recognition using very deep convolutional networks. In: MADIMA Workshop. pp. 41–49. ACM (2016)
7. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: CVPR. pp. 1026–1034 (2015)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
9. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural computation **9**(8), 1735–1780 (1997)
10. Huang, G., Liu, Z., Weinberger, K.Q., van der Maaten, L.: Densely connected convolutional networks. In: CVPR. vol. 1, p. 3 (2017)
11. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. In: NIPS. pp. 2017–2025 (2015)
12. Jetley, S., Lord, N.A., Lee, N., Torr, P.: Learn to pay attention. In: ICLR (2018)
13. Khosla, A., Jayadevaprakash, N., Yao, B., Li, F.F.: Novel dataset for fine-grained image categorization: Stanford dogs. In: FGVC. vol. 2, p. 1 (2011)
14. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: CVPR. pp. 554–561 (2013)
15. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images (2009)
16. LeCun, Y.: The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/> (1998)
17. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. Proceedings of the IEEE **86**(11), 2278–2324 (1998)
18. Lee, C.Y., Xie, S., Gallagher, P., Zhang, Z., Tu, Z.: Deeply-supervised nets. In: AISTATS. pp. 562–570 (2015)
19. Liu, X., Xia, T., Wang, J., Lin, Y.: Fully convolutional attention localization networks: Efficient attention localization for fine-grained recognition. arXiv preprint arXiv:1603.06765 (2016)
20. Matsuda, Y., Hoashi, H., Yanai, K.: Recognition of multiple-food images by detecting candidate regions. In: ICME (2012)
21. Mnih, V., Heess, N., Graves, A., et al.: Recurrent models of visual attention. In: NIPS. pp. 2204–2212 (2014)
22. Paszke, A., Gross, S., Chintala, S., Chanan, G.: Pytorch (2017)
23. Rodriguez, P., Cucurull, G., Gonfau, J.M., Roca, F.X., Gonzalez, J.: Age and gender recognition in the wild with deep attention. PR (2017)
24. Rothe, R., Timofte, R., Gool, L.V.: Deep expectation of real and apparent age from a single image without facial landmarks. IJCV (July 2016)

25. Russakovsky, O., Deng, J., Krause, J., Berg, A., Fei-Fei, L.: The imagenet large scale visual recognition challenge 2012 (ilsvrc2012) (2012)
26. Sermanet, P., Frome, A., Real, E.: Attention for fine-grained categorization. In: ICLR (2015)
27. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
28. Uijlings, R., van de Sande, A., Gevers, T., Smeulders, M., et al.: Selective search for object recognition. *IJCV* **104**(2), 154 (2013)
29. Ungerleider, S.K., G, L.: Mechanisms of visual attention in the human cortex. *Annual review of neuroscience* **23**(1), 315–341 (2000)
30. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NIPS. pp. 5998–6008 (2017)
31. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The Caltech-UCSD Birds-200-2011 Dataset. Tech. rep. (2011)
32. Wang, F., Jiang, M., Qian, C., Yang, S., Li, C., Zhang, H., Wang, X., Tang, X.: Residual attention network for image classification. *CVPR* (2017)
33. Xiao, T., Xu, Y., Yang, K., Zhang, J., Peng, Y., Zhang, Z.: The application of two-level attention models in deep convolutional neural network for fine-grained image classification. In: *CVPR*. pp. 842–850 (2015)
34. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: *CVPR*. pp. 5987–5995. IEEE (2017)
35. Zagoruyko, S., Komodakis, N.: Wide residual networks. In: *BMVC* (2016)
36. Zhang, N., Donahue, J., Girshick, R., Darrell, T.: Part-based r-cnns for fine-grained category detection. In: *ECCV*. pp. 834–849. Springer (2014)
37. Zhao, B., Feng, J., Wu, X., Yan, S.: A survey on deep learning-based fine-grained object classification and semantic segmentation. *IJAC* pp. 1–17 (2017)
38. Zhao, B., Wu, X., Feng, J., Peng, Q., Yan, S.: Diversified visual attention networks for fine-grained object classification. *IEEE Transactions on Multimedia* **19**(6), 1245–1256 (2017)
39. Zheng, H., Fu, J., Mei, T., Luo, J.: Learning multi-attention convolutional neural network for fine-grained image recognition. In: *ICCV* (2017)