

Multi-Attention Multi-Class Constraint for Fine-grained Image Recognition

Ming Sun^{[0000-0002-5948-2708]¹}, Yuchen Yuan^{[0000-0002-5251-9942]¹},
Feng Zhou^{[0000-0002-1132-5877]²}, and Errui Ding^{[0000-0002-1867-5378]¹}

¹ Department of Computer Vision Technology (VIS), Baidu Inc.

² Baidu Research

{sunming05, yuanyuchen02, zhoulfeng09, dingerrui}@baidu.com

Abstract. Attention-based learning for fine-grained image recognition remains a challenging task, where most of the existing methods treat each object part in isolation, while neglecting the correlations among them. In addition, the multi-stage or multi-scale mechanisms involved make the existing methods less efficient and hard to be trained end-to-end. In this paper, we propose a novel attention-based convolutional neural network (CNN) which regulates multiple object parts among different input images. Our method first learns multiple attention region features of each input image through the one-squeeze multi-excitation (OSME) module, and then apply the multi-attention multi-class constraint (MAMC) in a metric learning framework. For each anchor feature, the MAMC functions by pulling same-attention same-class features closer, while pushing different-attention or different-class features away. Our method can be easily trained end-to-end, and is highly efficient which requires only one training stage. Moreover, we introduce Dogs-in-the-Wild, a comprehensive dog species dataset that surpasses similar existing datasets by category coverage, data volume and annotation quality. Extensive experiments are conducted to show the substantial improvements of our method on four benchmark datasets.

Keywords: Fine-grained Classification, Metric Learning, Visual Attention, Multi-Attention Multi-Class Constraint, One-Squeeze Multi-Excitation

1 Introduction

In the past few years, the performances of generic image recognition on large-scale datasets (*e.g.*, ImageNet [8], Places [56]) have undergone unprecedented improvements, thanks to the breakthroughs in the design and training of deep neural networks (DNNs). Such fast-pacing progresses in research have also drawn attention of the related industries to build software like Google Lens on smartphones to recognize everything snapshotted by the user. Yet, recognizing the fine-grained category of daily objects such as car models, animal species or food dishes is still a challenging task for existing methods. The reason is that the global geometry and appearances of fine-grained classes can be very similar, and

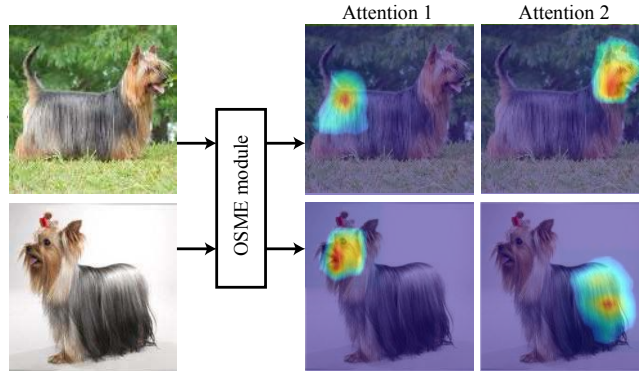


Fig. 1: Two distinct dog species from the proposed Dogs-in-the-Wild dataset. Our method is capable of capturing the subtle differences on the head and tail without manual part annotations.

how to identify their subtle differences on the key parts is of vital importance. For instance, to differentiate the two dog species in Figure 1, it is important to consider their discriminative features on the ear, tail and body length, which is extremely difficult to notice even for human without domain expertise.

Thus the majority of efforts in the fine-grained community focus on how to effectively integrate part localization into the classification pipeline. In the pre-DNN era, various parametric [9,24,29] and non-parametric [25] part models have been employed to extract discriminative part-specific features. Recently, with the popularity of DNNs, the tasks of object part localization and feature representation can be both learned in a more effective way [22,18,2,48,49]. The major drawback of these strongly-supervised methods, however, is that they heavily rely on manual object part annotations, which is too expensive to be prevalently applied in practice. Therefore, weakly-supervised frameworks have received increasing attention in recent researches. For instance, the attention mechanism can be implemented as sequential decision processes [27] or multi-stream part selections [10] without the need of part annotations. Despite the great progresses, these methods still suffer several limitations. First, their additional steps, such as the part localization and feature extraction of the attended regions, can incur expensive computational cost. Second, their training procedures are sophisticated, requiring multiple alternations or cascaded stages due to the complex architecture designs. More importantly, most works tend to detect the object parts in isolation, while neglect their inherent correlations. As a consequence, the learned attention modules are likely to focus on the same region and lack the capability to localize multiple parts with discriminative features that can differentiate between similar fine-grained classes.

From extensive experimental studies, we observe that an effective visual attention mechanism for fine-grained classification should follow three criteria: 1) The detected parts should be well spread over the object body to extract non-

correlated features; 2) Each part feature alone should be discriminative for separating objects of different classes; 3) The part extractors should be lightweight in order to be scaled up for practical applications. To meet these demands, this paper presents a novel framework that contains two major improvements. First, we propose one-squeeze multi-excitation module (OSME) to localize different parts inspired by the latest ImageNet winner SENet [13]. It is fully differentiable and can directly extract part features with budgeted computational cost. Second, inspired by metric learning loss, we propose the multi-attention multi-class constraint (MAMC) to coherently enforce the correlations among different parts in training. In addition, we have released a large scale dog species dataset named Dogs-in-the-Wild, which exhibits higher category coverage, data volume and annotation quality than similar public datasets. Experimental results show that our method achieves substantial improvements on four benchmark datasets. Moreover, our method can be easily trained end-to-end, and unlike many existing methods that require multiple feedforward processes for feature extraction [41,52] or multiple alternative training stages [10,31], only one stage and one feedforward are required for each training step.

2 Related Work

2.1 Fine-Grained Image Recognition

In the task of fine-grained image recognition, since the inter-class differences are subtle, more specialized techniques, including discriminative feature learning and object parts localization, need to be applied. A straightforward way is supervised learning with manual object part annotations, which has shown promising results in classifying birds [2,9,48,49], dogs [16,29,25,48], and cars [17,24,20]. However, it is usually laborious and expensive to obtain object part annotations, which severely restricts the effectiveness of such methods.

Consequently, more recently proposed methods tend to localize object parts with weakly-supervised mechanisms, such as the combination of pose alignment and co-segmentation [18], dynamic spatial transformation of the input image for better alignment [14], and parallel CNNs for bilinear feature extraction [23].

Compared with previous works, our method also takes a weakly-supervised mechanism, but can directly extract the part features without cropping them out, and is highly efficient to be scaled up with multiple parts.

In recent years, more advanced methods emerge with improved results. For instance, the bipartite-graph labeling [57] leverages the label hierarchy on the fine-grained classes, which is less expensive to obtain. The work in [51] exploit unified CNN framework with spatially weighted representation by the Fisher vector [30]. [3] and [45] incorporate human knowledge and various types of computer vision algorithms into a human-in-the-loop framework for the complementary strengths of both ends. In [34], the average and bilinear pooling are combined to learn the pooling strategy during training. [6] uses the dataset bootstrapping with the help of human. And the work in [50], the structures of label are ex-

ploited. These techniques can also be potentially combined with our method for further works.

2.2 Visual Attention

The aforementioned part-based methods have shown strong performances in fine-grained image recognition. Nevertheless, one of their major drawbacks is that they need meaningful definitions of the object parts, which are hard to obtain for non-structured objects such as flowers [28] and food dishes [1]. Therefore, the methods enabling CNN to attend loosely defined regions for general objects have emerged as a promising direction.

For instance, the soft proposal network [58] combines random walk and CNN for object proposals. The works in [52] and [26] introduce long short-term memory [12] and reinforcement learning to attention-based classification, respectively. And the class activation mapping [55] generates the heatmap of the input image, which provides a better way for attention visualization. On the other hand, the idea of multi-scale feature fusion or recurrent learning has become increasingly popular in recent works. For instance, the work in [31] extends [55] and establishes a cascaded multi-stage framework, which refines the attention region by iteration. The residual attention network [41] obtains the attention mask of input image by up-sampling and down-sampling, and a series of such attention modules are stacked for feature map refinement. And the recurrent attention CNN [10] alternates between the optimization of softmax and pairwise ranking losses, which jointly contribute to the final feature fusion. Even an acceleration method [21] with reinforcement learning is proposed particularly for the recurrent attention models above.

In parallel to these efforts, our method not only automatically localizes the attention regions, but also directly captures the corresponding features without explicitly cropping the ROI and feedforwarding again for the feature, which makes our method highly efficient.

2.3 Metric Learning

Apart from the techniques above, deep metric learning aims at the learning of appropriate similarity measurements between sample pairs, which provides another promising direction to fine-grained image recognition. The pioneer work of Siamese network [4] formulates the deep metric learning with a contrastive loss that minimizes distance between positive pairs while keeps negative pairs apart. Despite its great success on face verification [33], contrastive embedding requires that training data contains real-valued precise pair-wise similarities or distances. The triplet loss [32] addresses this issue by optimizing the relative distance of the positive pair and one negative pair from three samples. It has been proven that triplet loss is extremely effective for fine-grained product search [43]. Later, triplet loss is improved to automatically search for discriminative patches [44]. Nevertheless, compared with softmax loss, triplet loss is difficult to train due to its slow convergence. To alleviate this issue, the N-pair loss [37] is introduced to

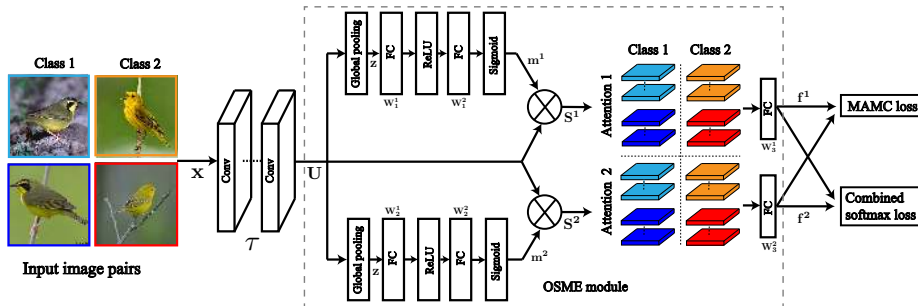


Fig. 2: Overview of our network architecture. Here we visualize the case of learning two attention branches given a training batch with four images of two classes. The MAMC and softmax losses would be replaced by a softmax layer in testing. Unlike hard-attention methods like [10], we do not explicitly crop the parts out. Instead, the feature maps (S^1 and S^2) generated by the two branches provide soft response for attention regions such as the birds’ head or torso, respectively.

consider multiple negative samples in training, and exhibits higher efficiency and performance. More recently, the angular loss [42] enhances N-pair loss by integrating high-order constraint that captures additional local structure of triplet triangles.

Our method differs previous metric learning works in two aspects: First, we take object parts instead of the whole images as instances in the feature learning process; Second, our formulation simultaneously considers the part and class labels of each instance.

3 Proposed Method

In this section, we present our proposed method which can efficiently and accurately attend discriminative regions despite being trained only on image-level labels. As shown in Figure 2, the framework of our method is composed by two parts: 1) A differentiable one-squeeze multi-excitation (OSME) module that extracts features from multiple attention regions with a slight increase in computational burden. 2) A multi-attention multi-class (MAMC) constraint that enforces the correlation of the attention features in favor of the fine-grained classification task. In contrast to many prior works, the entire network of our method can be effectively trained end-to-end in one stage.

3.1 One-Squeeze Multi-Excitation Attention Module

There have been a number of visual attention models exploring weakly supervised part localization, and the previous works can be roughly categorized in two groups. The first type of attention is also known as part detection, *i.e.*,

each attention is equivalent to a bounding box covering a certain area. Well-known examples include the early work of recurrent visual attention [27], the spatial transformer networks [14], and the recent method of recurrent attention CNN [10]. This hard-attention setup can benefit a lot from the object detection community in the formulation and training. However, its architectural design is often cumbersome as the part detection and feature extraction are separated in different modules. The second type of attention can be considered as imposing a soft mask on the feature map, which origins from activation visualization [46,54]. Later, people find it can be extended for localizing parts [55,31] and improving the overall recognition performance [41,13]. Our approach also falls into this category. We adopt the idea of SENet [13], the latest ImageNet winner, to capture and describe multiple discriminative regions in the input image. Compared to other soft-attention works [55,41], we build on SENet because of its superiority in performance and scalability in practice.

As shown in Figure 2, our framework is a feedforward neural network where each image is first processed by a base network, *e.g.*, ResNet-50 [11]. Let $\mathbf{x} \in \mathbb{R}^{W' \times H' \times C'}$ denote the input fed into the last residual block τ . The goal of SENet is to re-calibrate the output feature map,

$$\mathbf{U} = \tau(\mathbf{x}) = [\mathbf{u}_1, \dots, \mathbf{u}_C] \in \mathbb{R}^{W \times H \times C}, \quad (1)$$

through a pair of squeeze-and-excitation operations. In order to generate P attention-specific feature maps, we extend the idea of SENet by performing one-squeeze but multi-excitation operations.

In the first one-squeeze step, we aggregate the feature maps \mathbf{U} across spatial dimensions $W \times H$ to produce a channel-wise descriptor $\mathbf{z} = [z_1, \dots, z_C] \in \mathbb{R}^C$. The global average pooling is adopted as a simple but effective way to describe each channel statistic:

$$z_c = \frac{1}{WH} \sum_{w=1}^W \sum_{h=1}^H \mathbf{u}_c(w, h). \quad (2)$$

In the second multi-excitation step, a gating mechanism is independently employed on \mathbf{z} for each attention $p = 1, \dots, P$:

$$\mathbf{m}^p = \sigma\left(\mathbf{W}_2^p \delta(\mathbf{W}_1^p \mathbf{z})\right) = [m_1^p, \dots, m_C^p] \in \mathbb{R}^C, \quad (3)$$

where σ and δ refer to the Sigmoid and ReLU functions respectively. We adopt the same design of SENet by forming a pair of dimensionality reduction and increasing layers parameterized with $\mathbf{W}_1^p \in \mathbb{R}^{\frac{C}{r} \times C}$ and $\mathbf{W}_2^p \in \mathbb{R}^{C \times \frac{C}{r}}$. Because of the property of the Sigmoid function, each \mathbf{m}^p encodes a non-mutually-exclusive relationship among channels. We therefore use it to re-weight the channels of the original feature map \mathbf{U} ,

$$\mathbf{S}^p = [m_1^p \mathbf{u}_1, \dots, m_C^p \mathbf{u}_C] \in \mathbb{R}^{W \times H \times C}. \quad (4)$$

To extract attention-specific features, we feed each attention map \mathbf{S}^p to a fully connected layer $\mathbf{W}_3^p \in \mathbb{R}^{D \times WHC}$:

$$\mathbf{f}^p = \mathbf{W}_3^p \text{vec}(\mathbf{S}^p) \in \mathbb{R}^D, \quad (5)$$

where the operator $\text{vec}(\cdot)$ flattens a matrix into a vector.

In a nutshell, the proposed OSME module seeks to extract P feature vectors $\{\mathbf{f}^p\}_{p=1}^P$ for each image \mathbf{x} by adding a few layers on top of the last residual block. Its simplicity enables the use of relatively deep base networks and an efficient one-stage training pipeline.

It is worth to clarify that the SENet is originally not designed for learning visual attentions. By adopting the key idea of SENet, our proposed OSME module implements a lightweight yet effective attention mechanism that enables an end-to-end one-stage training on large-scale fine-grained datasets.

3.2 Multi-Attention Multi-Class Constraint

Apart from the attention mechanism introduced in Section 3.1, the other crucial problem is how to guide the extracted attention features to the correct class label. A straightforward way is to directly evaluate the softmax loss on the concatenated attention features [14]. However, the softmax loss is unable to regulate the correlations between attention features. As an alternative, another line of research [27,26,10] tends to mimic human perception with a recurrent search mechanism. These approaches iteratively generate the attention region from coarse to fine by taking previous predictions as references. The limitation of them, however, is that the current prediction is highly dependent on the previous one, thereby the initial error could be amplified by iteration. In addition, they require advanced techniques such as reinforcement learning or careful initialization in a multi-stage training. In contrast, we take a more practical approach by directly enforcing the correlations between parts in training. There has been some prior works like [44] that introduce geometrical constraints on local patches. Our method, on the other hand, explores much richer correlations of object parts by the proposed multi-attention multi-class constraint (MAMC).

Suppose that we are given a set of training images $\{(\mathbf{x}, y), \dots\}$ of K fine-grained classes, where $y = 1, \dots, K$ denotes the label associated with the image \mathbf{x} . To model both the within-image and inter-class attention relations, we construct each training batch, $\mathcal{B} = \{(\mathbf{x}_i, \mathbf{x}_i^+, y_i)\}_{i=1}^N$, by sampling N pairs of images³ similar to [37]. For each pair $(\mathbf{x}_i, \mathbf{x}_i^+)$ of class y_i , the OSME module extracts P attention features $\{\mathbf{f}_i^p, \mathbf{f}_i^{p+}\}_{p=1}^P$ from multiple branches according to Eq. 5.

Given $2N$ samples in each batch (Figure 3a), our intuition comes from the natural clustering of the $2NP$ features (Figure 3b) extracted by the OSME modules. By picking \mathbf{f}_i^p , which corresponds to the i^{th} class and p^{th} attention region as the anchor, we divide the rest features into four groups:

- same-attention same-class features, $\mathcal{S}_{\text{asc}}(\mathbf{f}_i^p) = \{\mathbf{f}_i^{p+}\}$;
- same-attention different-class features, $\mathcal{S}_{\text{sadc}}(\mathbf{f}_i^p) = \{\mathbf{f}_j^p, \mathbf{f}_j^{p+}\}_{j \neq i}$;
- different-attention same-class features, $\mathcal{S}_{\text{dasc}}(\mathbf{f}_i^p) = \{\mathbf{f}_i^q, \mathbf{f}_i^{q+}\}_{q \neq p}$;

³ N stands for the number of sample pairs as well as the number of classes in a mini-batch. Limited by GPU memory, N is usually much smaller than K , the total number of classes in the entire training set.

- different-attention different-class features $\mathcal{S}_{dadc}(\mathbf{f}_i^p) = \{\mathbf{f}_j^q, \mathbf{f}_j^{q+}\}_{j \neq i, q \neq p}$.

Our goal is to excavate the rich correlations among the four groups in a metric learning framework. As summarized in Figure 3c, we compose three types of triplets according to the choice of the positive set for the anchor \mathbf{f}_i^p . To keep notation concise, we omit \mathbf{f}_i^p in the following equations.

Same-attention same-class positives. The most similar feature to the anchor \mathbf{f}_i^p is \mathbf{f}_i^{p+} , while all the other features should have larger distance to the anchor. The positive and negative sets are then defined as:

$$\mathcal{P}_{sasc} = \mathcal{S}_{sasc}, \mathcal{N}_{sasc} = \mathcal{S}_{sadc} \cup \mathcal{S}_{dasc} \cup \mathcal{S}_{dadc}. \quad (6)$$

Same-attention different-class positives. For the features from different classes but extracted from the same attention region, they should be more similar to the anchor than the ones also from different attentions:

$$\mathcal{P}_{sadc} = \mathcal{S}_{sadc}, \mathcal{N}_{sadc} = \mathcal{S}_{dadc}. \quad (7)$$

Different-attention same-class positives. Similarly, for the features from same class but extracted from different attention regions, we have:

$$\mathcal{P}_{dasc} = \mathcal{S}_{dasc}, \mathcal{N}_{dasc} = \mathcal{S}_{dadc}. \quad (8)$$

For any positive set $\mathcal{P} \in \{\mathcal{P}_{sasc}, \mathcal{P}_{sadc}, \mathcal{P}_{dasc}\}$ and negative set $\mathcal{N} \in \{\mathcal{N}_{sasc}, \mathcal{N}_{sadc}, \mathcal{N}_{dasc}\}$ combinations, we expect the anchor to be closer to the positive than to any negative by a distance margin $m > 0$, *i.e.*,

$$\|\mathbf{f}_i^p - \mathbf{f}^+\|^2 + m \leq \|\mathbf{f}_i^p - \mathbf{f}^-\|^2, \forall \mathbf{f}^+ \in \mathcal{P}, \mathbf{f}^- \in \mathcal{N}. \quad (9)$$

To better understand the three constraints, let's consider the synthetic example of six feature points shown in Figure 4. In the initial state (Figure 4a), the \mathcal{S}_{sasc} feature point (green hexagon) stays further away from the anchor \mathbf{f}_i^p at the center than the others. After applying the first constraint (Eq. 6), the underlying feature space is transformed to Figure 4b, where the \mathcal{S}_{sasc} positive point (green \checkmark) has been pulled towards the anchor. However, the four negative features (cyan rectangles and triangles) are still in disordered positions. In fact, \mathcal{S}_{sadc} and \mathcal{S}_{dasc} should be considered as the positives compared to \mathcal{S}_{dadc} given the anchor. By further enforcing the second (Eq. 7) and third (Eq. 8) constraints, a better embedding can be achieved in Figure 4c, where \mathcal{S}_{sadc} and \mathcal{S}_{dasc} are regularized to be closer to the anchor than the ones of \mathcal{S}_{dadc} .

3.3 Training Loss

To enforce the triplet constraint in Eq. 9, a common approach is to minimize the following hinge loss:

$$\left[\|\mathbf{f}_i^p - \mathbf{f}^+\|^2 - \|\mathbf{f}_i^p - \mathbf{f}^-\|^2 + m \right]_+. \quad (10)$$

Despite being broadly used, optimizing Eq. 10 using standard triplet sampling leads to slow convergence and unstable performance in practice. Inspired by the

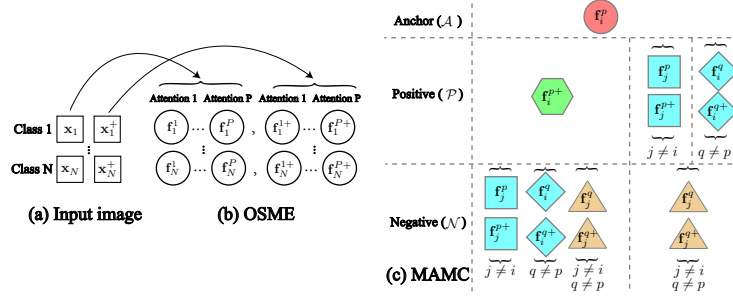


Fig. 3: Data hierarchy in training. (a) Each batch is composed by $2N$ input images in N-pair style. (b) OSME extracts P features for each image according to Eq. 5. (c) The group of features for three MAMC constraints by picking one feature \mathbf{f}_i^p as the anchor.




Fig. 4: Feature embedding of a synthetic batch. (a) Initial embedding before learning. (b) The result embedding by applying Eq. 6. (c) The final embedding by enforcing Eq. 7 and Eq. 8. See text for more details.

recent advance in metric learning, we enforce each of the three constraints by minimizing the N-pair loss⁴ [37],

$$L^{np} = \frac{1}{N} \sum_{\mathbf{f}_i^p \in \mathcal{B}} \left\{ \sum_{\mathbf{f}^+ \in \mathcal{P}} \log \left(1 + \sum_{\mathbf{f}^- \in \mathcal{N}} \exp(\mathbf{f}_i^{pT} \mathbf{f}^- - \mathbf{f}_i^{pT} \mathbf{f}^+) \right) \right\}. \quad (11)$$

In general, for each training batch \mathcal{B} , MAMC jointly minimizes the softmax loss and the N-pair loss with a weight parameter λ :

$$L^{mamc} = L^{softmax} + \lambda (L_{sasc}^{np} + L_{sadc}^{np} + L_{dasc}^{np}). \quad (12)$$

Given a batch of N images and P parts, MAMC is able to generate $2(PN - 1) + 4(N - 1)^2(P - 1) + 4(N - 1)(P - 1)^2$ constraints of three types (Eq. 6 to Eq. 8), while the N-pair loss can only produce $N - 1$. To put it in perspective, we are able to generate $130\times$ more constraints than N-pair loss with the same

⁴ It is worth to point out that the implementation of MAMC is independent to the use of N-pair loss, as MAMC is a general framework that can be combined with other triplet-based metric learning loss as well. The N-pair loss is taken as a reference because of its robustness and good convergence in practice.

data under the normal setting where $P = 2$ and $N = 32$. This implies that MAMC leverages much richer correlations among the samples, and is able to obtain better convergence than either triplet or N-pair loss.

4 The Dogs-in-the-Wild Dataset

Large image datasets (such as ImageNet [8]) with high-quality annotations enables the dramatic development in visual recognition. However, most datasets for fine-grained recognition are out-dated, non-natural and relatively small (as shown in Table 1). Recently, there are several attempts such as Goldfinch [19] and the iNaturalist Challenge [38] in building large-scale fine-grained benchmarks. However, there still lacks a comprehensive dataset with large enough data volume, highly accurate data annotation, and full tag coverage of common dog species. We hence introduce the Dogs-in-the-Wild dataset with 299,458 images of 362 dog categories⁵, which is $15\times$ larger than Stanford Dogs [16]. We generate the list of dog species by combining multiple sources (*e.g.*, Wikipedia), and then crawl the images with search engines (*e.g.*, Google, Baidu). The label of each image is then checked with crowd sourcing. We further prune small classes with less than 100 images, and merge extremely similar classes by applying confusion matrix and manual validation. The whole annotation process is conducted three times to guarantee the annotation quality. Last but not least, since most of the experimental baselines are pre-trained on ImageNet, which has substantial category overlap with our dataset, we exclude any image of ImageNet from our dataset for fair evaluation.

Figure 5a and Figure 5b qualitatively compare our dataset with the two most relevant benchmarks, Stanford Dogs [16] and the dog section of Goldfinch [19]. It can be seen that our dataset is more challenging in two aspects: (1) The intra-class variation of each category is larger. For instance, almost all common patterns and hair colors of Staffordshire Bull Terriers are covered in our dataset, as illustrated in Figure 5a. (2) More surrounding environment types are covered, which includes but is not limited to, natural scenes, indoor scenes and even artificial scenes; and the dog itself could either be in its natural appearance or dressed up, such as the first Boston Terrier in Figure 5a. Another feature of our dataset is that all of our images are manually examined to minimize annotation errors. Although Goldfinch has comparable class number and data volume, it is common to find noisy images inside, as shown in Figure 5b.

We then demonstrate the statistics of the three datasets in Figure 5c and Table 1. It is observed that our dataset is significantly more imbalanced in term of images per category, which is more consistent with real-life situations, and notably increases the classification difficulty. Note that the curves in Figure 5c are smoothed for better visualization. On the other hand, the average images per category of our dataset is higher than the other two datasets, which contributes to its high intra-class variation, and makes it less vulnerable to overfitting.

⁵ <http://ai.baidu.com/broad/subordinate?dataset=canine>

Table 1: Statistics of the related datasets.

Dataset	#Class	#Train	#Test	#Avg. Train/Class
CUB-200-2011	200	5,994	5,794	30
Stanford Dogs	120	12,000	8,580	100
Stanford Cars	196	8,144	8,041	42
Goldfinch	515	342,632	-	665
Dogs-in-the-Wild	362	258,474	40,984	714

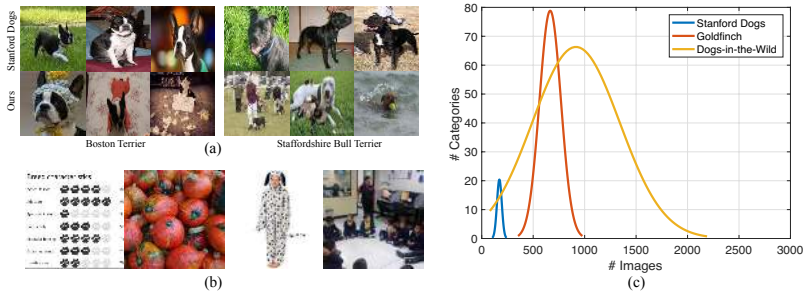


Fig. 5: Qualitative and quantitative comparison of dog datasets. (a) Example images from Stanford Dogs and Dogs-in-the-Wild; (b) Common bad cases from Goldfinch that are completely non-dog. (c) Images per category distribution.

5 Experimental Results

We conduct our experiments on four fine-grained image recognition datasets, including three publicly available datasets CUB-200-2011 [39], Stanford Dogs [16] and Stanford Cars [20], and the proposed Dogs-in-the-Wild dataset. The detailed statistics including class numbers and train/test distributions are summarized in Table 1. We adopt top-1 accuracy as the evaluation metric.

In our experiments, the input images are resized to 448×448 for both training and testing. We train on each dataset for 60 epochs; the batch size is set to 10 ($N=5$), and the base learning rate is set to 0.001, which decays by 0.96 for every 0.6 epoch. The reduction ratio r of \mathbf{W}_1^p and \mathbf{W}_2^p in Eq. 3 is set to 16 in reference to [13]. The weight parameter λ is empirically set to 0.5 as it achieves consistently good performances. And for the FC layers, we set the channels $C = 2048$ and $D = 1024$. Our method is implemented with Caffe [15] and one Tesla P40 GPU.

5.1 Ablation Analysis

To fully investigate our method, Table 2a provides a detailed ablation analysis on different configurations of the key components.

Base networks. To extract convolutional feature before the OSME module, we choose VGG-19 [36], ResNet-50 and ResNet-101 [11] as our candidate base-lines. Based on Table 2a, ResNet-50 and ResNet-101 are selected given their good

balance between performance and efficiency. We also note that although a better ResNet-50 baseline on CUB is reported in [21] (84.5%), it is implemented in Torch [5] and tuned with more advanced data augmentation (*e.g.*, color jittering, scaling). Our baselines, on the other hand, are trained with simple augmentation (*e.g.*, mirror and random cropping) and meet the Caffe baselines of other works, such as 82.0% in [26] and 78.4% in [7].

Importance of OSME. OSME is important in attending discriminative regions. For ResNet-50 without MAMC, using OSME solely with $P = 2$ can offer 3.2% performance improvement compared to the baseline (84.9% vs. 81.7%). With MAMC, using OSME boosts the accuracy by 0.5% than without OSME (using two independent FC layers instead, 86.2% vs. 85.7%). We also notice that two attention regions ($P = 2$) lead to promising results, while more attention regions ($P = 3$) provide slightly better performance.

MAMC constraints. Applying the first MAMC constraint (Eq. 6) achieves 0.5% better performance than the baseline with ResNet-50 and OSME. Using all of the three MAMC constraints (Eq. 6 to Eq. 8) leads to another 0.8% improvement. This indicates the effectiveness of each of the three MAMC constraints.

Complexity. Compared with the ResNet-50 baseline, our method provides significantly better result (+4.5%) with only 30% more time, while a similar method [10] offers less optimal result but takes $3.6\times$ more time than ours.

5.2 Comparison with State-of-the-Art

Quantitative experimental results are shown in Table 2b-2e.

We first analyze the results on the CUB-200-2011 dataset in Table 2b. It is observed that with ResNet-101, our method achieves the best overall performance (tied with MACNN) against state-of-the-art. Even with ResNet-50, our method exceeds the second best method using extra annotation (PN-CNN) by 0.8%, and exceeds the second best method without extra annotation (RAM) by 0.2%. For the weakly supervised methods without extra annotation, PDFR and MG-CNN conduct feature combination from multiple scales, and RACNN is trained with multiple alternative stages, while our method is trained with only one stage to obtain all the required features. Yet our method outperforms all of the the three methods by 2.0%, 4.8% and 1.2%, respectively. The methods B-CNN and RAN share similar multi-branch ideas with the OSME in our method, where B-CNN connects two CNN features with outer product, and RAN combines the trunk CNN feature with an additional attention mask. Our method, on the other hand, applies the OSME for multi-attention feature extraction in one step, which surpasses B-CNN and RAN by 2.4% and 3.7%, respectively.

Our method exhibits similar performances on Stanford Dogs and Stanford Cars, as shown in Table 2c and Table 2d. On Stanford Dogs, our method exceeds all of the comparison methods except RACNN, which requires multiple stages for feature extraction and is hard to be trained end-to-end. On Stanford Cars, our method obtains 93.0% accuracy, outperforming all of the comparison methods.

Finally, on the Dogs-in-the-Wild dataset, our method still achieves the best result with remarkable margins. Since this dataset is newly proposed, the re-

Table 2: Experimental results. ‘‘Anno.’’ stands for using extra annotation (bounding box or part) in training. ‘‘1-Stage’’ indicates whether the training can be done in one stage. ‘‘Acc.’’ denotes the top-1 accuracy in percentage.

Method	#Attention(P)	1-Stage	Acc.	Time(ms)
VGG-19	-	✓	79.0	79.8
ResNet-50	-	✓	81.7	48.6
ResNet-101	-	✓	82.5	82.7
ResNet-50 + OSME	2	✓	84.9	63.3
RACNN [10]	3	×	85.3	229
ResNet-50 + OSME + MAMC (Eq. 6)	2	✓	85.4	63.3
ResNet-50 + FC + MAMC (Eq. 6~8)	2	✓	85.7	60.3
ResNet-50 + OSME + MAMC (Eq. 6~8)	2	✓	86.2	63.3
ResNet-50 + OSME + MAMC (Eq. 6~8)	3	✓	86.3	68.1
ResNet-101 + OSME + MAMC (Eq. 6~8)	2	✓	86.5	102.1

(a) Ablation analysis of our method on CUB-200-2011.

Method	Anno.	1-Stage	Acc.
DVAN [52]	×	×	79.0
DeepLAC [22]	✓	✓	80.3
NAC [35]	×	✓	81.0
Part-RCNN [48]	✓	×	81.6
MG-CNN [40]	×	×	81.7
ResNet-50 [11]	×	✓	81.7
PA-CNN [18]	✓	✓	82.8
RAN [41]	×	×	82.8
MG-CNN [40]	✓	×	83.0
B-CNN [23]	×	×	84.1
ST-CNN [14]	×	×	84.1
FCAN [26]	×	✓	84.3
PDFR [51]	×	×	84.5
ResNet-101 [11]	×	✓	84.5
FCAN [26]	✓	✓	84.7
SPDA-CNN [47]	✓	✓	85.1
RACNN [10]	×	×	85.3
PN-CNN [2]	✓	×	85.4
RAM [21]	×	×	86.0
MACNN [53]	×	✓	86.5
Ours (ResNet-50)	×	✓	86.2
Ours (ResNet-101)	×	✓	86.5

(b) CUB-200-2011.

Method	Anno.	1-Stage	Acc.
PDFR [51]	×	×	72.0
ResNet-50 [11]	×	✓	81.1
DVAN [52]	×	×	81.5
RAN [41]	×	×	83.1
FCAN [26]	×	✓	84.2
ResNet-101 [11]	×	✓	84.9
RACNN [10]	×	×	87.3
Ours (ResNet-50)	×	✓	84.8
Ours (ResNet-101)	×	✓	85.2

(c) Stanford Dogs.

Method	Anno.	1-Stage	Acc.
DVAN [52]	×	×	87.1
FCAN [26]	×	✓	89.1
ResNet-50 [11]	×	✓	89.8
RAN [41]	×	×	91.0
B-CNN [23]	×	×	91.3
FCAN [26]	✓	✓	91.3
ResNet-101 [11]	×	✓	91.9
RACNN [10]	×	×	92.5
PA-CNN [18]	✓	✓	92.8
MACNN [53]	×	✓	92.8
Ours (ResNet-50)	×	✓	92.8
Ours (ResNet-101)	×	✓	93.0

(d) Stanford Cars.

Method	Anno.	1-Stage	Acc.
ResNet-50 [11]	×	✓	74.4
ResNet-101 [11]	×	✓	75.6
RAN [41]	×	×	75.7
RACNN [10]	×	×	76.5
Ours (ResNet-50)	×	✓	77.9
Ours (ResNet-101)	×	✓	78.5

(e) Dogs-in-the-Wild.

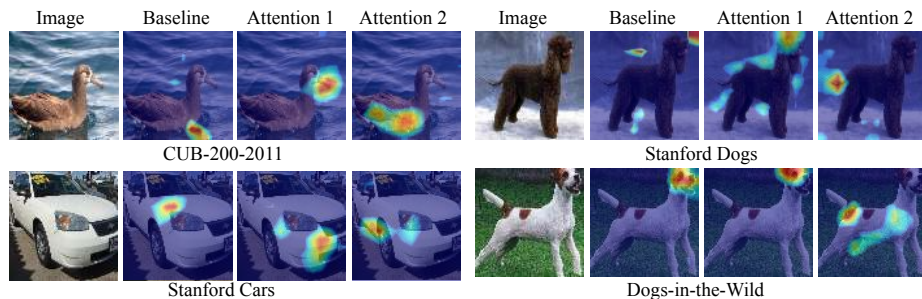


Fig. 6: Visualization of the attention regions detected by the OSME. For each dataset, the first column shows the input image, the second column shows the heatmap from the last conv layer of the baseline ResNet-101; the third and fourth columns show the heatmaps of the two detected attention regions via OSME.

sults in Table 2e can be used as baselines for future explorations. Moreover, by comparing the overall performances in Table 2c and Table 2e, we find that the accuracies on Dogs-in-the-wild are significantly lower than those on Stanford Dogs, which witness the relatively higher classification difficulty of this dataset.

By adopting our network with ResNet-101, we visualize the \mathbf{S}^p in Eq. 4 of each OSME branch (which corresponds to an attention region) as its channel-wise average heatmap, as shown in the third and fourth columns of Figure 6. In comparison, we also show the outputs of the last conv layer of the baseline network (ResNet-101) as heatmaps in the second column. It is seen that the highlighted regions of OSME outputs reveal more meaningful parts than those of the baseline, that we humans also rely on to recognize the fine-grained label, *e.g.*, the head and wing for birds, the head and tail for dogs, and the headlight/grill and frame for cars.

6 Conclusion

In this paper, we propose a novel CNN with the multi-attention multi-class constraint (MAMC) for fine-grained image recognition. Our network extracts attention-aware features through the one-squeeze multi-excitation (OSME) module, supervised by the MAMC loss that pulls positive features closer to the anchor, while pushing negative features away. Our method does not require bounding box or part annotation, and can be trained end-to-end in one stage. Extensive experiments against state-of-the-art methods exhibit the superior performances of our method on various fine-grained recognition tasks on birds, dogs and cars. In addition, we have collected and released the Dogs-in-the-Wild, a comprehensive dog species dataset with the largest data volume, full category coverage, and accurate annotation compared with existing similar datasets.

References

1. Bossard, L., Guillaumin, M., Gool, L.V.: Food-101 - mining discriminative components with random forests. In: ECCV (2014)
2. Branson, S., Van Horn, G., Belongie, S., Perona, P.: Bird species categorization using pose normalized deep convolutional nets. In: BMVC (2014)
3. Branson, S., Van Horn, G., Wah, C., Perona, P., Belongie, S.: The ignorant led by the blind: A hybrid human-machine vision system for fine-grained categorization. *Int. J. Comput. Vis.* **108**(1-2), 3–29 (2014)
4. Bromley, J., Guyon, I., LeCun, Y., Säckinger, E., Shah, R.: Signature verification using a “Siamese” time delay neural network. In: NIPS (1994)
5. Collobert, R., Kavukcuoglu, K., Farabet, C.: Torch7: A matlab-like environment for machine learning. In: BigLearn, NIPS workshop (2011)
6. Cui, Y., Zhou, F., Lin, Y., Belongie, S.: Fine-grained categorization and dataset bootstrapping using deep metric learning with humans in the loop. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1153–1162 (2016)
7. Cui, Y., Zhou, F., Wang, J., Liu, X., Lin, Y., Belongie, S.: Kernel pooling for convolutional neural networks. In: CVPR (2017)
8. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR (2009)
9. Farrell, R., Oza, O., Zhang, N., Morariu, V.I., Darrell, T., Davis, L.S.: Birdlets: Subordinate categorization using volumetric primitives and pose-normalized appearance. In: ICCV (2011)
10. Fu, J., Zheng, H., Mei, T.: Look closer to see better: recurrent attention convolutional neural network for fine-grained image recognition. In: CVPR (2017)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
12. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
13. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. arXiv preprint arXiv:1709.01507 (2017)
14. Jaderberg, M., Simonyan, K., Zisserman, A.: Spatial transformer networks. In: NIPS (2015)
15. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T.: Caffe: Convolutional architecture for fast feature embedding. In: ACM MM (2014)
16. Khosla, A., Jayadevaprakash, N., Yao, B., Li, F.F.: Novel dataset for fine-grained image categorization: Stanford dogs. In: CVPR Workshops on Fine-Grained Visual Categorization (2011)
17. Krause, J., Gebu, T., Deng, J., Li, L.J., Fei-Fei, L.: Learning features and parts for fine-grained recognition. In: ICPR (2014)
18. Krause, J., Jin, H., Yang, J., Fei-Fei, L.: Fine-grained recognition without part annotations. In: CVPR (2015)
19. Krause, J., Sapp, B., Howard, A., Zhou, H., Toshev, A., Duerig, T., Philbin, J., Fei-Fei, L.: The unreasonable effectiveness of noisy data for fine-grained recognition. In: ECCV (2016)
20. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3D object representations for fine-grained categorization. In: ICCV Workshops on 3D Representation and Recognition (2013)

21. Li, Z., Yang, Y., Liu, X., Wen, S., Xu, W.: Dynamic computational time for visual attention. arXiv preprint arXiv:1703.10332 (2017)
22. Lin, D., Shen, X., Lu, C., Jia, J.: Deep LAC: Deep localization, alignment and classification for fine-grained recognition. In: CVPR (2015)
23. Lin, T.Y., RoyChowdhury, A., Maji, S.: Bilinear CNN models for fine-grained visual recognition. In: ICCV (2015)
24. Lin, Y., Morariu, V.I., Hsu, W.H., Davis, L.S.: Jointly optimizing 3D model fitting and fine-grained classification. In: ECCV (2014)
25. Liu, J., Kanazawa, A., Jacobs, D.W., Belhumeur, P.N.: Dog breed classification using part localization. In: ECCV (2012)
26. Liu, X., Xia, T., Wang, J., Yang, Y., Zhou, F., Lin, Y.: Fully convolutional attention networks for fine-grained recognition. arXiv preprint arXiv:1603.06765 (2017)
27. Mnih, V., Heess, N., Graves, A., Kavukcuoglu, K.: Recurrent models of visual attention. In: NIPS (2014)
28. Nilsback, M., Zisserman, A.: Automated flower classification over a large number of classes. In: ICVGIP (2008)
29. Parkhi, O.M., Vedaldi, A., Jawahar, C., Zisserman, A.: The truth about cats and dogs. In: ICCV (2011)
30. Perronnin, F., Larlus, D.: Fisher vectors meet neural networks: A hybrid classification architecture. In: CVPR (2015)
31. Rosenfeld, A., Ullman, S.: Visual concept recognition and localization via iterative introspection. In: ACCV (2016)
32. Salakhutdinov, R., Hinton, G.E.: Learning a nonlinear embedding by preserving class neighbourhood structure. In: AISTATS (2007)
33. Schroff, F., Kalenichenko, D., Philbin, J.: FaceNet: A unified embedding for face recognition and clustering. In: CVPR (2015)
34. Simon, M., Gao, Y., Darrell, T., Denzler, J., Rodner, E.: Generalized orderless pooling performs implicit salient matching. arXiv preprint arXiv:1705.00487 (2017)
35. Simon, M., Rodner, E.: Neural activation constellations: Unsupervised part model discovery with convolutional networks. In: ICCV (2015)
36. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
37. Sohn, K.: Improved deep metric learning with multi-class n-pair loss objective. In: NIPS (2016)
38. Van Horn, G., Mac Aodha, O., Song, Y., Shepard, A., Adam, H., Perona, P., Belongie, S.: The iNaturalist challenge 2017 dataset. arXiv preprint arXiv:1707.06642 (2017)
39. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The Caltech-UCSD birds-200-2011 dataset. Tech. Rep. CNS-TR-2011-001, California Institute of Technology (2011)
40. Wang, D., Shen, Z., Shao, J., Zhang, W., Xue, X., Zhang, Z.: Multiple granularity descriptors for fine-grained categorization. In: ICCV (2015)
41. Wang, F., Jiang, M., Qian, C., Yang, S., Li, C., Zhang, H., Wang, X., Tang, X.: Residual attention network for image classification. In: CVPR (2017)
42. Wang, J., Zhou, F., Wen, S., Liu, X., Lin, Y.: Deep metric learning with angular loss. In: ICCV (2017)
43. Wang, J., Song, Y., Leung, T., Rosenberg, C., Wang, J., Philbin, J., Chen, B., Wu, Y.: Learning fine-grained image similarity with deep ranking. In: CVPR (2014)
44. Wang, Y., Choi, J., Morariu, V., Davis, L.S.: Mining discriminative triplets of patches for fine-grained classification. In: CVPR (2016)

45. Wilber, M., Kwak, I.S., Kriegman, D., Belongie, S.: Learning concept embeddings with combined human-machine expertise. In: ICCV (2015)
46. Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. In: ECCV (2014)
47. Zhang, H., Xu, T., Elhoseiny, M., Huang, X., Zhang, S., Elgammal, A., Metaxas, D.: SPDA-CNN: Unifying semantic part detection and abstraction for fine-grained recognition. In: CVPR (2016)
48. Zhang, N., Donahue, J., Girshick, R., Darrell, T.: Part-based R-CNNs for fine-grained category detection. In: ECCV (2014)
49. Zhang, N., Paluri, M., Ranzato, M., Darrell, T., Bourdev, L.: Panda: Pose aligned networks for deep attribute modeling. In: CVPR (2014)
50. Zhang, X., Zhou, F., Lin, Y., Zhang, S.: Embedding label structures for fine-grained feature representation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1114–1123 (2016)
51. Zhang, X., Xiong, H., Zhou, W., Lin, W., Tian, Q.: Picking deep filter responses for fine-grained image recognition. In: CVPR (2016)
52. Zhao, B., Wu, X., Feng, J., Peng, Q., Yan, S.: Diversified visual attention networks for fine-grained object classification. *IEEE Trans. Multimedia* **19**(6), 1245–1256 (2017)
53. Zheng, H., Fu, J., Mei, T., Luo, J.: Learning multi-attention convolutional neural network for fine-grained image recognition. In: ICCV (2017)
54. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Object detectors emerge in deep scene CNNs. In: ICLR (2014)
55. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: CVPR (2016)
56. Zhou, B., Lapedriza, A., Xiao, J., Torralba, A., Oliva, A.: Learning deep features for scene recognition using places database. In: NIPS (2014)
57. Zhou, F., Lin, Y.: Fine-grained image classification by exploring bipartite-graph labels. In: CVPR (2016)
58. Zhu, Y., Zhou, Y., Ye, Q., Qiu, Q., Jiao, J.: Soft proposal networks for weakly supervised object localization. In: ICCV (2017)