

Boosting Few-Shot Fine-Grained Recognition With Background Suppression and Foreground Alignment

Zican Zha¹, Hao Tang¹, Yunlian Sun¹, and Jinhui Tang¹, *Senior Member, IEEE*

Abstract—Few-shot fine-grained recognition (FS-FGR) aims to recognize novel fine-grained categories with the help of limited available samples. Undoubtedly, this task inherits the main challenges from both few-shot learning and fine-grained recognition. First, the lack of labeled samples makes the learned model easy to overfit. Second, it also suffers from high intra-class variance and low inter-class differences in the datasets. To address this challenging task, we propose a two-stage background suppression and foreground alignment framework, which is composed of a background activation suppression (BAS) module, a foreground object alignment (FOA) module, and a local-to-local (L2L) similarity metric. Specifically, the BAS is introduced to generate a foreground mask for localization to weaken background disturbance and enhance dominative foreground objects. The FOA then reconstructs the feature map of each support sample according to its correction to the query ones, which addresses the problem of misalignment between support-query image pairs. To enable the proposed method to have the ability to capture subtle differences in confused samples, we present a novel L2L similarity metric to further measure the local similarity between a pair of aligned spatial features in the embedding space. What's more, considering that background interference brings poor robustness, we infer the pairwise similarity of feature maps using both the raw image and the refined image. Extensive experiments conducted on multiple popular fine-grained benchmarks demonstrate that our method outperforms the existing state of the art by a large margin. The source codes are available at: <https://github.com/CSer-Tanghao/BSFA-FSFG>.

Index Terms—Few-shot learning, fine-grained recognition, background suppression, foreground alignment.

I. INTRODUCTION

FEW-SHOT learning (FSL) [1], [2], [3], [4], [5] has received widespread attention in the computer vision and multimedia fields because it mimics the ability of humans to learn new concepts with few available examples. Compared with the traditional classification paradigm, FSL does not depend on the large scale of labeled datasets and can be easily applied to many real-world scenarios where only very

Manuscript received 5 October 2022; revised 14 December 2022; accepted 2 January 2023. Date of publication 16 January 2023; date of current version 4 August 2023. This work was supported in part by the National Key Research and Development Program of China under Grant 2018AAA0102002, and in part by the National Natural Science Foundation of China under Grant 62076131. This article was recommended by Associate Editor F. Díaz-de-María. (Corresponding author: Yunlian Sun.)

The authors are with the School of Computer Science and Engineering, Nanjing University of Science and Technology, Nanjing 210094, China (e-mail: zhazican@njust.edu.cn; tanghao0918@njust.edu.cn; yunlian.sun@njust.edu.cn; jinhuitang@njust.edu.cn).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TCSVT.2023.3236636>.

Digital Object Identifier 10.1109/TCSVT.2023.3236636

1051-8215 © 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See <https://www.ieee.org/publications/rights/index.html> for more information.



Fig. 1. Key challenges of FS-FGR task. Here, rows represent different backgrounds and columns represent different species. Observing horizontally, different sub-categories have low *inter-class variance*. Observing longitudinally, every sub-category has high *intra-class variance*.

sparse training samples are available. Fine-grained recognition (FGR) [6], [7], [8], [9], [10], a popular and challenging problem, aims to recognize images of multiple sub-categories belonging to a super-category (*e.g.*, birds, dogs, cars). Considering the manual annotation for fine-grained images requires domain-specific knowledge, it is labor- and time-consuming to collect high quality and fully labeled large-scale datasets [11], [12], [13], thus FGR is a suitable application scenario of FSL. In this paper, we will study a more challenging and practical task, namely few-shot fine-grained recognition (FS-FGR), which aims to recognize fine-grained objects under few-shot settings.

Recently, to effectively learn from limited data, researchers have explored many general FSL algorithms for generic categories. In general, the main FSL methods can be roughly divided into two groups, *i.e.*, optimization-based methods [17], [18], [19], [20] and metric-learning based methods [1], [15], [21]. Optimization-based methods often adopt the “learning to learn” paradigm [22] to learn a meta-learner that can generate a robust model. In this way, the model can easily generalize to a new unseen task with a few training samples. The model of metric-learning based methods normally consists of two parts: a feature embedding module and a comparator. The feature embedding module maps images to an embedding space. Then the distance (*e.g.*, Euclidean distance and Cosine distance) between the support image and the query image is calculated

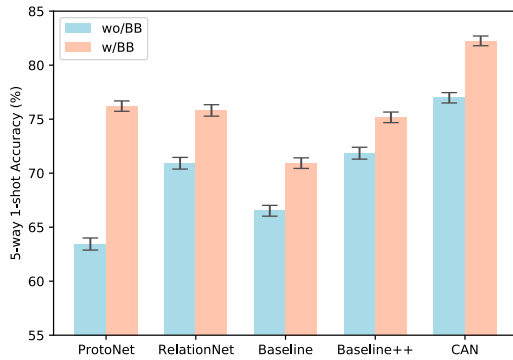


Fig. 2. Performance comparison of five models with and without human-annotated bounding boxes on CUB-200-2011 [14]: ProtoNet [1], RelationNet [15], Baseline [16], Baseline++ [16], CAN [2]. “w/ BB” denotes using the cropped images as input, where the raw images are pre-processed by human-annotated bounding boxes in advance.

directly in the embedding space. The comparator recognizes the query image based on the distance from each support image. Both methods have greatly promoted the performance of FSL on general datasets but few of them can obtain ideal performance on fine-grained datasets because of high intra-class variance and low inter-class difference.

In order to migrate the model from general datasets to fine-grained datasets, current FS-FGR models generally attempt to capture key distinctions, but can easily ignore the negative effects of background. As shown in Fig. 1, the habitats of birds are basically in the sky, land, sea, and grass. Although these birds are from different species, the same background makes them look very similar and difficult to classify. Therefore, removing the influence of the background (*e.g.*, by manually annotating bounding boxes) is significant to boost current few-shot methods. In order to show the performance boost brought by manually annotated bounding boxes, we did an experiment on CUB-200-2011 benchmark [14]. The results are reported in Fig. 2, which clearly show the significance of removing background influence. However, achieving better performance with human-annotated bounding boxes runs counter to the original aspiration of FSL to free people from burdensome and boring annotation tasks. Another key point that is easily overlooked in Fig. 1 is the large variation across viewpoints of the same fine-grained object. Semantic patterns of fine-grained objects are mainly determined by visual appearance. Thus, aligning semantically relevant local regions remains non-trivial for the FS-FGR task, which could effectively reduce the negative effects of feature inconsistency caused by the changes of object pose or viewing angle. Therefore, our goal is to remove cluttered backgrounds and align semantically relevant foregrounds when only image-level labels and limited training images are available.

In this paper, to tackle these problems with only image-level labels available. We propose a two-stage background suppression and foreground alignment framework for few-shot fine-grained recognition, which can be trained in an end-to-end manner. As shown in Fig. 3, the proposed method mainly consists of four modules: a feature extractor, a Background Activation Suppression (BAS) module, a Foreground Object

Alignment (FOA) module, and a Local to Local (L2L) similarity metric. Firstly, a feature extractor is used to extract the features for subsequent matching and localization learning. The BAS aims to generate a foreground mask map for localization based on the activation maps since the position in the activation map with a higher value are often where the interesting parts are located. Specifically, without adding extra trainable parameters, we only process the feature map to generate object location coordinates, and the BAS is supervised by the global classification loss. With the help of the generated bounding box information, we further obtained the finer scale of the object image by cropping and zooming in to remove the cluttered background. Different from the conventional method, we incorporate both the original images and the refined ones obtained by BAS into the model for learning, which can effectively generate fine-grained tailored representations for few-shot recognition. Secondly, following DN4 [23], we use local descriptors to avoid losing spatial information and subtle features. Thus, to alleviate the misalignment between the support image and the query image in the embedding space due to variations in the scale, position, and posture of the foreground objects in the image, an object-aware FOA is introduced to transform support features with respect to query features by calculating a semantic correlation matrix. Thirdly, we propose an L2L similarity metric to measure the local similarity between the aligned spatial features of a given pair of samples. Experiments show that the best results are obtained by combining the similarity of raw and refined support-query pairs. Extensive experiments on three common benchmarks (*i.e.*, CUB-200-2011 [14], StanfordDogs [24], and StanfordCars [25]) confirm the effectiveness of the proposed method. Our main contributions can be summarized as follows:

- We empirically reveal two crucial points to significantly improve the performance of the FS-FGR task, *i.e.*, weakening background disturbance and aligning foreground response.
- We develop a novel two-stage weakly-supervised framework to address the FS-FGR task in a meta-learning fashion. The proposed paradigm can be optimized in an end-to-end manner with only the image-level label available.
- Three well-designed modules are instantiated in our framework to achieve the background suppression and foreground alignment, *i.e.*, background activation suppression module, foreground object alignment module, and local-to-local similarity metric.
- We conduct comprehensive experiments on three popular fine-grained datasets, and our proposed model achieves state-of-the-art performance. On basis of the empirical evaluation of these datasets, we also provide insights for the elegant implementation of our solution.

The remainder of this paper is organized as follows. Sect. II introduces the related work, including Few-shot Learning, Fine-grained Recognition, and Few-shot Fine-grained Recognition. Sect. III gives the definition of the problem, the proposed framework, and the details of each module. Sect. IV presents the introduction of the datasets, the experimental

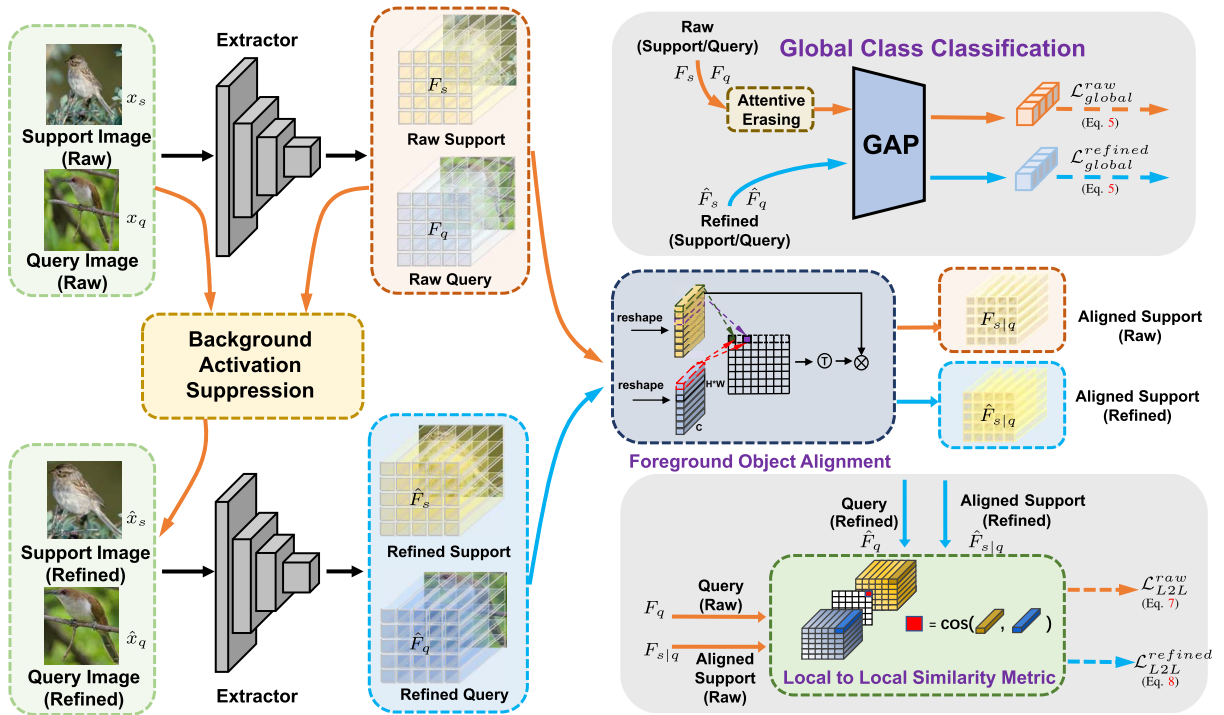


Fig. 3. The framework of our proposed method. Note that the red line is the process of the raw stage, and the blue line is the process of the refined stage. Note that the input of the refined stage is generated by the background activation suppression module based on the input of the raw stage. Both stages have the same structure and all parameters are shared except for the global classifiers. Each stage consists of a foreground object alignment (FOA) module, a global classifier, and a similarity comparator based on L2L. In the inference process, we calculate the similarity between the query-support aligned image pair in both the raw stage and the refined stage, and the similarities of both image pairs are integrated to support the final decision.

results and analysis in detail, and some visualizations. The paper is concluded in Sect. V.

II. RELATED WORK

A. Few-Shot Learning

Few-shot learning (FSL), which aims to learn a well-generalized model with only a few samples, has attracted intensive attention in recent years. As one of the standard benchmarks in meta-learning, existing work on FSL can be roughly categorized into three groups according to their innovation: metric-based methods, optimization-based methods, and augmentation-based methods. Metric-based methods [1], [15], [21] aim to learn a suitable model that projects all images to a metric space, and then predicts query images according to their distances to the labeled support samples. Optimization-based methods [17], [18], [26] follow the “learning to learn” paradigm to learn a good initialization for the model or learn how to update the parameters of the model within a few iterations for quick adaptation. Augmentation-based methods [27], [28], [29] aim to learn a generator from base categories and use it with additional techniques to increase the number of novel samples or features for data augmentation. In this work, we are particularly interested in the metric-based methods owing to their intuitive and effective characteristics, which are favored by state-of-the-art methods on many tasks.

Metric-based methods mainly focus on two key factors, *i.e.*, *feature extraction* [3], [23] and *similarity measurement*

[2], [30], [31]. The former is responsible for mapping the samples to an ideal embedding space, where the samples belonging to the same category are close together and the samples belonging to different categories are far apart. The latter serves as a criterion for the nearest neighbor search to perform recognition. However, most of the existing metric-based FSL methods [3], [32], [33], [34] mainly focus on learning a discriminative global embedding for coarse-grained generic object recognition, which are less suitable to address the few-shot fine-grained recognition (FS-FGR) task well. The main reason is that fine-grained recognition requires to emphasis on distinguished appearance details of local parts, but global features are incapable of mining spatially local discriminative parts, especially under the few-shot scenario. By comparison, the proposed method proposes a foreground alignment module to emphasize the semantically correlated parts of the input support-query pair via spatially aligning dense local features for FS-FGR.

B. Fine-Grained Recognition

Fine-grained recognition [35] aims to recognize multiple subordinate categories belonging to the same super-category (*e.g.*, bird species [14], car models [24] and dog breeds [25]), which has gained much attraction of the research community. Compared to the coarse-grained recognition [36], [37], [38], fine-grained recognition is a more challenging task, since fine-grained objects are usually distinguished by local feature variations or subtle feature differences.

Earlier work [39], [40] mainly relies on part annotations or hand-crafted bounding boxes to locate discriminative specific parts, which exacerbates the cost of prior information or additional annotations. Benefiting from the remarkable progress of powerful deep neural networks and large-scale annotated datasets, some deep learning-based approaches [8], [9], [41] attempt to learn discriminative features or locate the distinguishable parts in a weakly supervised manner, where only image-level category labels are available. These methods can be roughly categorized into two main groups: *feature encoding*-based methods [6], [42], [43] and *part localization*-based methods [7], [44], [45]. For example, Guo et al. [46] introduce a lightweight attention module to localize critical regions and learn fine-grained feature representation. More recently, MMAL [9] and AP-CNN [41] propose to first find strong discriminative regions, and then re-input the original image or feature maps to the network by cropping and resizing them, which further enhance the object representation's discrimination. The above novel paradigm also serves as some inspiration for our proposed method.

Despite the success brought by the above methods, they still rely on large-scale datasets, which are less practical in real scenarios, due to the large-scale annotated datasets being hard to attain in some cases. Compared with previous work, we study fine-grained recognition under the challenging few-shot learning setting, namely few-shot fine-grained recognition (FS-FGR), where only a few labeled samples are provided to recognize novel fine-grained categories.

C. Few-Shot Fine-Grained Recognition

Recently, inspired by the rapid development of FSL, there are some emerging work [47], [48], [49] that starts to explore the intractable few-shot fine-grained recognition (FS-FGR) task, where the goal is to distinguish the novel sub-categories in a super-category given a limited number of labeled samples. Wei et al. [47] propose the first solution for FS-FGR by introducing a bilinear pooling network to encode features and multiple sub-classifiers to predict the label. LRPABN [50] proposes to learn multiple transformations to match the embedded input pairs. MattML [20] introduces multi-attention mechanisms to learn a task-specified classifier initialization, which is helpful to capture diverse discriminative parts. However, the above methods tackle the FS-FGR task by learning and matching diverse and informative global features, which follow the same perspective as the generic FSL methods. Limited efforts have been devoted to distinguishing the subtle difference from the perspective of dense local features. In contrast, CPSN [51] introduces two coupled branches to compute the similarity scores between input pairs from patch level to capture subtle and local differences.

More recently, considering object location differences in the support and query set cause performance degradation, some works have introduced feature alignment to solve the FS-FGR problem, which has achieved very inspiring results. TOAN [52] proposes a target-oriented matching mechanism to learn explicit feature transformations to reduce the intra-class variance. Wu et al. [53] introduce a dual correlation

attention mechanism to explore position-wise relationships of two compared features to capture their semantic information. Wu et al. [10] propose an object-aware long-short-range spatial alignment method to align features between various support and query features. Although the above cross-correction attention-based workflow seems to make sense, they are all pretty sophisticated and obscure.

We can see that existing FS-FGR methods usually focus on how to extract discriminative features and how to strengthen the relation matching process of input pairs by different metric functions. However, inherent characteristics (*e.g.*, pose variations and similar backgrounds) of fine-grained images are also important factors for the FS-FGR task that are often ignored by existing FS-FGR methods. In this paper, motivated by aforementioned work in the FS-FGR community, we argue that **background suppression** and **foreground alignment** are essential to recognize fine-grained objects, and they have been demonstrated to be very appropriate for the FS-FGR task and could significantly improve the performance.

We are aware that recently AGPF [49] proposed an attention-guided refinement strategy to enhance the dominative object and conducted a two-stage meta-learning framework to capture attention-guided pyramidal features. Although AGPF also adopts a two-stage framework to weaken the background disturbance, our approach obviously differs from AGPF in these aspects, at least: (1) Our method encodes each input into a set of dense local features and introduces a feature alignment technique for fine-grained objects, which is more suitable for capturing subtle appearance differences of local parts than a single informative global feature used in AGPF. (2) The attention-guided refinement proposed in AGPF utilizes sophisticated attention mechanisms to locate the dominative foreground. However, ours only relies on the feature activation map to predict and disentangle the foreground object and background noise, which is parameter-free. (3) AGPF aggregates complementary features discovered from different scales and stages to form a unified global embedding, but our method computes the spatial similarity between aligned feature pairs in two stages to jointly support the final decision. (4) Compared with AGPF, our results are significantly better and have significant advantages in efficiency.

III. METHODOLOGY

A. Problem Formulation

In the standard setting of few-shot learning, given a large-scale labeled dataset with base categories C_{base} , the goal is to learn a standard embedding network that can be easily adapted to a family of unseen tasks $\{\mathcal{T}\}_1^n$ with novel categories C_{novel} , where $C_{base} \cap C_{novel} = \emptyset$. Considering a specific image recognition task \mathcal{T} , it usually contains a *support set* \mathbf{S} and a *query set* \mathbf{Q} where both are sampled from the same novel categories C_{novel} , our goal is to determine the category of samples in the unlabeled \mathbf{Q} according to the labeled \mathbf{S} . In this work, we study a “ N -way K -shot” task, where the support set \mathbf{S} contains N categories and each category has K images. Concretely, N is set to 5 and K is set to 1 or 5 in most previous studies, and so do we. To have a good generalization on novel

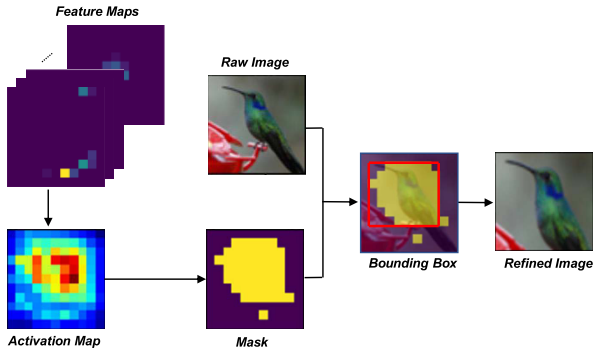


Fig. 4. Background Activation Suppression (BAS) Module. For the input feature, BAS first generates a class-agnostic activation map for localization. The foreground prediction map is then obtained by suppressing the background activation value.

categories \mathcal{C}_{novel} and avoid the overfitting on base categories \mathcal{C}_{base} , we also adopt the episodic training mechanism [21] widely used in most literature.

B. The Proposed Framework

Given a “ N -way K -shot” fine-grained recognition task, we need to divide it into a support set $\mathbf{S} = \{(x_i^j, y_i^j) \mid i = 1 \dots N, j = 1 \dots K\}$ and a query set $\mathbf{Q} = \{(x_i, y_i) \mid i = 1 \dots |\mathbf{Q}|\}$, and then map all samples in a task into an embedding space by a feature extractor Θ , *i.e.*, $F_i = \Theta(x_i) \in \mathbb{R}^{c \times h \times w}$, where $x_i \in \mathbf{S} \cup \mathbf{Q}$. Here, c is the number of channels, h and w denote the size of the feature map. Note that the proposed method borrows the idea of metric learning and needs to match each query sample to support prototypes by the nearest neighbor strategy. Therefore, the i_{th} support prototype \mathbf{C}_i can be formulated as $\mathbf{C}_i = \frac{1}{K} \sum_{j=1}^K \Theta(x_i^j)$, where x_i^j is a sample labeled with category $i \in \{1, \dots, N\}$.

The overview of the proposed method is shown in Fig. 3. Given one support sample $x_s \in \mathbf{S}$ and one query sample $x_q \in \mathbf{Q}$, we first extract the corresponding feature maps F_s and F_q , respectively. Next, the *background activation suppression* module is used to remove the cluttered background based on the activation value of the feature map and then zoom in the cropped region to the size of the raw image. \hat{x}_s and \hat{x}_q represent the refined support image and the refined query image, respectively. We extract the refined feature maps \hat{F}_s and \hat{F}_q from the same backbone Θ . Here, the feature extractors in the upper and lower branches are parameter-shared. Subsequent steps first use the proposed *foreground object alignment* module to align two semantic feature pairs (*i.e.*, $\{F_s, F_q\}$ and $\{\hat{F}_s, \hat{F}_q\}$) spatially, and then compute the total spatial similarity between two aligned pairs by introducing the proposed “*local to local*” similarity metric. Finally, we integrate the similarity scores of both pairs to classify the novel query set \mathbf{Q} .

C. Background Activation Suppression Module

1) *Motivation*: As we all know, traditional fine-grained recognition is a challenging problem due to the small inter-category variations and the large intra-category variations that exist in the fine-grained dataset. As shown in Fig. 1, the

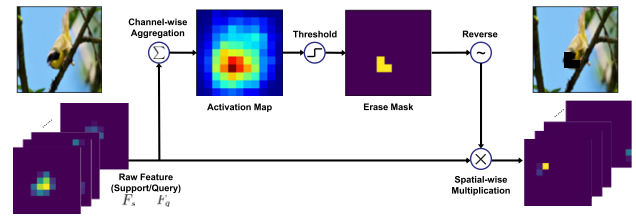


Fig. 5. Attentive erasing diagram. In the raw stage, to encourage the network to minimize the ratio of background activation and overall activation generated by the entire image, we erase the region of raw feature with activation higher than a threshold before being fed into the global classifier, which is helpful to explore the entire extent of the foreground.

visual patterns of different sub-categories are similar to some extent, and the backgrounds in the same sub-categories may be different but appear slightly similar among different sub-categories. Therefore, the image background plays a negative role in the fine-grained recognition task, especially in the few-shot scenario. Inspired by the weakly-supervised localization methods [7], [9] for the fine-grained task, the proposed method attempts to eliminate the negative effect of background by the BAS module, without additional annotations except for image-level annotation. BAS directly generates class-agnostic activation maps to disentangle the foreground object and background in an image.

For example, given an input query image x_q , the corresponding feature map $F_q = \Theta(x_q) \in \mathbb{R}^{c \times h \times w}$ is generated by the feature extractor. Note that each channel of a feature map can be regarded as a pattern detector [54], denoted by f_i ($i = 1, \dots, c$). If a certain spatial position in the feature map has high activation for most channels, it is likely to correspond to an informative region. Therefore, we obtain the activation map $A_F \in \mathbb{R}^{h \times w}$ by aggregating F along the channel dimension as $A_F = \sum_{i=1}^c f_i$. Then, we introduce an adaptive threshold θ_{A_F} for the activation map A_F to determine which position is a part of the object regions, *i.e.*, $\theta_{A_F} = \frac{\sum_{i=1}^w \sum_{j=1}^h A_F(i, j)}{h \times w}$. Finally, we generate a foreground mask $M_{A_F} \in \{0, 1\}^{h \times w}$ by comparing each element of A_F with threshold θ_{A_F} . Concretely, for a particular position (i, j) , if the activation value of $A_F(i, j)$ is larger than the threshold θ_{A_F} , the corresponding $M_{A_F}(i, j)$ is set to 1 otherwise 0. Formulaically,

$$M_{A_F}(i, j) = \begin{cases} 1, & \text{if } A_F(i, j) > \theta_{A_F} \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Here, we search the largest connected component of the generated M_{A_F} to obtain the smallest bounding box for the foreground prediction. Based on the location coordinates, we can perform background suppression by cropping interesting foreground on raw image x_q and enlarging it to the same size as x_q for subsequent classification, denoted as \hat{x}_q . The processing steps of the BAS module are illustrated in Fig. 4.

Due to the unbalanced distribution between the foreground features and background features, fine-grained recognition models incline to identify patterns from small discriminative object regions. To address this issue, attentive erasing strategy [55], [56] as an efficient solution has been widely used in

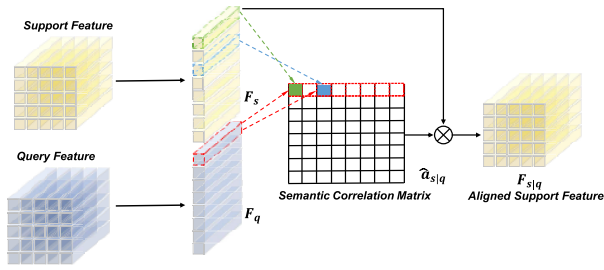


Fig. 6. Foreground Object Alignment (FOA) Module. The input support-query image pair is initially misaligned, due to the arbitrary pose and position variations of foreground between the two fine-grained images. By leveraging the proposed FOA module, semantic features of the original support image could be aligned *w.r.t* the query image.

weakly supervised object localization. However, these methods have limited potential to derive complete foreground activation maps of unseen objects. To further improve the completeness of the foreground prediction map with image-level labels, we propose to promote the learning of BAS by erasing the prominent regions of the feature map with activation higher than a threshold as shown in Fig. 5. Firstly, we generate an activation map $M_{att} \in \mathbb{R}^{h \times w}$ as same as A_F in the BAS module by aggregating raw feature (*i.e.* F_s or F_q) along the channel dimension. Then, the spatial distribution could be searched for the most discriminative part based on M_{att} , due to the intensity of each position being proportional to the discriminative power. To obtain the erase mask $M_{era} \in \mathbb{R}^{h \times w}$, a threshold θ_{M_F} is introduced by setting a ratio γ of maximum intensity of M_{att} . Thus, M_{era} is produced by setting each pixel to 1 if it is larger than θ_{M_F} and 0 if it is smaller, and the size of the erased region decreases as γ increases. For the erasing threshold, we set hyper-parameter γ to 85% in our recommended settings. Finally, the mask M_{era} is applied to the input feature map by inversion and spatial-wise multiplication. This attentive erasing strategy facilitates the generation of a meaningful class-agnostic activation map to explore the entire foreground object by indirectly increasing the foreground activation value. This also encourages the model to learn the less discriminative parts.

D. Foreground Object Alignment Module

1) *Motivation*: Although global embedding generated by global average pooling or global max pooling operations is capable of capturing discriminative information in general few-shot learning tasks, it might suppress some local discriminative characteristics while overemphasizing the irrelevant background features for the fine-grained object [57], hurting the FG-FSR performance. As illustrated in Fig. 1, we can see that capturing the subtle differences in discriminative parts remains non-trivial to distinguish similar sub-categories. To this end, our proposed method attempts to exploit dense local features to calculate the similarity between a given pair of samples, which can further capture spatial structure information of the foreground for FG-FSR. In addition, the lower intra-category visual consistency of fine-grained objects usually causes different similarities in both global appearance

and various local regions. Therefore, an ideal solution to FS-FGR should not only be sensitive to the subtle discrepancy among instances from different sub-categories but also be invariant to the arbitrary poses, scales, and appearances of the same subcategory. In light of this, our proposed method goes two steps further to compute the similarity of each support-query feature pair: (1) aligning semantic features between a query image and various support images according to their correlations; (2) modeling the similarity measurement between each aligned support-query pair as a local-to-local (patch-to-patch) matching process.

Specifically, before locally comparing a support feature $F_s \in \mathbb{R}^{c \times h \times w}$ and query feature $F_q \in \mathbb{R}^{c \times h \times w}$, we first need to address the issue of semantic misalignment, due to intra-class variance (*e.g.*, object posture and position variation) and cluttered background. In detail, we first reshape a pair of spatial features from $\mathbb{R}^{c \times h \times w}$ to $\mathbb{R}^{hw \times c}$, *i.e.*, $F_s = \{s_i\}_{i=1}^{h \times w}$ and $F_q = \{q_i\}_{i=1}^{h \times w}$, where s_i and $q_i \in \mathbb{R}^{c \times 1 \times 1}$ represent the i th deep local descriptor in F_s and F_q , respectively. Then, we use the reshaped F_s and F_q to compute a semantic correlation matrix $a_{s|q} \in \mathbb{R}^{hw \times hw}$, which is calculated by the cosine similarity as $a_{s|q}(i, j) = \frac{q_i^T s_j}{\|q_i\| \|s_j\|}$. Furthermore, a normalized operation $\text{softmax}(\cdot)$ is applied to $a_{s|q}$ row-by-row, so that the sum of each row in the normalized $a_{s|q}$ is 1. The resulting semantic correlation matrix is formulated as follows:

$$\hat{a}_{s|q} = \text{softmax}(a_{s|q}) = \frac{\exp(a_{s|q}(i, j))}{\sum_{m=1}^{hw} \exp(a_{s|q}(i, m))}. \quad (2)$$

As shown in Fig. 6, we employ the FOA module to compute the semantic correlation matrix $\hat{a}_{s|q}$, which is used to align the support feature F_s with respect to query feature F_q . Formulaically, the aligned feature $F_{s|q}$ can be computed by an alignment function $\mathcal{A}(\cdot, \cdot)$ as:

$$F_{s|q} = \mathcal{A}(F_s, F_q) = \hat{a}_{s|q} \cdot F_s. \quad (3)$$

Considering that the semantic features aligned by the FOA module are no longer affected by the position deviation, we further propose a patch-level similarity metric, indicated by $L2L(\cdot, \cdot)$, which reformulates similarity measurement as a local-to-local matching process. As shown in Fig. 7, given a pair of aligned semantic features $\{F_{s|q}, F_q\} \in \mathbb{R}^{hw \times c}$, the $L2L$ metric can be formulated as a general function of $F_{s|q}$ and F_q :

$$L2L(F_{s|q}, F_q) = \sum_{i=1}^h \sum_{j=1}^w \cos(F_{s|q}^{i,j}, F_q^{i,j}), \quad (4)$$

where $F^{i,j} \in \mathbb{R}^c$ is a deep descriptor of position (i, j) in the semantic feature F and $\cos(\cdot, \cdot)$ denotes the cosine similarity function.

E. Overall Loss Function

During the training phase, given a “ N -way K -shot” fine-grained recognition task \mathcal{T} , each sample x_i in \mathcal{T} has two types of labels: $y^g \in \{1, 2, \dots, G\}$ and $y^l \in \{1, 2, \dots, N\}$, where G is the overall number of base categories. Firstly, to learn a class-agnostic activation map for achieving foreground prediction and background suppression, we use cross-entropy loss as the classification loss to explicitly explore

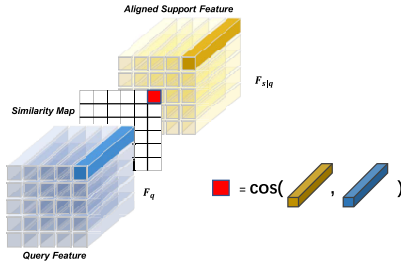


Fig. 7. Local to Local (L2L) Similarity Metric. We compute one similarity map by measuring the cosine similarity element-wisely between the two aligned feature maps of compared images.

rich inter-class relationships by learning features related to the whole base category space. As shown in Fig. 3, there are two independent global classifiers in our proposed framework to classify features of the raw image x_i and the refined image \hat{x}_i to its corresponding base category y^g . The corresponding global classification loss $\mathcal{L}_{global}^{raw}$ and $\mathcal{L}_{global}^{refined}$ are formulated as follows:

$$\begin{aligned}\mathcal{L}_{global}^{raw} &= - \sum_{(x_i, y^g) \in \mathcal{T}} y^g \log(\text{softmax}(W_1 \text{GAP}(\Theta(x_i))))), \\ \mathcal{L}_{global}^{refined} &= - \sum_{(\hat{x}_i, y^g) \in \mathcal{T}} y^g \log(\text{softmax}(W_2 \text{GAP}(\Theta(\hat{x}_i)))).\end{aligned}\quad (5)$$

Here, W_1 , W_2 are the weight parameters of the global classifiers, $\text{GAP}(\cdot)$ denotes the global average pooling, y^g is the true base category label for x^{raw} and $x^{refined}$. Finally, the global classification loss \mathcal{L}_{global} can be formulated as follow:

$$\mathcal{L}_{global} = \alpha \mathcal{L}_{global}^{raw} + \beta \mathcal{L}_{global}^{refined}, \quad (6)$$

where α and β are two weighting factors. Note that the global classification branch is only a part of the whole framework and will be activated in the training phase to constrain the generation of class-agnostic activation maps.

Our goal is to identify each query sample with a few labeled samples in the support set. We compute the similarity between the query and each support prototype and then the matching process is guided by the nearest-neighbor strategy. For task \mathcal{T} with N categories, given a prototype set $\{\mathbf{C}_j\}_{j=1}^N$ constructed from raw support images, the local-level few-shot loss between the feature F_i of raw query image $x_i \in \mathbf{Q}$ (i.e., $F_i = \Theta(x_i)$) and the corresponding prototype \mathbf{C}_j can be represented by

$$\mathcal{L}_{L2L}^{raw} = - \sum_{i=1}^{|\mathbf{Q}|} \log \frac{\exp(-L2L(\mathcal{A}(\mathbf{C}_j, F_i), F_i))}{\sum_{j'=1}^N \exp(-L2L(\mathcal{A}(\mathbf{C}_{j'}, F_i), F_i))}. \quad (7)$$

Similar to \mathcal{L}_{global} in Eq. 5, we replace the raw query image x_i in Eq. 7 with the refined query image \hat{x}_i to obtain $\mathcal{L}_{L2L}^{refined}$ as follow:

$$\mathcal{L}_{L2L}^{refined} = - \sum_{i=1}^{|\mathbf{Q}|} \log \frac{\exp(-L2L(\mathcal{A}(\hat{\mathbf{C}}_j, \hat{F}_i), \hat{F}_i))}{\sum_{j'=1}^N \exp(-L2L(\mathcal{A}(\hat{\mathbf{C}}_{j'}, \hat{F}_i), \hat{F}_i))}, \quad (8)$$

where \hat{F}_i indicates the feature of refined query image \hat{x}_i generated by BAS, and $\hat{\mathbf{C}}_{j'}$ is the corresponding prototype computed by refined support images. Therefore, the final local-level few-shot loss \mathcal{L}_{local} can be represented as follow:

$$\mathcal{L}_{local} = \alpha \mathcal{L}_{L2L}^{raw} + \beta \mathcal{L}_{L2L}^{refined}, \quad (9)$$

where α and β are two weighting factors as same as hyper-parameters in Eq. 6 to control the contributions of two stages. Here, two scale factors α and β are both set to 0.5 for all experiments, while optimizing them can further improve the performance.

In summary, the cooperative loss is formulated as follows:

$$\mathcal{L}_{total} = \mathcal{L}_{global} + \lambda \mathcal{L}_{local}, \quad (10)$$

where λ is a weighting factor. In the inference phase, only the branch of local-level few-shot loss \mathcal{L}_{local} is used to compute the similarity to recognize the query samples.

IV. EXPERIMENTS

A. Datasets

We conduct experiments on three widely used benchmarks for FG-FSR, including CUB-200-2011 [14], Stanford Dogs [24], and Stanford Cars [25]. We follow the same evaluation protocol as the previous work [10], [23], [49], and all the input images are resized to 84×84 .

- **CUB-200-2011** contains 200 species of birds with 11, 788 images in total [14]. Following the evaluation protocol of [10], we randomly split the dataset into 100 base categories for training, 50 categories for validation, and 50 novel categories for evaluation.
- **Stanford Dogs** contains 20, 580 images from 120 subclasses of dogs in total [24]. Following the evaluation protocol of [23], we randomly split the dataset into 70 base categories for training, 20 categories for validation, and 30 novel categories for evaluation.
- **Stanford Cars** contains 16, 185 images spanning 196 car models in total [25]. Following the evaluation protocol of [23], we randomly split the dataset into 130 base categories for training, 17 categories for validation, and 49 novel categories for evaluation.

Furthermore, considering that prior work [30], [32], [58], [59] on CUB-200-2011 pre-processes images with human-annotated bounding box as the input, we also conduct experiments under this setting for a fair comparison.

B. Implementation Details

To have a fair and broad comparison with the previous work [10], [30], [48], [52], we adopt the widely-used backbone (i.e., ResNet-12) as the feature extractor. ResNet-12 consists of 4 consecutive basic blocks and the number of filters is set to 64-128-256-512. To achieve more accurate background suppression and foreground alignment, we remove the last pooling layer of ResNet-12, thus the size of the output feature map is $512 \times 11 \times 11$. Without bells and whistles, no additional parameters are introduced in our model except two fully connected (FC) layers are used as global classifiers.

TABLE I

THE ACCURACY (%) OF 5-WAY 1-SHOT AND 5-SHOT TASKS ON THREE POPULAR BENCHMARKS. WE REPORT THE MEAN ACCURACY WITH 95% CONFIDENCE INTERVALS ON 2000 TEST EPISODES. THE BEST RESULTS ARE SHOWN IN BOLD. * DENOTES THE RESULTS OBTAINED FROM THE ORIGINAL PAPER. WE REPRODUCE THE OTHER RESULTS UNDER THE SAME EXPERIMENTAL SETTINGS ACCORDING TO THE OPEN SOURCE CODE

Method	CUB-200-2011		Stanford Dogs		Stanford Cars	
	5-way 1-shot	5-way 5-shot	5-way 1-shot	5-way 5-shot	5-way 1-shot	5-way 5-shot
ProtoNet [1]	63.44 ± 0.56	83.17 ± 0.35	41.61 ± 0.50	76.78 ± 0.36	45.01 ± 0.49	87.19 ± 0.31
RelationNet [15]	70.92 ± 0.54	84.90 ± 0.35	61.21 ± 0.51	80.27 ± 0.37	78.04 ± 0.53	90.03 ± 0.30
DN4 [23]	64.95 ± 0.99	83.18 ± 0.62	49.70 ± 0.85	71.59 ± 0.68	75.79 ± 0.84	94.14 ± 0.35
Baseline [16]	66.52 ± 0.50	83.31 ± 0.34	59.19 ± 0.48	78.21 ± 0.34	69.48 ± 0.46	87.81 ± 0.27
Baseline++ [16]	71.85 ± 0.55	82.14 ± 0.62	58.53 ± 0.51	73.12 ± 0.38	76.79 ± 0.53	87.18 ± 0.32
CAN [2]	76.98 ± 0.48	87.77 ± 0.30	64.73 ± 0.52	77.93 ± 0.35	86.90 ± 0.42	93.93 ± 0.22
TOAN* [52]	66.10 ± 0.86	82.27 ± 0.60	49.77 ± 0.86	69.29 ± 0.70	75.28 ± 0.72	87.45 ± 0.48
BSNet* [48]	73.48 ± 0.92	83.84 ± 0.59	61.95 ± 0.97	79.62 ± 0.63	71.07 ± 1.03	88.38 ± 0.62
OLSA* [10]	77.77 ± 0.44	89.87 ± 0.24	64.15 ± 0.49	78.28 ± 0.32	77.03 ± 0.46	88.85 ± 0.46
AGPF [49]	78.73 ± 0.84	89.77 ± 0.47	72.34 ± 0.86	84.02 ± 0.57	85.34 ± 0.74	94.79 ± 0.35
Ours	82.27 ± 0.46	90.76 ± 0.26	69.58 ± 0.50	82.59 ± 0.33	88.93 ± 0.38	95.20 ± 0.20

During the training stage, standard data augmentation strategies are adopted, *e.g.*, random cropping and random horizontal flipping. For all experiments, we utilize SGD with an initial learning rate of 0.1 as the optimizer to train our model on the Titan RTX GPU for 90 epochs. The learning rate decays to 0.06 in the 60th epoch and then times 0.2 every 10 epochs. The λ in the overall loss is experimentally set to 0.1. Significantly, our proposed framework is trained from scratch in an end-to-end manner, and there is no need to pre-train the feature extractor on other datasets. During the testing stage, evaluation is performed under standard ‘5-way 1-shot’ and ‘5-way 5-shot’ settings, we test 2,000 episodes and report the mean accuracy with 95% confidence intervals as the final result. Our experiments are implemented in PyTorch.

C. Comparison With State-of-the-Art Methods

1) *Results on General Fine-Grained Datasets:* Tab. I shows the performance evaluations on three aforementioned fine-grained benchmark datasets. We demonstrate ten representative methods that have reported evaluation results on the corresponding datasets, including six typical FSL methods (*i.e.*, ProtoNet [1], RelationNet [15], DN4 [23], Baseline [16], Baseline++ [16], and CAN [2]) and four specialized FS-FGR methods (*i.e.*, TOAN [52], BSNet [48], OLSA [10], and AGPF [49]). For a fair comparison with the above methods, we display the results of our proposed method based on the ResNet-12 backbone. Compared with the above typical FSL and FS-FGR methods, our proposed method achieves the highest performance under both ‘5-way 1-shot’ and ‘5-way 5-shot’ settings on the most benchmark datasets. The observations can be summarized as follows:

- On the CUB-200-2011 dataset, our proposed method achieves stable and excellent performance and surpasses all existing methods by a large margin in terms of ‘5-way 1-shot’ setting. It’s clear that our method exceeds

six typical FSL methods (*i.e.* [1], [2], [15], [16], [23]) significantly. Compared with the global feature based FS-FGR method TOAN [52], our proposed method achieves 16.17% and 8.49% performance improvement under 1-shot and 5-shot settings, respectively. In addition, the current state-of-the-art accuracy of 1-shot is achieved by AGPF [49] with 78.73%, which shares the same insights with our method. Our method outperforms AGPF by 3.54% with an accuracy of 82.27%.

- On the Stanford Dogs dataset, our method overall exceeds all typical FSL methods significantly in terms of 1-shot and 5-shot settings. But compared with AGPF [49], our method is lower than it by 2.76% and 1.43% under 1-shot and 5-shot settings, respectively. Note that the AGPF model currently gets the highest performance, which mainly benefits from additional structures used to construct the feature pyramid and attention pyramid. In contrast, the operations of background suppression and foreground alignment in our method are *parameter-free*, which has significant advantages in efficiency.
- On the Stanford Cars dataset, our method again outperforms all other compared methods with the best performance of 88.93% for 1-shot and 95.20% for 5-shot. Compared to the leading results obtained by AGPF [49] and CAN [2], the relative performance gains are 2.03% and 0.41% for 1-shot and 5-shot, respectively. It’s worth noting that cars are rigid objects and the corresponding intra-category variances are not as significant as birds and dogs. Our methods can reach the best performance of 95.20% under 5-shot setting, which further confirms the significance of our method for FG-FSR.

2) *Results on CUB-200-2011 With Bounding Boxes:* Recently, some state-of-the-art FSL methods (*e.g.* DeepEMD [30], RENet [32], DMN4 [58], *etc*) also demonstrate their generalization ability on the fine-grained dataset, where the images from CUB-200-2011 are pre-processed by

TABLE II

THE ACCURACY (%) OF 5-WAY 1-SHOT AND 5-SHOT TASKS ON THE CROPPED CUB-200-2011, WHERE ALL IMAGES ARE PRE-PROCESSED WITH THE PROVIDED HUMAN-ANNOTATED BOUNDING BOXES. WE REPORT THE MEAN ACCURACY WITH 95% CONFIDENCE INTERVALS ON 2000 TEST EPISODES. NOTE THAT "BB." DENOTES WHETHER TO LEVERAGE THE PROVIDED HUMAN-ANNOTATED BOUNDING BOXES TO CROP THE IMAGES. ° DENOTES THE RESULTS ARE REPORTED BY [30]. * DENOTES THE RESULTS OBTAINED FROM THE ORIGINAL PAPER

Method	BB.	Backbone	CUB-200-2011	
			5-way 1-shot	5-way 5-shot
MatchingNet° [21]	✓	ResNet-12	71.87 ± 0.85	85.08 ± 0.57
ProtoNet [1]	✓	ResNet-12	76.21 ± 0.48	89.70 ± 0.27
RelationNet [15]	✓	ResNet-12	75.81 ± 0.53	88.63 ± 0.29
Baseline [16]	✓	ResNet-12	70.93 ± 0.49	88.57 ± 0.28
Baseline++ [16]	✓	ResNet-12	75.17 ± 0.49	85.58 ± 0.34
DeepEMD* [30]	✓	ResNet-12	75.65 ± 0.83	88.69 ± 0.50
RENet* [32]	✓	ResNet-12	79.49 ± 0.44	91.11 ± 0.24
DMN4* [58]	✓	ResNet-12	82.95 ± 0.75	90.46 ± 0.46
FRN [59]	✓	ResNet-12	82.12 ± 0.85	92.49 ± 0.43
Ours	✗	ResNet-12	82.27 ± 0.46	90.76 ± 0.26
Ours	✓	ResNet-12	86.00 ± 0.41	92.53 ± 0.23

human-annotated bounding boxes. In this section, we also carry out extra experiments to compare our method with these state-of-the-art methods under the same benchmark (*i.e.*, cropped CUB-200-2011) for a fair comparison. The cropped CUB-200-2011 is less challenging because object images cropped by human-annotated bounding boxes have a finer scale and less background, which is helpful to study the fine-grained characteristics for recognition. As shown in Tab. II, our method overall exceeds all compared methods in terms of 1-shot and 5-shot setting and establishes a novel state-of-the-art. Significantly, it can be observed that:

- Compared with the current state-of-the-art methods trained on the cropped CUB-200-2011, our method trained on the uncropped CUB-200-2011 (*i.e.*, w/o BB.) can achieve competitive performance. It is well-known that RENet [32] and DMN4 [58] are the two latest methods in the FSL community. Our method achieves excellent results in the case of not using the cropped dataset, even 2.78% higher than RENet [32] under 1-shot setting and 0.30% higher than DMN4 [58] under 5-shot setting. In fact, cropping object images with human-annotated bounding boxes is equivalent to perfectly removing background noises. This proves the great contribution of *background suppression* to improve the performance of FS-FGR, which also proves the advantages of our method.
- When introducing an additional bounding box to pre-process the CUB-200-2011, our method surpasses all existing methods by a large margin, especially under the 1-shot setting. What's more, compared with the results obtained from the uncropped dataset, our method (*i.e.*, w/ BB.) improves 3.73% and 1.77% in terms of 1-shot and 5-shot setting, respectively. For the cropped CUB-200-2011, the performance of our method is further improved significantly without the interference of

cluttered background. In this case, the excellent performance mainly benefits from two aspects: (1) BAS can further enable the model to crop out the key region of the image with the most discrimination and less redundancy; (2) The *foreground alignment* operation of FOA can enable the model to learn fine-grained features of different parts with different importance degree. Note that the FRN [59] model currently gets the state-of-the-art performance for FS-FGR. Our method, however, still surpasses FRN by 3.88% under the 1-shot setting, which further demonstrates that the proposed method has a significant superiority under the 1-shot setting.

D. Ablation Studies

In this section, we conduct ablation studies to investigate the effectiveness of each individual component of our method. The following experiments are all conducted on the three fine-grained datasets with ResNet-12 as the backbone if not particularly mentioned. We first build a group of baseline models as follows:

- **B0: Global features for FS-FGR** refers to employing backbone to extract global features from raw images, which are directly used to compute the *cosine distances* between the query images and support prototypes. This aims to test the basic performance of our full model by removing BAS, FOA, and L2L. Here, we also retain an auxiliary global classifier to learn in the whole base category space.
- **B1: Raw Local features for FS-FGR** refers to employing the backbone to extract local features only from raw images, which are used to compute the similarity between the query images and support prototypes. This aims to test the superiority of raw local features by adding L2L on the basis of **B0**.
- **B2: Refined Local features for FS-FGR** refers to employing the backbone to extract local features only from refined images, which are used to compute the similarity between the query images and support prototypes. This aims to test the superiority of refined local features by adding L2L and BAS on the basis of **B0**.
- **B3: Dual Local features for FS-FGR** refers to employing backbone to extract local features from raw images and refined images, where both of them are used to compute the similarity between the query images and support prototypes. This aims to test the superiority of our two-stage workflow by adding L2L and BAS on the basis of **B0**.

The comparison results among the above baselines are shown in Table III. In this part, we first investigate the significance of the local feature-based similarity metric for FS-FGR. Obviously, B1-B3 (using local features) outperform B0 (using global features) and lead to consistent performance gains on three datasets. This validates that using local features to compute image-to-image similarity can avoid losing subtle information. With the proposed L2L, we have observed that the performance of B2 is a little lower than B1. That is to say, the baseline model using only the refined image obtained

TABLE III

COMPARISON RESULTS ON BACKBONES WITH/WITHOUT TWO-STAGE LOCAL FEATURE-BASED SIMILARITY METRIC. LOCAL: REMOVING GAP AND USING LOCAL FEATURES. RAW: EMPLOYING FEATURES FROM RAW IMAGES. Refined: EMPLOYING FEATURES FROM REFINED IMAGES

Model	LOCAL	RAW	REFINED	CUB-200-2011		Stanford Dogs		Stanford Cars	
				5-way 1-shot	5-way 5-shot	5-way 1-shot	5-way 5-shot	5-way 1-shot	5-way 5-shot
B0		✓		77.44 ± 0.50	86.82 ± 0.32	65.80 ± 0.51	79.66 ± 0.36	85.20 ± 0.44	92.87 ± 0.24
B1	✓	✓		78.75 ± 0.49	88.96 ± 0.28	67.18 ± 0.50	81.27 ± 0.34	85.87 ± 0.43	93.85 ± 0.24
B2	✓		✓	78.35 ± 0.50	87.62 ± 0.30	64.56 ± 0.51	78.24 ± 0.35	84.04 ± 0.45	92.50 ± 0.25
B3	✓	✓	✓	80.69 ± 0.48	90.21 ± 0.27	68.89 ± 0.49	82.06 ± 0.33	87.65 ± 0.40	94.85 ± 0.21

TABLE IV

CONTRIBUTION OF EACH INDIVIDUAL COMPONENT IN OUR METHOD. BAS: BACKGROUND ACTIVATION SUPPRESSION. L2L: LOCAL TO LOCAL SIMILARITY METRIC. FOA: FOREGROUND OBJECT ALIGNMENT. AE: ATTENTIVE ERASING

Model	BAS	L2L	FOA	AE	CUB-200-2011		Stanford Dogs		Stanford Cars	
					5-way 1-shot	5-way 5-shot	5-way 1-shot	5-way 5-shot	5-way 1-shot	5-way 5-shot
C0	✓				78.18 ± 0.50	87.28 ± 0.31	66.85 ± 0.52	79.68 ± 0.36	86.03 ± 0.42	92.81 ± 0.25
C1	✓	✓			80.69 ± 0.48	90.21 ± 0.27	68.89 ± 0.49	82.06 ± 0.33	87.65 ± 0.40	94.85 ± 0.21
C2	✓	✓	✓		81.71 ± 0.47	90.49 ± 0.27	69.09 ± 0.50	82.25 ± 0.33	88.93 ± 0.39	95.00 ± 0.20
C3	✓	✓	✓	✓	82.27 ± 0.46	90.76 ± 0.26	69.58 ± 0.50	82.59 ± 0.33	88.93 ± 0.38	95.20 ± 0.20
C4		✓	✓	✓	79.75 ± 0.48	89.44 ± 0.28	67.80 ± 0.51	81.59 ± 0.33	88.19 ± 0.39	94.37 ± 0.22

by BAS is inferior to the model using the raw image on each benchmark. The underlying reason could be that the low resolution (*i.e.*, 84×84) of the input image results in inaccurate localization of the completed region for the foreground as shown in Fig. 9. In fact, although BAS can highlight foreground objects, it also leads to some suppression results on some parts of objects, due to the imbalanced feature distribution between foreground and background. However, B3 integrated with local features from both the raw images and refined images achieves the best performance, this verifies the necessity and effectiveness of background suppression for FS-FGR. In addition, a substantial improvement achieved by B3 reveals that our method equipped with BAS considers both the raw image and refined image to alleviate the problem of insufficient samples and makes the model robust to varieties of object scales.

As mentioned above, BAS, L2L, and FOA are three key components of our proposed method. Thus, we also conduct the ablation experiments to analyze the contribution of each component by setting four baselines, as follows:

- **C0: Background suppression for FS-FGR** refers to employing BAS on the backbone to extract global features from raw images and refined images. Both of images are used to compute the *cosine distances* between the query images and support prototypes. This aims to test the superiority of background suppression by adding BAS based on **B0**.
- **C1: Local to local similarity metric for FS-FGR** refers to employing the backbone to extract local features from raw and refined images. Both images are used to compute the similarity between the query images and support prototypes. This aims to test the superiority of the local feature-based similarity metric by adding L2L based on **C0**.
- **C2: Foreground alignment for FS-FGR** refers to employing the backbone to first extract local features

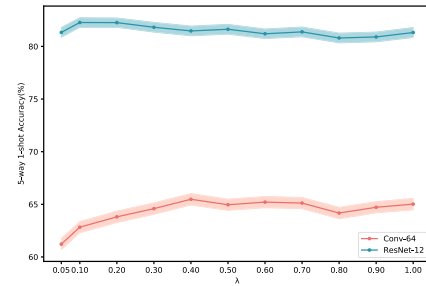


Fig. 8. Performance with respect to different values of λ on CUB-200-2011 under 5-way 1-shot setting. The blue line shows the results of ResNet-12, and the red line shows the results of Conv-64.

from raw images and refined images and then align them. Both images are used to compute the similarity between the query images and support prototypes. This aims to test the superiority of aligned local features by adding FOA based on **C1**.

- **C3: Attentive erasing for FS-FGR** refers to erasing the region with activation higher than a threshold to further improve the completeness of foreground based on BAS. This aims to test the superiority of erasing strategy in our method by adding attentive erasing to the basis of **C2**.
- **C4: Foreground alignment w/o background suppression for FS-FGR** refers to only employing backbone to extract local features from raw images and then align them to compute the similarity between the query images and support prototypes. This aims to verify the complementarity of background suppression and foreground alignment by removing BAS based on **C3**.

Table IV lists the comparison results of the ablation study when using our proposed modules. We have the following observations. First, the addition of BAS solely can only contribute to limited improvement in performance when using global features to compute similarity (*i.e.*, C0 vs. B0). Second,

when introducing local features to measure the distances between the query images and support prototypes, L2L significantly improves the localization accuracy under 1-shot and 5-shot settings (*i.e.*, C1 vs. C0). Third, a substantial improvement of performance on each dataset is gained by integrating FOA to align support features and query features (*i.e.*, C2 vs. C1). This verifies the necessity of considering variations in the object’s pose and position for the FS-FGR task. Finally, erasing the regions with activation higher than a threshold during the training phase alleviates the problem of the model focusing on the most discriminative regions (*i.e.*, C3 vs. C2). The aforementioned experiments and comparisons have demonstrated that the addition of FOA based on BAS can significantly improve the performance of FS-FGR. Moreover, we further demonstrate that BAS and FOA are complementary to each other in improving performance by removing BAS from the full model (*i.e.*, C4 vs. C3). Meanwhile, we observe that C4 outperforms B1 by a large margin, demonstrating the superiority of our feature alignment in the absence of background suppression.

E. Model Complexity Analysis

To comprehensively assess the performance, we compare the efficiency of our method with some publicly available local feature based few-shot learning methods. The results of model parameters (Params.) and floating-point operations (FLOPs) are reported in Table VI. From the aforementioned experiments and comparisons in Table I and Table II, although the proposed method adopts a two-stage framework that would double the amount of computation and storage, we can see that our method still achieves better FS-FGR performance with competitive computational efficiency. Our proposed background activation suppression operation can be theoretically conducted on any position of the network by generating a class-agnostic activation map for localization. With the consideration of both accuracy and efficiency, we consequently choose the high-level features from the last layer of the feature extraction. In summary, without bells and whistles, the efficiency of the proposed two-stage method can benefit from three aspects: (1) the background suppression and foreground alignment operations are only directly conducted on the feature maps without introducing additional parameters; (2) the raw stage and the refined stage are based on the same network with shared parameters, which limits the increase of the number of model parameters; (3) the background suppression and foreground alignment are conducted on high-level features with minimum resolutions.

F. Backbones Analysis

In the previous experiments, we used ResNet-12 as the backbone to extract local features and compute similarities. To further investigate the effectiveness of the proposed method, we switch to another two embedding backbones from shallow to deep, namely Conv-64 and ResNet-18, which are commonly used in the FSL community. Compared to ResNet-18, Conv-64 is shallower with only four layers. Each convolution layer is designed with 3×3 convolution of

TABLE V
THE ACCURACY (%) OF 5-WAY 1-SHOT AND 5-SHOT TASKS ON CUB-200-2011 WITH VARIOUS BACKBONES (*i.e.*, CONV-64 AND RESNET-18). THE RESULTS ARE OBTAINED FROM THE ORIGINAL PAPER. WE REPORT THE MEAN ACCURACY WITH 95% CONFIDENCE INTERVALS ON 2000 TEST EPISODES. THE BEST RESULTS ARE SHOWN IN BOLD

Method	Backbone	CUB-200-2011	
		5-way 1-shot	5-way 5-shot
DN4 [23]	Conv-64	53.15 ± 0.84	81.90 ± 0.60
MattML [20]	Conv-64	66.29 ± 0.56	80.34 ± 0.30
BSNet [48]	Conv-64	65.89 ± 1.00	80.99 ± 0.63
TOAN [52]	Conv-64	65.34 ± 0.75	80.43 ± 0.60
Ours	Conv-64	65.48 ± 0.51	76.02 ± 0.41
ProtoNet [1]	ResNet-18	72.99 ± 0.88	86.64 ± 0.51
Neg-Margin [60]	ResNet-18	72.66 ± 0.85	89.40 ± 0.43
Centroid [61]	ResNet-18	74.22 ± 1.09	88.65 ± 0.55
RAP-Neg-Margin [62]	ResNet-18	75.37 ± 0.81	90.61 ± 0.39
Ours	ResNet-18	84.19 ± 0.43	92.46 ± 0.23

TABLE VI
COMPARISON OF THE EFFICIENCY OF SOME PUBLICLY AVAILABLE LOCAL FEATURE BASED FEW-SHOT LEARNING METHODS. SMALLER PARAMS. AND SMALLER FLOPs INDICATE BETTER EFFICIENCY

Method	Backbone	Model Complexity	
		Params. (M)	FLOPs (G)
DN4 [23]	ResNet-12	12.42	67.39
CAN [2]	ResNet-12	8.04	12.75
DeepEMD [30]	ResNet-12	12.42	35.23
RENet [32]	ResNet-12	12.63	35.70
DMN4 [58]	ResNet-12	12.42	35.23
FRN [59]	ResNet-12	12.42	35.18
AGPF [49]	ResNet-12	8.77	51.53
Ours (1-stage)	ResNet-12	8.04	16.88
Ours (2-stage)	ResNet-12	8.04	33.76

64 filters, followed by batch normalization, a **RELU** activation function, and a 2×2 max-pooling. To keep consistent with the previous experimental settings, and to preserve the spatial resolution of features, we remove the last max-pooling layer of Conv-64. Thus, the size of the output feature maps becomes $64 \times 10 \times 10$. The top half of Table V shows that our method achieves stable and competitive performance under 1-shot setting when using Conv-64. Our method, however, performs worse than some recent methods under the 5-shot setting. The underlying reason could be that Conv-64 is too shallow and the poor feature expression ability can not approximate the spatial distribution of the foreground efficiently, which results in the model not being able to generate good foreground object coordinates for background suppression as shown in Fig. 11. Thus, sub-optimal refined images cannot promote the performance of our method and even bring some negative effects. The existing state-of-the-art FS-FGR models also embed additional complex structures in the basic feature extractor, where redundant parameters can significantly promote the performance of shallow backbones. Compared with them, our model has no redundant structure and parameter

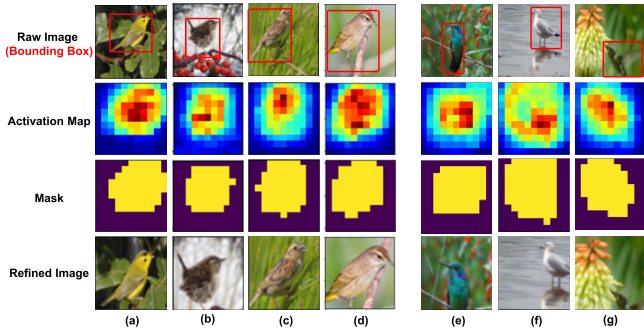


Fig. 9. Visualization of background activation suppression under 5-way 1-shot setting on CUB-200-2011 dataset, including some success and failure cases. For better comparison, the ground-truth bounding box provided by the original dataset is marked in each raw image.

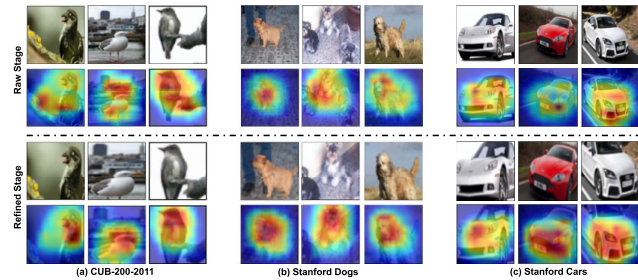


Fig. 10. Visualization results of heatmaps in the raw stage and refined stage. We can observe that the fine-grained visual cues of the attended regions from the refined stage are more clear and discriminative to the corresponding categories, and are easier to be recognized than those in the raw stage (zooming in is recommended for better visualization).

except the two global classifiers. Significantly, the operations of background suppression and foreground alignment in our method are *parameter-free*, which has significant advantages in efficiency. As shown in the bottom half of Table V, when using a deeper backbone (*i.e.*, ResNet-18, the standard unchanged architecture used in [63]) as feature extractor, our method overall exceeds all compared methods and further establishes a novel state-of-the-art. Compared with ResNet-12, ResNet-18 achieves a significant improvement of 1.92% under a 5-way 1-shot setting and 1.70% under a 5-way 5-shot setting, respectively. The reason may be that the deeper feature extractor can significantly improve the accuracy of background activation suppression and enhance the semantic correlation of foreground object alignment.

G. Sensitivity Analysis

As shown in Section III-E, there is a hyper-parameter λ in the final loss function to control the contribution of $\mathcal{L}_{\text{local}}$. We conduct a comparative experiment on CUB-200-2011 under 5-way 1-shot setting to study the effect of this factor across different backbones. From Fig. 8, when achieving the best performance, we have observed that the value of λ on Conv-64 is relatively larger compared to the ResNet-12. The results show the fact that the local loss accounts for a larger proportion of the total loss for the shallower network and delays the appearance of performance saturation points.

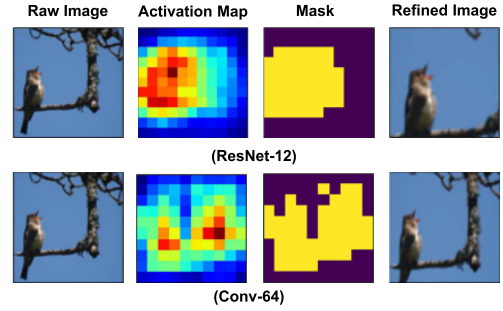


Fig. 11. Visualization comparison of background activation suppression with different backbones on CUB-200-2011 dataset. It can be found that a more compact activation map, complete mask, and accurate refined image can be generated when using deeper ResNet-12 as a feature extractor.

TABLE VII

THE ACCURACY (%) COMPARISON OF 5-WAY 1-SHOT AND 5-SHOT TASKS ON CUB-200-2011 BY ADJUSTING TWO WEIGHTING FACTORS IN THE LOSS FUNCTIONS. WE REPORT THE MEAN ACCURACY WITH 95% CONFIDENCE INTERVALS ON 2000 TEST EPISODES

α	β	Backbone	CUB-200-2011	
			5-way 1-shot	5-way 5-shot
0.1	0.9	Conv-64	61.79 \pm 0.55	76.39 \pm 0.41
0.3	0.7	Conv-64	63.84 \pm 0.54	75.73 \pm 0.42
0.5	0.5	Conv-64	65.48 \pm 0.51	76.02 \pm 0.41
0.7	0.3	Conv-64	62.97 \pm 0.53	75.95 \pm 0.43
0.9	0.1	Conv-64	62.02 \pm 0.53	75.62 \pm 0.43
0.1	0.9	ResNet-12	81.06 \pm 0.48	89.68 \pm 0.28
0.3	0.7	ResNet-12	81.73 \pm 0.46	90.34 \pm 0.26
0.5	0.5	ResNet-12	82.27 \pm 0.46	90.76 \pm 0.26
0.7	0.3	ResNet-12	80.76 \pm 0.48	90.16 \pm 0.27
0.9	0.1	ResNet-12	79.81 \pm 0.48	89.07 \pm 0.28

This is because, during the training phase, the shallow feature extractor has a weak feature expression ability and is unable to explore the inter-category relationship for FS-FGR. In all our experiments, we set the value of λ to 0.1 for the deeper network (*i.e.*, ResNet-12 and ResNet-18), and 0.4 for the shallow network (*i.e.*, Conv-64) to achieve the final results.

There are two weighting factors α and β in global loss (*i.e.* Eq. 6) and local loss (*i.e.* Eq. 9) to control the contributions of raw image and refined image. The sensitivity analyses of two parameters according to $\alpha + \beta = 1$ are performed on the CUB-200-2011 dataset and the results are presented in Table VII. We can observe a significant gap in the 1-shot performance between different settings of the weights. Meanwhile, although different backbones have different feature expression abilities and foreground predictions, the highest mean value of performance is achieved when the weights of the two factors both equal to 0.5. The underlying reason could be that the raw stage and the refined stage are complementary to each other in improving the performance, which is also proved in Fig. 12. Thus, we simply select equal weights for the two weighting factors in the training and inference stage.

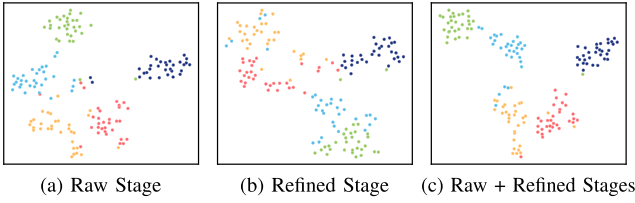


Fig. 12. The t-SNE visualization of feature distribution in different stages on CUB-200-2011. We randomly select five classes (one for each color) with 30 query images in each class (one for each point). (a) Raw Stage: features are from only the raw stage. (b) Refined Stage: features are from only the refined stage. (c) Raw + Refined Stages: features are from both stages.

H. Visualization Analysis

In this section, we first visualize the results of background suppression on the CUB-200-2011. As presented in Fig. 9(a)–(d), the proposed BAS module can cover the object region more completely and effectively remove the cluttered background, which enables the model to see more clearly with only image-level annotations. But as shown in Fig. 9(e)–(g), we also find that the proposed BAS module performs poorer effectiveness for localizing smaller objects, especially the small object that is far away from the center of the image.

We gracefully argue that taking an amplified attended region without distraction from the background as input in the refine stage can distill more discriminative fine-grained characteristics in low-data scenarios. The insight behind it is that “see closer, see better” because the background noises can distract attention from the subtle differences. To verify this idea, we visualize the feature heatmaps of the two stages in Figure 10, it is clear that our model focuses on more discriminative visual cues in the refined stage, which confirms the significance of our two-stage paradigm for FS-FGR.

In Fig. 12, we visualize the feature distribution in the embedding space by t-SNE [64] on CUB-200-2011 under 5-way 1-shot setting, where 30 query images are randomly sampled from each category. It can be seen that the features obtained from our full model with both raw and refined stages have more compact and separable clusters than the other two models using only the raw stage or refined stage, which indicates that our method has good adaptability to the object’s scale. Furthermore, the visualization of the model using only raw images is better than the one using only refined images. This further substantiates the correlation analysis in Section IV-D and indicates that the raw stage and the refined stage are complementary to each other in improving performance.

Lastly, we investigate the effects of the attentive erasing strategy in foreground localization and background suppression. Fig. 13 provides the visual results on CUB-200-2011, which proves the advantages of the proposed attentive erasing strategy. We observe that the proposed attentive erasing strategy facilitates the generation of foreground maps by indirectly increasing the foreground activation value, and the masked prediction map achieves better coverage of the localization results on the foreground object.

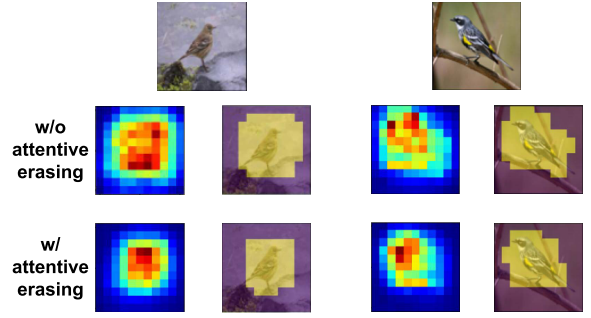


Fig. 13. Visualization of the influence of attentive erasing strategy for foreground localization and background suppression.

TABLE VIII
THE ACCURACY (%) OF 5-WAY 1-SHOT AND 5-SHOT TASKS ON CUB-200-2011 WITH VARIOUS SETTINGS

Method	Backbone	CUB-200-2011	
		5-way 1-shot	5-way 5-shot
Ours(w/o BAS)	ResNet-12	79.75 ± 0.48	89.44 ± 0.28
Ours(w/ BAS)	ResNet-12	82.27 ± 0.46	90.76 ± 0.26
Ours(w/ BAS×2)	ResNet-12	81.62 ± 0.47	90.48 ± 0.27
Ours(w/ BB.)	ResNet-12	86.29 ± 0.42	93.23 ± 0.22

I. Discussion

In this paper, we explore a two-stage weakly-supervised framework by introducing two parameter-free adjustments: *background suppression* and *foreground alignment*. Here, we also conduct some analyses to intuitively demonstrate the effectiveness and limitation of the proposed method.

To deal with the cluttered background, we propose a BAS module to disentangle the foreground object and background distraction in an image. As shown in Table VIII, ‘Ours(w/ BAS)’ achieves significant and consistent improvement on both 1-shot and 5-shot tasks when compared to ‘Ours(w/o BAS).’ But as seen in Fig. 9, we note that our BAS is inconsistent for localizing objects of different sizes, especially poorer localization for small objects in the corner. Although the refined image generated by BAS covers the entire area of the foreground, it also contains some low-activated background regions when compared to the foreground region cropped by the ground-truth bounding box.

Going further, we attempt to achieve better localization of the foreground object by using BAS twice, *i.e.*, Ours(w/ BAS×2), which is devised to maximize the ratio of foreground activation and overall activation. It is a pity that the performance does not achieve further improvement as expected, which also indicates that the proposed BAS is insufficient to resist background distractions in some cases. The underlying reason could be that the limitation of feature map resolution causes imprecise localization, especially since the distribution between the foreground and background is unbalanced.

Lastly, we keep the raw image unchanged and replace the refined image with the image cropped by the ground-truth bounding box. As shown in Table VIII, ‘Ours(w/ BB.)’ can

be seen as an oracle and the upper bound of our method. Not surprisingly, ‘Ours(w/ BB.)’ surpasses all existing methods by a large margin, which further demonstrates the potential of our proposed two-stage framework. Thus, how to overcome the limitation of the BAS module to achieve similar results with the ground-truth bounding box will be researched in future work.


V. CONCLUSION

In this paper, we show that background suppression and foreground alignment matter for few-shot fine-grained recognition (FS-FGR). Thus, we propose a novel two-stage weakly-supervised framework for FS-FGR, where the background suppression and foreground alignment are jointly learned in an end-to-end manner. In particular, a background activation suppression module is introduced in the raw stage to weaken background disturbance and enhance dominative foreground objects for the refined stage. In each stage, a foreground object alignment module is proposed to align the foreground object of support features concerning query features for addressing the problem of spatial misalignment. Finally, a local-to-local similarity metric is introduced to enable the model to capture subtle differences of confused sample pairs on two stages for the final decision. Extensive experiments on three fine-grained benchmarks demonstrate that our proposed model can be trained in an end-to-end manner with only the image-level label, and achieve state-of-the-art performance. What calls for special attention is that this is the first research to reveal the impact of background suppression and foreground alignment, and the first attempt to integrate both aspects into a standard pipeline for FS-FGR task. In addition, some designs in our designed generic framework include some heuristics, which could be further studied in detail.

REFERENCES

- J. Snell, K. Swersky, and R. Zemel, “Prototypical networks for few-shot learning,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 4077–4087.
- R. Hou, H. Chang, B. Ma, S. Shan, and X. Chen, “Cross attention network for few-shot classification,” in *Proc. NIPS*, 2019, pp. 4005–4016.
- H. Tang, Z. Li, Z. Peng, and J. Tang, “Blockmix: Meta regularization and self-calibrated inference for metric-based meta-learning,” in *Proc. 28th ACM Int. Conf. Multimedia*, 2020, pp. 610–618.
- Y. Fu, L. Zhang, J. Wang, Y. Fu, and Y. Jiang, “Depth guided adaptive meta-fusion network for few-shot video recognition,” in *Proc. 28th ACM Int. Conf. Multimedia*, 2020, pp. 1142–1151.
- Y. Fu, Y. Fu, and Y.-G. Jiang, “Meta-FDMixup: Cross-domain few-shot learning guided by labeled target data,” in *Proc. 29th ACM Int. Conf. Multimedia*, Oct. 2021, pp. 5326–5334.
- T.-Y. Lin, A. RoyChowdhury, and S. Maji, “Bilinear CNN models for fine-grained visual recognition,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 1449–1457.
- X.-S. Wei, C.-W. Xie, J. Wu, and C. Shen, “Mask-CNN: Localizing parts and selecting descriptors for fine-grained bird species categorization,” *Pattern Recognit.*, vol. 76, pp. 704–714, Apr. 2018.
- H. Zheng, J. Fu, Z. Zha, and J. Luo, “Looking for the devil in the details: Learning trilinear attention sampling network for fine-grained image recognition,” in *Proc. CVPR*, Jun. 2019, pp. 5012–5021.
- F. Zhang, M. Li, G. Zhai, and Y. Liu, “Multi-branch and multi-scale attention learning for fine-grained visual categorization,” in *Proc. Int. Conf. Multimedia Modeling*, vol. 12572, 2021, pp. 136–147.
- Y. Wu et al., “Object-aware long-short-range spatial alignment for few-shot fine-grained image classification,” in *Proc. ACM Multimedia*, 2021, pp. 107–115.
- R. Yan, J. Tang, X. Shu, Z. Li, and Q. Tian, “Participation-contributed temporal dynamic model for group activity recognition,” in *Proc. 26th ACM Int. Conf. Multimedia*, 2018, pp. 1292–1300.
- D. Wang, J. Liu, X. Fan, and R. Liu, “Unsupervised misaligned infrared and visible image fusion via cross-modality image generation and registration,” in *Proc. IJCAI*, 2022, pp. 3508–3515.
- D. Wang, L. Ma, R. Liu, and X. Fan, “Semantic-aware texture-structure feature collaboration for underwater image enhancement,” in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2022, pp. 4592–4598.
- C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie, “The caltech-UCSD birds-200–2011 dataset,” California Inst. Technol., Pasadena, CA, USA, Tech. Rep. CNS-TR-2011-001, 2011.
- F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. S. Torr, and T. M. Hospedales, “Learning to compare: Relation network for few-shot learning,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 1199–1208.
- W. Chen, Y. Liu, Z. Kira, Y. F. Wang, and J. Huang, “A closer look at few-shot classification,” in *Proc. ICLR*, 2019, pp. 1–17.
- C. Finn, P. Abbeel, and S. Levine, “Model-agnostic meta-learning for fast adaptation of deep networks,” in *Proc. 34th Int. Conf. Mach. Learn.*, vol. 70, Aug. 2017, pp. 1126–1135.
- K. Lee, S. Maji, A. Ravichandran, and S. Soatto, “Meta-learning with differentiable convex optimization,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 10657–10665.
- Q. Cai, Y. Pan, T. Yao, C. Yan, and T. Mei, “Memory matching networks for one-shot image recognition,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 4080–4088.
- Y. Zhu, C. Liu, and S. Jiang, “Multi-attention meta learning for few-shot fine-grained image recognition,” in *Proc. 29th Int. Joint Conf. Artif. Intell.*, Jul. 2020, pp. 1090–1096.
- O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra, “Matching networks for one shot learning,” in *Proc. NIPS*, 2016, pp. 3630–3638.
- M. Andrychowicz et al., “Learning to learn by gradient descent by gradient descent,” in *Proc. NIPS*, 2016, pp. 3981–3989.
- W. Li, L. Wang, J. Xu, J. Huo, Y. Gao, and J. Luo, “Revisiting local descriptor based image-to-class measure for few-shot learning,” in *Proc. CVPR*, Jun. 2019, pp. 7260–7268.
- A. Khosla, N. Jayadevaprakash, B. Yao, and F.-F. Li, “Novel dataset for fine-grained image categorization: Stanford dogs,” in *Proc. CVPR Workshop*, 2011, pp. 1–2.
- J. Krause, M. Stark, J. Deng, and L. Fei-Fei, “3D object representations for fine-grained categorization,” in *Proc. ICCV Workshops*, 2013, pp. 554–561.
- A. A. Rusu et al., “Meta-learning with latent embedding optimization,” in *Proc. ICLR*, 2019, pp. 1–17.
- R. Zhang, T. Che, Z. Ghahramani, Y. Bengio, and Y. Song, “MetaGAN: An adversarial approach to few-shot learning,” in *Proc. NIPS*, 2018, pp. 2371–2380.
- B. Hariharan and R. Girshick, “Low-shot visual recognition by shrinking and hallucinating features,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 3037–3046.
- K. Li, Y. Zhang, K. Li, and Y. Fu, “Adversarial feature hallucination networks for few-shot learning,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 13467–13476.
- C. Zhang, Y. Cai, G. Lin, and C. Shen, “DeepEMD: Few-shot image classification with differentiable Earth Mover’s distance and structured classifiers,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 12200–12210.
- F. Hao, F. He, J. Cheng, and D. Tao, “Global-local interplay in semantic alignment for few-shot learning,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 32, no. 7, pp. 4351–4363, Jul. 2022.
- D. Kang, H. Kwon, J. Min, and M. Cho, “Relational embedding for few-shot classification,” in *Proc. ICCV*, Oct. 2021, pp. 8802–8813.
- W. Jiang, K. Huang, J. Geng, and X. Deng, “Multi-scale metric learning for few-shot learning,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 31, no. 3, pp. 1091–1102, Mar. 2021.
- C. Zhang, C. Li, and J. Cheng, “Few-shot visual classification using image pairs with binary transformation,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 30, no. 9, pp. 2867–2871, Sep. 2020.
- X.-S. Wei et al., “Fine-grained image analysis with deep learning: A survey,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 12, pp. 8927–8948, Dec. 2022.
- H. Li, S. Yan, Z. Yu, and D. Tao, “Attribute-identity embedding and self-supervised learning for scalable person re-identification,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 30, no. 10, pp. 3472–3485, Oct. 2020.

- [37] H. Li, N. Dong, Z. Yu, D. Tao, and G. Qi, "Triple adversarial learning and multi-view imaginative reasoning for unsupervised domain adaptation person re-identification," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 32, no. 5, pp. 2814–2830, May 2022.
- [38] S. Yan, Y. Zhang, M. Xie, D. Zhang, and Z. Yu, "Cross-domain person re-identification with pose-invariant feature decomposition and hyper-graph structure alignment," *Neurocomputing*, vol. 467, pp. 229–241, Jan. 2022.
- [39] N. Zhang, J. Donahue, R. Girshick, and T. Darrell, "Part-based R-CNNs for fine-grained category detection," in *Proc. Eur. Conf. Comput. Vis.*, vol. 8689, Sep. 2014, pp. 834–849.
- [40] M. Lam, B. Mahasseni, and S. Todorovic, "Fine-grained recognition as HSnet search for informative image parts," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 6497–6506.
- [41] Y. Ding et al., "AP-CNN: Weakly supervised attention pyramid convolutional neural network for fine-grained visual classification," *IEEE Trans. Image Process.*, vol. 30, pp. 2826–2836, 2021.
- [42] Y. Gao, O. Beijbom, N. Zhang, and T. Darrell, "Compact bilinear pooling," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 317–326.
- [43] Y. Hu, Y. Yang, J. Zhang, X. Cao, and X. Zhen, "Attentional kernel encoding networks for fine-grained visual categorization," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 31, no. 1, pp. 301–314, Jan. 2021.
- [44] J. Fu, H. Zheng, and T. Mei, "Look closer to see better: Recurrent attention convolutional neural network for fine-grained image recognition," in *Proc. CVPR*, Jun. 2017, pp. 4476–4484.
- [45] X. He, Y. Peng, and J. Zhao, "Fast fine-grained image classification via weakly supervised discriminative localization," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 29, no. 5, pp. 1394–1407, May 2019.
- [46] J. Guo, G. Qi, S. Xie, and X. Li, "Two-branch attention learning for fine-grained class incremental learning," *Electronics*, vol. 10, no. 23, p. 2987, Dec. 2021.
- [47] X.-S. Wei, P. Wang, L. Liu, C. Shen, and J. Wu, "Piecewise classifier mappings: Learning fine-grained learners for novel categories with few examples," *IEEE Trans. Image Process.*, vol. 28, no. 12, pp. 6116–6125, Dec. 2019.
- [48] X. Li, J. Wu, Z. Sun, Z. Ma, J. Cao, and J.-H. Xue, "BSNet: Bi-similarity network for few-shot fine-grained image classification," *IEEE Trans. Image Process.*, vol. 30, pp. 1318–1331, 2021.
- [49] H. Tang, C. Yuan, Z. Li, and J. Tang, "Learning attention-guided pyramidal features for few-shot fine-grained recognition," *Pattern Recognit.*, vol. 130, Oct. 2022, Art. no. 108792.
- [50] H. Huang, J. Zhang, J. Zhang, J. Xu, and Q. Wu, "Low-rank pairwise alignment bilinear network for few-shot fine-grained image classification," *IEEE Trans. Multimedia*, vol. 23, pp. 1666–1680, 2021.
- [51] S. Tian, H. Tang, and L. Dai, "Coupled patch similarity network FOR one-shot fine-grained image recognition," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2021, pp. 2478–2482.
- [52] H. Huang, J. Zhang, L. Yu, J. Zhang, Q. Wu, and C. Xu, "TOAN: Target-oriented alignment network for fine-grained image categorization with few labeled samples," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 32, no. 2, pp. 853–866, Feb. 2022.
- [53] Z. Wu, Y. Li, L. Guo, and K. Jia, "PARN: Position-aware relation networks for few-shot learning," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 6658–6666.
- [54] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, vol. 8689, 2014, pp. 818–833.
- [55] Q. Hou, P. Jiang, Y. Wei, and M. Cheng, "Self-erasing network for integral object attention," in *Proc. NIPS*, 2018, pp. 547–557.
- [56] Y. Wei, J. Feng, X. Liang, M. Cheng, Y. Zhao, and S. Yan, "Object region mining with adversarial erasing: A simple classification to semantic segmentation approach," in *Proc. CVPR*, Jun. 2017, pp. 6488–6496.
- [57] S. Yan, H. Tang, L. Zhang, and J. Tang, "Image-specific information suppression and implicit local alignment for text-based person search," 2022, *arXiv:2208.14365*.
- [58] Y. Liu, T. Zheng, J. Song, D. Cai, and X. He, "DMN4: Few-shot learning via discriminative mutual nearest neighbor neural network," in *Proc. AAAI Conf. Artif. Intell.*, vol. 36, no. 2, Jun. 2022, pp. 1828–1836.
- [59] D. Wertheimer, L. Tang, and B. Hariharan, "Few-shot classification with feature map reconstruction networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 8012–8021.
- [60] B. Liu et al., "Negative margin matters: Understanding margin in few-shot classification," in *Proc. ECCV*, pp. 12349, 2020, pp. 438–455.
- [61] A. Afrasiyabi, J. Lalonde, and C. Gagné, "Associative alignment for few-shot image classification," in *Proc. ECCV*, vol. 12350, 2020, pp. 18–35.
- [62] J. Hong et al., "Reinforced attention for few-shot learning and beyond," in *Proc. CVPR*, Jun. 2021, pp. 913–923.
- [63] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. CVPR*, Jun. 2016, pp. 770–778.
- [64] L. Van der Maaten and G. Hinton, "Visualizing data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, no. 11, pp. 1–27, 2008.



Zican Zha received the B.E. degree from the Jiangsu University of Science and Technology, Zhenjiang, China, in 2019. He is currently pursuing the master's degree with the School of Computer Science and Engineering, Nanjing University of Science and Technology, Nanjing, China. His research interests include multimedia computing, computer vision, and deep learning.



Hao Tang received the B.E. degree from Harbin Engineering University in 2018. He is currently pursuing the Ph.D. degree with the School of Computer Science and Engineering, Nanjing University of Science and Technology. His research interests include multimedia computing, computer vision, and deep learning.



Yunlian Sun received the M.E. degree in computer science and technology from the Harbin Institute of Technology, China, in 2010, and the Ph.D. degree in ingegneria elettronica, informatica e delle telecomunicazioni from the University of Bologna, Italy, in 2014. After the Ph.D. study, she worked as a Post-Doctoral Researcher at the National Laboratory of Pattern Recognition (NLPR), Institute of Automation, Chinese Academy of Sciences. She is currently an Associate Professor at the School of Computer Science and Engineering, Nanjing University of Science and Technology, China. Her research interests include biometrics, pattern recognition, and computer vision.



Jinhui Tang (Senior Member, IEEE) received the B.E. and Ph.D. degrees from the University of Science and Technology of China, Hefei, China, in 2003 and 2008, respectively. He is currently a Professor with the Nanjing University of Science and Technology. He has authored more than 150 papers in top-tier journals and conferences. His research interests include multimedia analysis and computer vision. He was a recipient of the Best Paper Awards from ACM MM 2007, PCM 2011, and ICIMCS 2011, the Best Paper Runner-up Award from ACM MM 2015, and the Best Student Paper Awards from MMM 2016 and ICIMCS 2017. He has served as an Associate Editor for the IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS (TNNLS), the IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING (TKDE), and the IEEE TRANSACTIONS ON MULTIMEDIA (TMM).