

Robust Saliency-Aware Distillation for Few-shot Fine-grained Visual Recognition

Haiqi Liu, C. L. Philip Chen *Fellow, IEEE*, Xinrong Gong, Tong Zhang *Senior Member, IEEE*

Abstract—Recognizing novel sub-categories with scarce samples is an essential and challenging research topic in computer vision. Existing literature addresses this challenge by employing local-based representation approaches, which may not sufficiently facilitate meaningful object-specific semantic understanding, leading to a reliance on apparent background correlations. Moreover, they primarily rely on high-dimensional local descriptors to construct complex embedding space, potentially limiting the generalization. To address the above challenges, this article proposes a novel model, Robust Saliency-aware Distillation (RSaD), for few-shot fine-grained visual recognition. RSaD introduces additional saliency-aware supervision via saliency detection to guide the model toward focusing on the intrinsic discriminative regions. Specifically, RSaD utilizes the saliency detection model to emphasize the critical regions of each sub-category, providing additional object-specific information for fine-grained prediction. RSaD transfers such information with two symmetric branches in a mutual learning paradigm. Furthermore, RSaD exploits inter-regional relationships to enhance the informativeness of the representation and subsequently summarize the highlighted details into contextual embeddings to facilitate the effective transfer, enabling quick generalization to novel sub-categories. The proposed approach is empirically evaluated on three widely used benchmarks, demonstrating its superior performance.

Index Terms—Few-shot Fine-grained Visual Recognition, Few-shot Learning, Saliency Detection, Mutual Learning

I. INTRODUCTION

FINE-GRAINED visual recognition (FGVR) is crucial and challenging research that aims to distinguish visually similar objects within a specific category, such as different species of birds or types of cars. With the significant advancements in computer vision [1]–[3], extensive research [4]–[6] has been conducted to improve the performance of FGVR. However, due to the considerable expenses entailed in fine-grained annotation and the relative scarcity of rare categories, the extensive applicability of such approaches is

This work was funded in part by the National Natural Science Foundation of China grant under numbers 62076102, 62222603, and 92267203, in part by the National Key Research and Development Program of China under number 2019YFA0706200, in part by the STI2030-Major Projects grant from the Ministry of Science and Technology of the People’s Republic of China under number 2021ZD0200700, in part by the Key-Area Research and Development Program of Guangdong Province under number 2023B0303030001, in part by the Guangdong Natural Science Funds for Distinguished Young Scholar under number 2020B1515020041, and in part by the Program for Guangdong Introducing Innovative and Entrepreneurial Teams (2019ZT08X214). (*Corresponding author: Tong Zhang*)

The authors are with the School of Computer Science and Engineering, South China University of Technology, Guangzhou 510006, China, and is with the Brain and Affective Cognitive Research Center, Pazhou Lab, Guangzhou 510335, China. C. L. Philip Chen and Tong Zhang are also with the Engineering Research Center of the Ministry of Education on Health Intelligent Perception and Paralleled Digital-Human, Guangzhou 510006, China. (e-mail: tony@scut.edu.cn)

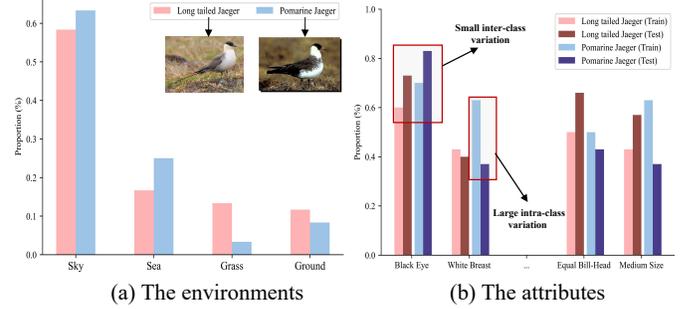


Fig. 1: The statistical analysis of both external environments and intrinsic attributes across two sub-categories within the CUB-200-2011 dataset. The results reveal that 1) environments and specific attributes across different sub-categories may exhibit significant similarities; 2) within the same sub-category, identical attributes can vary considerably.

limited by the annotated data scarcity issue. Unlike machines, humans can distinguish subtle variations among previously unseen sub-categories, even when presented with one sample. Therefore, improving the generalizability of FGVR models in data-scarce scenarios is imperative, thereby reducing the disparities between humans and machines. Thus, this paper focuses on the few-shot fine-grained visual recognition (FS-FGVR) task, which involves utilizing scarce training examples per category (typically 1 or 5 samples) to recognize unseen similar-looking sub-categories.

The main challenge of FS-FGVR is quickly capturing key regions between similar-looking sub-categories. Existing works [7]–[10] are mainly dedicated to solving this challenge in a local representation manner. In local representation-based approaches, the relationship between the support set (training data) and query set (testing data) is explored by the feature alignment [8], [9] or feature reconstruction [10] to highlight semantic-related local features as the key regions. Despite the remarkable accomplishments in these approaches, there has been a lack of substantial focus on exploring the intrinsic structure (specifically, background and foreground) of fine-grained images, which are crucial factors for FGVR [11]–[13]. Fig.1(a) illustrates that environments across different sub-categories may exhibit notable similarities. The local-representation-based approach may inadequately promote object-relevant semantic understanding, thereby shifting attention towards apparent background correlations, which causes misclassification. Recent studies [11]–[14] have attempted to address this issue by integrating saliency detection models intended

to redirect the model’s attention to object-related features. However, these models generally apply saliency information for data or feature enhancement at the pixel or feature level. Such approaches present a considerable challenge in understanding high-dimensional saliency information with limited samples. The Deutsch–Norman late-selection filter model [15] in neuroscience indicates that humans will amplify meaningful information and weaken irrelevant information after processing it for meaning. Considering the previously stated observations, this paper designs a saliency-aware guidance strategy to encourage the model to focus on the object-relevant regions during prediction (low-dimensional space), thereby effectively reducing the negative impact of background clutter.

Furthermore, when confronted with novel fine-grained visual concept tasks, humans tend to identify noticeable distinctions and memorable attributes between them. It is rational that simple salient patterns can easily generalize to unseen sub-categories [16]. Previous approaches [8], [9] have often relied on high-dimensional deep local descriptors to establish embedding spaces through Bilinear Pooling [17] and Relation Network [18] techniques. Despite the fine-grained nature of these embedding spaces, their inherent space complexity poses challenges in terms of generalization and computational demands. Moreover, as depicted in Fig.1(b), it is evident that partial characters may not provide reliable clues for FS-FGVR. Some works [19]–[21] indicate that capturing the difference in holistic view can also effectively distinguish unseen similar-like visual concepts. Based on the above observations, this paper designed the representation highlight&summarize module by learning a simplified contextual embedding space to ensure quick generalization towards novel sub-categories.

To address the above challenges, this paper proposes a Robust Saliency-aware Distillation model called RSaD to generate fine-grained intrinsic and transferable embedding for FS-FGVR. Specifically, RSaD consists of the Saliency-aware Guidance (SaG) strategy and Representation Highlight&Summarize (RHS) module. The SaG is designed to exploit explicit intrinsic relationship of sub-categories via the distillation technique on the salient region probability distribution. The alignment of salient region probability distribution reduces the impact of background clutter, which provides more pure object information for fine-grained prediction. Subsequently, the RHS module is designed to highlight the object-relevant deep descriptors and squeeze them as the contextual embedding. It can effectively capture discriminative, simple patterns across sub-categories, facilitating rapid generalization to novel sub-categories while maintaining a low model complexity. Extensive experiments are conducted on widely used benchmarks to evaluate the proposed approach. Results indicate that the proposed model outperforms current state-of-the-art FS-FGVR approaches, confirming its effectiveness. To sum up, this work makes the following contributions:

- 1) This article proposes a novel saliency-aware guidance strategy for FS-FGVR. This method first exploits object-specific information via the distillation technique on the salient region probability distribution, which provides more accurate clues for fine-grained prediction.
- 2) This article proposes the RHS module, which involves

learning a robust embedding space to facilitate rapid generalization towards novel sub-categories.

- 3) Extensive experiments are conducted to verify the effectiveness of the proposed approach. The results demonstrate that the proposed approach achieves comparable performance with state-of-the-art approaches.

The remainder of this article is as follows: Section II provides a brief review of related work. Section III presents the details of the RSaD model. Section IV presents and analyzes comparative experiment results. Then, Section V discusses the impact of different modules proposed in this article. Finally, Section VI concludes with a summary and outlook.

II. RELATED WORK

A. *Fined-Grained Visual Recognition*

Fine-grained Visual Recognition (FGVR) and fine-grained image retrieval [22] have emerged as critical components in fine-grained image analysis [23]. FGVR identifies visually similar objects within the same meta-category, typically at a sub-category or individual instance level. Early research [24]–[26] on FGVR was based on hand-crafted features such as attributes, SIFT [27] and part annotations. However, the ability of such methods was somewhat limited due to the drawbacks of hand-crafted features, i.e., lack of information and high labeling effort [28]. Benefiting from the success of convolutional neural networks (CNNs), research on FGVR has shifted towards learning deep feature representations. For example, Lin *et al.* [29] introduced bilinear CNN models to generate rich deep features for FGVR, while Kong *et al.* [30] designed a low-rank bilinear pooling strategy to reduce the computational effort required by high-dimensional bilinear features. Some previous works [31]–[34] have leveraged object detection models such as FRN and R-CNN to select informative regions. These part-based approaches are practical in locating and analyzing fine-grained local features. However, these methods often incur significant computational costs and substantial data requirements, primarily attributable to the necessity of detailed analysis using auxiliary sub-networks or selective strategies [35]. More recently, several studies [6], [36]–[38] utilized vision transformer [3] to improve model capacity in FGVR. Regardless of the demonstrated successes, large-scale data is necessary for model training. Learning with limited data remains an open problem in FGVR. Thus, this paper focuses on the FS-FGVR task, which entails effectively leveraging limited training examples per sub-categories.

B. *Few-Shot Fine-Grained Visual Recognition*

Recently, research [7], [39] has emerged focusing on a more practical and challenging setting, termed FS-FGVR (Few-Shot Fine-Grained Visual Recognition). FS-FGVR aims to distinguish novel sub-categories with limited samples. Thanks to the success of few-shot learning [40], FS-FGVR research has made significant progress in recent years and can be broadly classified into two groups:

Global representation-based methods aim to learn discriminative global features between sub-categories. Li *et al.* [41]

proposed the bi-similarity network, which considers various metric functions to find a better global feature space. Xu *et al.* [21] introduced a dual attention mechanism to provide a refined global embedding for recognition. Xu *et al.* [20] utilized β -TCVAE [42] to learn transferable intra-class variance to generate additional global features applicable to novel classes.

Local representation-based methods focus on learning the discriminative parts of the whole image. They argue that global representations obtained by pooling lack critical spatial information in FS-FGVR. In local representation-based approaches, the relationship between the support set and query set is primarily leveraged through local descriptor alignment or local descriptor reconstruction to emphasize semantically related local features as discriminative features. Wei *et al.* [39] proposed a novel exemplar-to-classifier mapping strategy that can learn a discriminative classifier within bilinear CNN features in a parameter-economic way. Li *et al.* [7] replaced the traditional image-to-class measure with a local descriptor-based measure to capture differences at a finer spatial level. Wu *et al.* [9] proposed an object-aware long-short-range spatial alignment strategy to align discriminative semantic parts between query and support sets. Zhang *et al.* [43] explored cross-image object semantic relations to distinguish subtle feature differences. Wertheimer *et al.* [10] investigated the relationship between support-query pairs through local feature reconstruction.

Despite the promising results of the methods above, they either lack sufficient supervision signal for recognition or require complex structures to capture local relations. Moreover, there has been limited emphasis on systematically exploring the intrinsic structure, particularly the differentiation between background and foreground components, in FS-FGVR. Recently, several studies [12], [44] have emerged that explicitly address the reduction of background clutter at the pixel level, leading to notable performance improvements. However, learning background noise filters at the pixel level, particularly with low model capacity, poses a significant challenge. Unlike the above work, this paper first leverages object-specific information in prediction by employing the distillation technique on the probability distribution of salient regions, thereby enhancing the precision of object information for fine-grained prediction.

C. Saliency Detection

Saliency detection is a computer vision technique that identifies regions of an image or video most visually significant to a human observer. Over the past few years, saliency detection has proven effective in various downstream tasks such as person re-identification [45] and visual tracking [46]. Several works have recently adopted saliency detection in few-shot learning research. Specifically, Zhang *et al.* [11] employed a saliency detection model to separate foreground and background regions and generate additional samples by combining viable foreground-background. Wang *et al.* [12] proposed a foreground object transformation strategy. They extracted the object's foreground by employing a saliency detection model and generating additional samples by the transformation learned in the foreground. Zhao *et al.* [13], [14] designed a complementary attention mechanism guided

by saliency, which employs saliency detection signals to learn interpretable representation. In contrast to the above works, this article efficiently utilizes saliency detection signals from the viewpoint of knowledge distillation.

D. Knowledge Distillation

Knowledge distillation [47] aims to transfer dark knowledge from a pre-trained large model to a smaller model, reducing its computational resource occupation. Existing works mainly include unidirectional knowledge distillation and bidirectional knowledge distillation. The former approach [48]–[51] typically requires designing a solid network to serve as the teacher, generating soft labels to facilitate the student network's learning. However, obtaining strong teachers requires considerable effort, and negative transfer frequently arises due to the significant capacity gap between the teacher and student networks [52]. Zhang *et al.* [53] proposed a deep mutual learning paradigm in a bidirectional view to address these issues. Bidirectional knowledge distillation has drawn increasing attention in the community due to its simplicity and effectiveness [54]–[56]. Despite the demonstrated achievements in knowledge distillation, previous works have primarily focused on scenarios with large sample sizes. In situations with limited data, the likelihood of negative transfer increases considerably due to high inductive bias. Ma *et al.* [57] proposed partner-assisted learning, a two-stage learning scheme that transfers the good embedding space from the pre-trained partner to the few-shot learner. Zhou *et al.* [58] proposed a binocular mutual learning approach that benefits few-shot learners by incorporating global and local views. Ye *et al.* [59] proposed the LastShot framework, which enables the few-shot learner to perform comparably to the pre-trained model. Differently, this article focuses on the distillation of saliency prior under the limited sample in a peer-teaching manner rather than relying on dark knowledge provided by the strong model.

III. THE ADOPTED METHODOLOGY

This section presents a robust saliency-aware distillation model for few-shot fine-grained visual recognition. The proposed model includes two core components: saliency-aware guidance and representation highlight&summarize module. The details are explained as follows.

A. Problem Definition

FS-FGVR aims to learn general feature representation on large-scale sub-categories with annotations that can be generalized to unseen sub-categories. To simulate this scenario, the dataset \mathcal{D} is typically divided into base set $\mathcal{D}_{\text{base}}$ and novel set $\mathcal{D}_{\text{novel}}$ by category, where $\mathcal{D}_{\text{base}} \cap \mathcal{D}_{\text{novel}} = \emptyset$. $\mathcal{D}_{\text{base}}$ represents the prior knowledge while $\mathcal{D}_{\text{novel}}$ represents the novel knowledge. For the typical FS-FGVR setup, models are trained (or tested) on N -way K -shot episodes sampled from $\mathcal{D}_{\text{base}}$ (or $\mathcal{D}_{\text{novel}}$). Each episode consists of a support set $\mathcal{D}^{\mathcal{S}}$ and a query set $\mathcal{D}^{\mathcal{Q}}$, sampled from the same N class. The model updates its parameters using the support set $\mathcal{D}^{\mathcal{S}}$ consisting of K samples per class, where K is set to 1 or 5 in few-shot setups. Moreover, the generalization performance of the

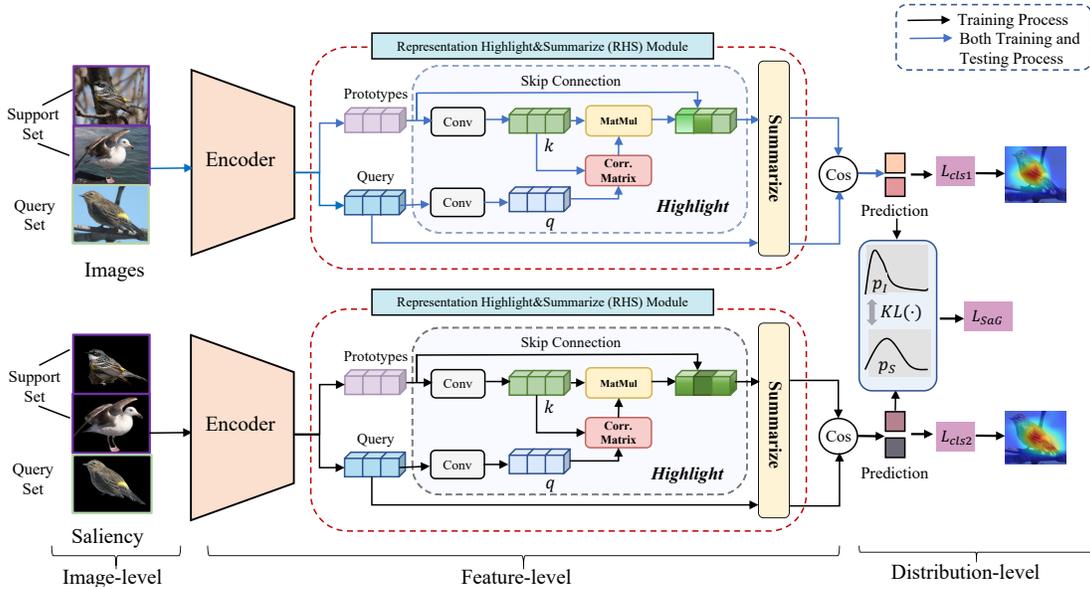


Fig. 2: The framework of the proposed RSaD for few-shot fine-grained visual recognition. This framework consists of three hierarchical levels of operations. At the image level, the framework generates the saliency serving as input for the symmetric branch. At the feature level, it highlights crucial features while aggregating significant information. Finally, the two branches independently optimize the cross-entropy (CE) loss while simultaneously providing complementary signals at the distribution level via mutual learning.

updated model is evaluated using the query set \mathcal{D}^Q . To be more specific, the general training objective can be expressed as

$$L^* = \arg \min_L \sum_{(x,y) \in \mathcal{D}_{\text{base}}^S} \ell(L(x; \mathcal{D}_{\text{base}}^S), y), \quad (1)$$

where L refers to the few-shot learner, whereas x and y denote the samples and corresponding labels from $\mathcal{D}_{\text{base}}^S$, the support set in the base set. Additionally, ℓ represents the loss function, i.e., cross-entropy loss.

B. Saliency-aware Guidance (SaG) for FS-FGVR

Backgrounds typically have a negative impact on FS-FGVR. The backgrounds probably amplify the variance within a class and diminish the variance between classes, which misleads the learning of inherent discriminative feature embeddings. Since the saliency detection model effectively highlights image regions associated with human visual attention, it is natural to utilize these models to provide crucial supplementary supervision in the meta-training phase of few-shot learners. To this end, this article presents the SaG strategy, which introduces additional saliency-aware supervision to guide the model toward focusing on the intrinsic discriminative features. SaG consists of two steps: augmented saliency generation and saliency-aware knowledge transfer.

1) *Augmented Saliency Generation*: To generate reliable saliency-aware supervision signals, it is crucial to ensure the quality of the generated saliency. Thus, this article employs multiple saliency detection models for image preprocessing and incorporates the resultant foreground as auxiliary data during model training. As shown in Fig. 3, two saliency detection models, BAS-Net [60] and U2-Net [61] are ensembled. Both

are pre-trained on the DUTS-TR [62] dataset, which does not overlap with FS-FGVR benchmarks. The models $h_1(\cdot)$ and $h_2(\cdot)$ process the given image I and produce saliency maps with pixel values ranging from [0, 1], where a higher value indicates a more critical region. Next, binarization operations are performed on the two maps, followed by an OR operation to produce a binary mask. Then, the saliency of the image can be obtained by taking the Hadamard product of the image and the mask, as shown below:

$$F_I = I \odot \text{Mask}(I), \quad (2)$$

where F_I represents the saliency prior. And $\text{Mask}(I)$ represents the mask of the image I calculated by the following formula:

$$\text{Mask}(I) = \sigma(h_1(I)) \mid \sigma(h_2(I)), \quad (3)$$

where the binary operation represented by \mid denotes logical OR, while the binarization process implemented by the activation function is denoted by σ . The activation function σ is defined as follows:

$$\sigma(h(I)_{ij}) = \begin{cases} 1, & \text{if } h(I)_{ij} \geq t \\ 0, & \text{otherwise} \end{cases}, \quad (4)$$

where the value of the saliency map matrix in row i and column j is denoted by $h(I)_{ij}$. The masked threshold, denoted as t , is set to 0.5 in this article. Notably, the cost associated with generating auxiliary data is acceptable as it involves a one-time consumption cost.

2) *Saliency-aware Knowledge Transfer*: In order to leverage the rich prior knowledge, knowledge distillation is a natural way. Moreover, according to the Deutsch–Norman late-selection filter model, critical information selection occurs exclusively after a comprehensive analysis of all inputs at

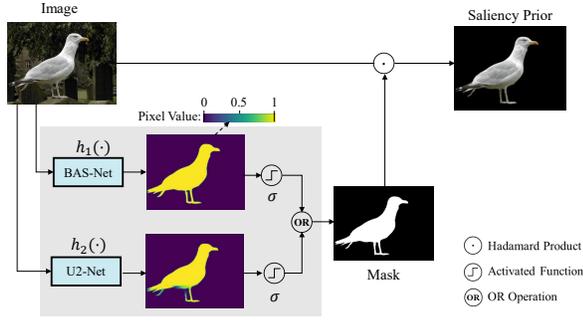


Fig. 3: Augmented Saliency Generation. For the input image, this module generates multiple saliency maps. Next, binarization operations are performed on these maps, followed by an OR operation. Then, the priors are synthesized through the Hadamard product of the input image and mask.

a higher level. Inspired by this, this article provides high-level supervision to enable the model to concentrate on objects during prediction, effectively mitigating the adverse effects of background clutter. As illustrated in Fig. 2, this paper introduces a symmetric structure of the main branch. Considering the saliency as the input for the additional branch, the resulting feature representation mapped to the latent space distinctly assumes an object-focused nature. While this feature exhibits an object-focused characteristic, transferring guidance from the feature dimension poses challenges, particularly when data availability is limited. Hence, this paper proposes to offer guidance from the standpoint of prediction distribution. The model prioritizes object relationships by minimizing the discrepancy between the distribution generated from raw images p_I and that produced from saliency p_S . The discrepancy between probability distributions can be measured using KL divergence as follows:

$$D_{KL}(p_I||p_S) = \sum_i p_I(i) \log \frac{p_I(i)}{p_S(i)}. \quad (5)$$

Nevertheless, employing traditional unidirectional supervision may not be optimal due to the large gap in representation between the saliency and the raw data and the potential risk of error amplification. After carefully considering these factors, this article utilized the deep mutual learning paradigm [53] to align the probabilities between the image and foreground bidirectionally. The loss function for saliency-aware guidance can be defined as follows:

$$\mathcal{L}_{SaG} = D_{KL}(p_I||p_S) + D_{KL}(p_S||p_I). \quad (6)$$

Notably, the article applies the same data augmentation to both images and their corresponding prior before feeding them into the network to ensure that the input difference remains within an acceptable range and maintains consistency in the probability.

C. Representation Highlight&Summarize (RHS) Module

The saliency-aware guidance described in the previous section facilitates the network to learn object-focused information

within the distribution perspective. However, it provides limited guidance for attention towards the representation produced by the backbone network. Moreover, the quality of representation also significantly impacts the effectiveness of guidance. To provide robust, deep information for fine-grained prediction, the RHS module is devised to simultaneously highlight and summarize crucial patterns within the representation level across multiple objects.

1) *Representation Highlight*: Due to significant differences in posture and scale, the global representation obtained by common FSL may not be sufficient for FS-FGVR. Inspired by previous work [9], [44], different sets' relationships are established to obtain informative representation. As illustrated in Fig. 2, the encoder $f_\theta(\cdot)$ produces high-level semantic representations for all images in the N -way K -shot episode $\mathcal{T}=(\mathcal{D}^S, \mathcal{D}^Q)$. Next, the semantic representations of each class in \mathcal{D}^S are aggregated to compute the prototypes by the following formula:

$$pt_i = \frac{1}{|\mathcal{D}_i^S|} \sum_{x \in \mathcal{D}_i^S} f_\theta(x), \quad (7)$$

where pt_i represents the prototype of the i -th class and \mathcal{D}_i^S refers to the i -th class within the support set. Then, the RHS module applies 1×1 convolution to project the prototype feature map pt and query feature map f^q to key k and query q , respectively, to filter out redundant information. Then, the semantic relation between k and q is explored. The relation matrix M between q and k can be calculated using the following formula:

$$M_{ij}(q, k) = \text{Softmax}(\langle q_i^T, k_j \rangle), \quad (8)$$

where $q_i \in \mathbb{R}^{c \times 1}$ and $k_j \in \mathbb{R}^{c \times 1}$ are the local descriptor of q and k , respectively, and $M \in \mathbb{R}^{hw \times hw}$. \langle, \rangle denotes the metric function. This article chooses cosine similarity as the metric function due to its superior ability to measure local representation similarity.

The prototype feature representation can be highlighted by computing a linear combination of the first-order and second-order representations by using the following equation:

$$\bar{pt} = \hat{k} \cdot M + \hat{pt}, \quad (9)$$

where \bar{pt} denotes the refined prototype feature representation. The first-order representation is represented by $\hat{pt}^s \in \mathbb{R}^{c \times hw}$, which is the spatially vectorized form of pt . The second-order representation is obtained by rearranging the first-order representation using a relation matrix M .

2) *Representation Summarize*: Existing research is mainly based on high-dimensional local descriptors, which may present a potential overfitting concern. Considering the potential risk for overfitting and computational efficiency, the RHS summarizes the highlighted information into contextual embedding. The RHS employs a combination of global max pooling and global average pooling to aggregate features while retaining the dominant feature, which can be expressed as

$$\begin{cases} \tilde{pt} = \frac{1}{h \times w} \sum_{i \in hw} \bar{pt}_i + \max_{i \in hw} \bar{pt}_i, \\ \tilde{f}^q = \frac{1}{h \times w} \sum_{i \in hw} f_i^q + \max_{i \in hw} f_i^q, \end{cases} \quad (10)$$

where \tilde{p}^t and \tilde{f}^q are the final refined representation of prototype feature representation p^t and query feature representation f^q .

D. Overall Learning Objective

Recalling the objectives for an FS-FGVR algorithm, the optimization problem arises as

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{cls1}} + \mathcal{L}_{\text{cls2}} + \alpha \mathcal{L}_{\text{SaG}}, \quad (11)$$

where $\mathcal{L}_{\text{cls1}}$ and $\mathcal{L}_{\text{cls2}}$ are the classification losses for two branches, respectively, calculated as follow:

$$P(y = c | x; \mathcal{D}^S) = \frac{\exp(\langle \tilde{f}^q, \tilde{p}^t_c \rangle)}{\sum_{n \in N} \exp(\langle \tilde{f}^q, \tilde{p}^t_n \rangle)}, \quad (12)$$

$$\mathcal{L}_{\text{cls}} = \sum_{(x,y) \in \mathcal{D}^S} -y \log P(y = c | x; \mathcal{D}^S), \quad (13)$$

where \langle, \rangle is the metric function. The article selects cosine similarity as the metric function in the probability computation to ensure consistency with the metric space of the RHS module. α is the hyper-parameter, which controls the interaction degree of saliency-aware guidance.

IV. PERFORMANCE EVALUATION

The benchmarks are introduced in Sec. IV-A, followed by the detailed implementation of the proposed approach in Sec. IV-B. The results of comparison experiments with different methods are presented in Sec. IV-C, IV-D, and IV-E, respectively.

A. Datasets

This article evaluates the performance of the proposed approach on three commonly used fine-grained benchmarks, including CUB-200-2011 [63], Stanford Dogs [64] and Stanford Cars [65]. Details of those benchmarks are described below:

- **CUB-200-2011** consists of 11,788 images depicting 200 species of birds. Following the setup of [44], the article randomly assigns 100, 50, and 50 classes to the training, validation, and testing. To ensure a fair comparison, none of the methods compared in this study utilize the bounding-box annotations provided by this benchmark.
- **Stanford Dogs** contains 20,580 images depicting 120 species of dogs. Following the setup of [44], this article randomly ad assigns 70, 20, and 30 classes for training, validation, and testing, respectively.
- **Stanford Cars** comprises 16,185 images depicting 196 car models. Following the setup of [44], this article randomly assigns 130, 17, and 49 classes to the training, validation, and testing, respectively.

B. Implementation Details

Network architecture: Following the previous research [44], this article employs ResNet-12 as the primary backbone to evaluate model performance. ResNet-12 consists of four residual blocks, each with three convolutional layers utilizing a 3×3 kernel and a 2×2 max-pooling layer. Unlike previous studies [20], [58] that employed wider-width ResNet-12 with a drop-block as a regularizer, this study uses a ResNet-12 with filter numbers 64-128-256-512 without introducing the drop-block to ensure a fair comparison. Moreover, the model's performance on the commonly-used shallower backbone Conv4 in FS-FGVR is discussed in Section IV-E.

Pre-training Stage: Some recent work [66]–[68] observed that pre-trained models for whole-classification have improved the transferability of novel classes in meta-learning models. Currently, almost all state-of-the-art methods [9], [20], [43] utilize a two-stage approach involving pre-training and episodic training for training in the context of FS-FGVR tasks. This article follows the pre-training setting in [68], [69]. During the pre-training stage, a fully-connected layer is appended at the end of the backbone for whole-class classification. All the images apply standard data augmentation, including random crop and random flip. For ResNet-12, the study utilizes the SGD optimizer with Nesterov acceleration, incorporating a learning rate of 0.001, weight decay of 0.0005, and momentum of 0.9. The model is trained from scratch for 300 epochs, with the learning rate decreasing by 0.1 at epoch 75 and epoch 150. As for Conv4, this study employs the SGD optimizer with a learning rate of 0.001, momentum of 0.9, and weight decay of 0.0005. The model is trained from scratch for 200 epochs, with the learning rate decreasing by 0.1 at epoch 85 and epoch 170.

Episodic Training Stage: After pre-training, the pre-trained model removed the last fully connected layer and fine-tuned the model following the standard meta-learning scheme. For CUB-200-2011 and Stanford Dogs under ResNet-12 backbone, this article utilizes the Adam optimizer with an initial learning rate of 0.001 during the episodic training stage and optimizes the model over 40,000 episodes, except for the 5-way 1-shot setting, which involves 60000 episodes. For the rest of the cases, this article employed the AdamW optimizer with an initial learning rate of 0.001 during the episodic training stage. This article optimized the proposed model over 40000 episodes in a 15-way 5-shot setting, evaluating its performance in both 1-shot and 5-shot settings for training stability.

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

The following experiments present a comparison of the proposed network with state-of-the-art approaches, comprising four generic FSL methods (ProtoNet [40], DN4 [7], Baseline++ [66], CAN [70]) and five specialized FS-FGVR methods (TOAN [8], BSNet [41], OLSA [9], AGPF [71] and BSFA [44]). The performance evaluation is conducted on three fine-grained benchmarks mentioned in Sec. IV-A. The performance evaluations of the proposed framework on different benchmarks are presented in Sec. IV-C1, Sec. IV-C2, and Sec. IV-C3, respectively.

TABLE I: This table reports the accuracy (%) of 5-way 1-shot and 5-shot tasks on three popular benchmarks using ResNet-12. The accuracy refers to the mean value with 95% confidence intervals on 600 episodes. The best results are presented in bold. * indicates that those results are obtained from the original paper, while † denotes that the results are reported in [44]. The remaining results are reproduced using the open-source code under the same experimental settings.

Method	Backbone	CUB-200-2011		Stanford Dogs		Stanford Cars	
		1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
ProtoNet [40]	ResNet-12	73.52 ± 0.98	87.57 ± 0.52	61.46 ± 0.94	83.03 ± 0.58	75.46 ± 0.89	91.18 ± 0.41
DN4 † [7]	ResNet-12	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++ [66]	ResNet-12	72.07 ± 0.91	86.16 ± 0.54	64.63 ± 0.82	82.37 ± 0.56	76.14 ± 0.82	91.15 ± 0.40
CAN † [70]	ResNet-12	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* [8]	ResNet-12	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† [41]	ResNet-12	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* [9]	ResNet-12	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† [71]	ResNet-12	78.73 ± 0.84	89.77 ± 0.47	72.34 ± 0.86	84.02 ± 0.57	85.34 ± 0.74	94.79 ± 0.35
BSFA [44]	ResNet-12	82.22 ± 0.85	90.49 ± 0.47	69.62 ± 0.92	82.50 ± 0.58	89.42 ± 0.68	95.36 ± 0.36
RSaD	ResNet-12	82.45 ± 0.79	92.02 ± 0.44	73.75 ± 0.93	86.65 ± 0.54	87.27 ± 0.70	95.01 ± 0.49

1) *Results on the CUB-200-2011 Dataset:* As presented in Table I, the RSaD outperforms FSL methods, and FS-FGVR methods on the CUB-200-2011 dataset. Specifically, RSaD achieves comparable performance with state-of-the-art method BSFA in 5-way 1-shot and 5-way 5-shot tasks, respectively. This indicates its ability to capture class correlations at various granularities effectively. One key factor contributing to the superior performance of RSaD lies in its ability to provide saliency guidance and region relations, resulting in better embedding spaces. In contrast to the proposed approach, BSFA incorporates a crop function aimed at discerning the foreground of images in both the training and testing phases. The crop function learned under the ResNet-12 backbone may be biased, potentially contributing to BSFA’s inferior performance relative to the proposed approach.

2) *Results on the Stanford Dogs Dataset:* In order to evaluate the capability of RSaD in handling more challenging FS-FGVR tasks, this section further presents its performance on the more complex dataset Stanford Dogs. As shown in Table I, RSaD outperforms state-of-the-art FS-FGVR approaches by a substantial margin in the Stanford Dogs dataset. Specifically, RSaD achieves 4.54%, 1.03%, 9.22%, 11.42% and 23.6% performance gain for the 1-shot setting and 3.45%, 2.06%, 7.80%, 6.46% and 16.79% performance gain for the 5-shot setting compared with the BSFA and four other FS-FGVR methods, respectively. An explanation for the remarkable performance of RSaD could be its ability to mitigate noise in complex environments, such as the background. This enables the model to allocate increased attention toward the intrinsic dissimilarities within the image.

3) *Results on the Stanford Cars Dataset:* To further evaluate the efficacy of the RSaD, this article conducted experiments on the simpler benchmark, Stanford Cars. As depicted in Table I, RSaD achieves competitive results compared to state-of-the-art methods. Specifically, Our method outperforms all typical FSL methods and some FS-FGVR methods under 1-shot and 5-shot. However, RSaD performs slightly worse than the state-of-the-art methods BSFA by 2.15% for 1-shot,

TABLE II: Comparison with other few-shot knowledge distillation based methods and saliency-guided approach on CUB-200-2011. The best results are shown in bold. * indicates that those results are obtained from the original paper.

Method	Backbone	CUB-200-2011	
		5-way 1-shot	5-way 5-shot
FOT* [12]	ResNet-18	80.40	89.68
SGCA* [14]	ResNet-12	79.84 ± 0.42	90.91 ± 0.22
SGCA++* [13]	ResNet-12	80.69 ± 0.42	91.43 ± 0.22
Protonet [40]	ResNet-12	73.52 ± 0.98	87.57 ± 0.52
Meta-Baseline [68]	ResNet-12	76.49 ± 0.87	87.86 ± 0.52
BML* [58]	ResNet-12	76.21 ± 0.63	90.45 ± 0.36
LST-Proto* [59]	ResNet-12	75.80 ± 0.21	90.22 ± 0.12
LST-FEAT* [59]	ResNet-12	80.20 ± 0.21	91.49 ± 0.12
RSaD	ResNet-12	82.45 ± 0.79	92.02 ± 0.44

and 0.35% for 5-shot, respectively. Their superior performance may be attributed to their additional structure and spatial descriptor comparison. Considering the trade-off between computational efficiency and performance, RSaD avoids incorporating excessive additional structure. Furthermore, RSaD employs spatial descriptor aggregation to avoid cost-expensive spatial descriptor comparison. Undoubtedly, RSaD provides significant computational efficiency advantages, as shown in Sec. X for details.

D. Comparison With Saliency-Guided Approaches and KD-Based Approaches

As the proposed approach is a hybrid of saliency-guided and knowledge-distillation (KD) based approaches, this subsection compares RSaD with the saliency-guided and KD-based methods. The comparison results are presented in Sec. IV-D1 and Sec. IV-D2, respectively.

1) *Comparison With Saliency-Guided Approaches:* This article compares RSaD with three other saliency-guided methods: SGCA [14], SGCA++ [13] and FOT [12]. SGCA and SGCA++ obtain transferable representations through learning

TABLE III: This table reports the accuracy (%) of 5-way 1-shot and 5-shot tasks on three popular benchmarks using Conv4. The best results are presented in bold. * indicates that those results are obtained from the original paper. The remaining results are reproduced using the open-source code under the same experimental settings.

Method	Backbone	CUB-200-2011		Stanford Dogs		Stanford Cars	
		1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
ProtoNet [40]	Conv4	51.78 ± 0.93	74.72 ± 0.69	41.60 ± 0.82	56.98 ± 0.75	46.38 ± 0.79	64.58 ± 0.67
LRPABN* [72]	Conv4	63.63 ± 0.77	76.06 ± 0.58	45.72 ± 0.75	60.94 ± 0.66	60.28 ± 0.76	73.29 ± 0.63
MattML* [73]	Conv4	66.29 ± 0.56	80.34 ± 0.30	54.84 ± 0.53	71.34 ± 0.38	66.11 ± 0.54	82.80 ± 0.28
SoSN* [74]	Conv4	64.56 ± 0.91	77.82 ± 0.57	48.21 ± 0.72	63.15 ± 0.67	62.88 ± 0.72	76.10 ± 0.28
TOAN* [8]	Conv4	65.34 ± 0.75	80.43 ± 0.60	49.30 ± 0.77	67.16 ± 0.49	65.90 ± 0.72	84.24 ± 0.48
BSFA [44]	Conv4	59.40 ± 0.97	74.42 ± 0.62	49.13 ± 0.84	63.27 ± 0.73	56.06 ± 0.89	73.28 ± 0.68
MLSO* [75]	Conv4	68.21 ± 0.78	82.18 ± 0.47	55.62 ± 0.58	71.98 ± 0.71	67.83 ± 0.63	84.83 ± 0.48
RSaD	Conv4	71.15 ± 0.92	84.03 ± 0.62	59.42 ± 0.95	75.30 ± 0.69	65.43 ± 1.29	81.75 ± 0.88

saliency-guided attention, while FOT synthesizes additional samples using a saliency map matching strategy. Table II presents the results. Compared to SGCA, RSaD achieves 2.61% and 1.11% performance gain in 1-shot and 5-shot tasks under backbone ResNet-12. Furthermore, RSaD outperforms SGCA, which is a refined version of SGCA. Compared to FOT, RSaD achieves 3.43% and 2.45% performance gain in 1-shot and 5-shot tasks under backbone ResNet-18. The superior performance of RSaD shows that the salient features obtained by RSaD are beneficial for learning a better embedding space.

2) *Comparison With KD-Based FSL Approaches:* This article compared RSaD with several recently published KD-based FSL approaches, including Meta-Baseline [68], BML [58] and LST [69]. Meta-Baseline employs a simple process of meta-learning over a whole-classification pre-trained model. BML utilizes mutual learning to align the feature distributions between local and global views. LST proposes a general learning strategy that allows the model to learn from the pre-trained strong classifier. Both their method and the proposed method are based on the typical FSL approach ProtoNet [40], resulting in improved learning of the embedding space. LST-FEAT is an improved version of LST-Proto that applies LST with self-attention to the ProtoNet. Performance comparison is shown in Table II. As can be seen, knowledge distillation can significantly improve the performance of FSL methods. The experiment results indicate that saliency prior may offer higher quality dark knowledge for knowledge distillation than using pre-trained models, potentially contributing to our model’s superiority over other approaches. Moreover, even if pre-trained dark knowledge is optimized, unidirectional knowledge distillation may not be optimal in the 1-shot scenario.

E. Few-shot Fine-Grained Classification with Shallow Backbone Network

This article conducts experiments on shallower backbone Conv4 to further evaluate the model performance. Conv4 is composed of four blocks, each containing a 3×3 convolution layer, a batch normalization layer, a ReLU activation layer, and a 2×2 max-pooling layer. The experiments provide a thorough comparison of the proposed network with state-of-the-art approaches, including three FSL methods (ProtoNet [40],

SoSN [74], MLSO [75]), as well as four specialized FS-FGVR methods (LRPABN [72], MattML [73], TOAN [72], BSFA [44]) using a Conv4 backbone. The results in Table III demonstrate that the proposed approach has a performance advantage over shallow backbones. Due to the low capacity of Conv4, BSFA fails to generate optimal foreground object coordinates for background suppression, resulting in significant performance degradation. However, our proposed approach maintains superior performance despite the low-capacity backbone due to effective probability distribution comparison in a lower-dimensional space determined by the number of ways. Overall, RSaD consistently performs well on multiple backbones, thus providing strong evidence of its effectiveness.

V. ANALYSIS AND DISCUSSION

This section presents a series of ablation studies and visualizations to investigate each module’s significance. Sec. V-A introduces the module-wise ablation experiment, while Sec. V-B, V-C, and V-D demonstrate the impact of saliency-aware guidance and RHS module. Additionally, Sec. V-E provides an analysis of the model complexity.

A. Module-Wise Ablation Study

Table IV displays the results of module-wise ablation experiments performed on ResNet-12. The baseline is a model obtained through meta-learning on a whole-classification pre-trained model. RHS refers to the representation highlight&summarize module, while SaG stands for saliency-aware guidance strategy. The results in Table IV demonstrate the effectiveness of the SaG and RHS module on the FS-FGVR task. SaG gains more significant improvement than the SaG module, indicating that prior saliency provides more valuable information for optimizing the model than second-order features. Combining SaG and RHS modules further improves the model’s performance. The main reason is that the RHS module optimizes representation and results in higher-quality supervision. In this way, the model leverages such supervision more effectively, improving overall performance.

TABLE IV: Ablation experiments of the proposed method on CUB-200-2011 and Stanford Dogs datasets under backbone ResNet-12. The best results are highlighted in bold.

SaG	RHS	CUB-200-2011		Stanford Dogs	
		1-shot	5-shot	1-shot	5-shot
		77.58±0.90	88.24±0.52	69.40±0.88	84.42±0.55
✓		80.28±0.87	91.18±0.46	72.03±0.87	85.42±0.52
	✓	78.50±0.86	88.87±0.52	72.66±0.87	86.13±0.52
✓	✓	82.45±0.79	92.02±0.44	73.75±0.93	86.65±0.54

B. Discussion on Impact of SaG strategy

The SaG strategy plays a vital role in learning discriminative features in the proposed method. This strategy involves two primary stages: augmented saliency generation and saliency-aware knowledge transfer. This subsection conducts experiments and Grad-CAM [76] visualization on the CUB-200-2011 dataset to investigate the impact of in SaG strategy.

1) *Impact of Augmented Saliency Generation*: This article integrates two distinct saliency detection models (U2-Net [61] and BASNet [60]) using an OR operation within augmented saliency generation to produce saliency maps. This subsection conducts quantitative comparisons to explore the individual contributions of these saliency detection models. As shown in Table V, utilizing a single saliency detection model can yield high performance, illustrating that both models can effectively generate saliency from images. Furthermore, employing an ensemble strategy leads to consistently superior classification performance in all scenarios. This suggests that the ensemble approach offers a more dependable prior for model training, thus alleviating the performance constraints encountered by individual models in certain instances.

TABLE V: Quantitative evaluation of the augmented saliency generation strategy on CUB-200-2011 and Stanford Dogs dataset. The best results are in bold.

Prior	CUB-200-2011				Stanford Dogs			
	5-way	1-shot	5-way	5-shot	5-way	1-shot	5-way	5-shot
U2-Net	82.09±0.81	91.88±0.42	72.71±0.88	85.49±0.53				
BASNet	81.72±0.79	91.86±0.43	72.42±0.87	86.36±0.50				
Ensemble	82.45±0.79	92.02±0.44	73.75±0.93	86.65±0.54				

Furthermore, this article presents a visualization of representative examples to demonstrate the efficacy of the ensemble strategy. As illustrated in Fig. 4, BASNet may produce incomplete segmentation results for the whole object, while U2-Net typically performs robustly. This inconsistency may originate from BASNet’s reliance on a predict-refine architecture and the potential loss of local information due to dilated convolution [77], [78]. Conversely, BASNet tends to provide comprehensive segmentation in scenarios with occlusion, attributable to its utilization of hybrid boundary loss. Hence, the ensemble of the two models can produce more accurate saliency maps for further training.

2) *Impact of Saliency-aware Knowledge Transfer*: To further assess the impact of the SaG strategy, this article conducts Grad-CAM visualizations. Figs.5(b) and 5(c) visualize the

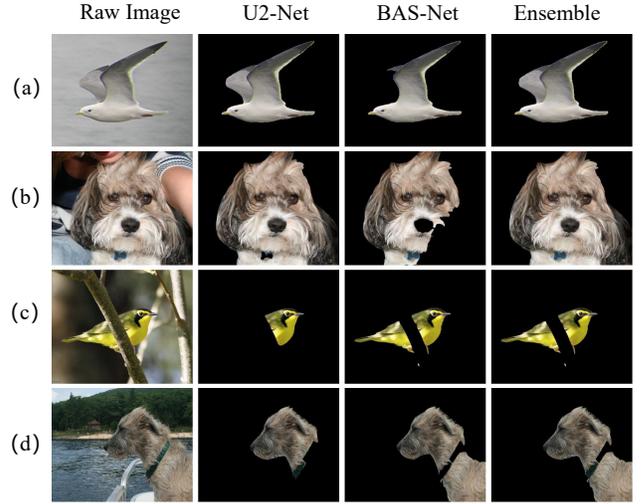


Fig. 4: Visual comparison of the ensemble strategy with the individual model. The saliency produced by the ensemble model is more precise and more reliable than others.

class activation regions on the baseline and the proposed approach. The baseline model is obtained through meta-learning on a pre-trained classification model without using saliency as auxiliary information. The RHS module is excluded from the proposed framework in this study to eliminate any potential influence of the RHS module. Fig. 5 illustrates that the class activation map generated by our proposed approach primarily focuses on the object itself while paying less attention to the background. In contrast, the baseline approach focuses on the part of the background in addition to the object itself. This indicates that introducing the SaG strategy can reduce interference from the background and capture the critical regions of the different sub-class. Thus, the robustness and generalization capabilities of the proposed model are improved.

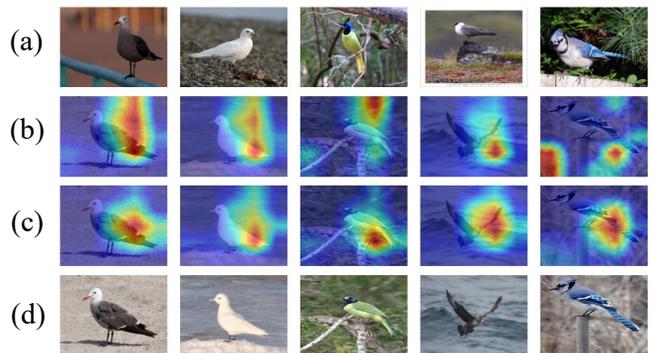


Fig. 5: The visualization of local regions obtained from Grad-CAM. (a) prototype, (b) baseline, (c) our RSaD w/o RHS, and (d) query. The Grad-CAM Map identifies significant regions in the input image that affect classification decisions, where darker regions indicate higher importance. The specimens within each column belong to the same subclass, while those in different columns belong to distinct classes. Compared with baseline, RSaD pays more attention to the distinguishing characteristics of objects themselves.

C. Discussion on Impact of RHS module

The RHS module is similar to the regular cross-attention module but lacks the value projection. To evaluate the effectiveness of the RHS module, this section conducts comparative experiments on the CUB-200-2011 and Stanford Dogs datasets. This study substituted the highlighted sub-module in the RHS with a cross-attention module. The results in Table VI indicate that the RHS approach outperforms the cross-attention mechanism, particularly in the 1-shot scenario. It is reasonable to assert that incorporating a value projection has the potential to introduce inconsistencies with the key projection, thereby leading to inaccurate alignment between prototypes and queries. Additionally, omitting a value projection enhances parameter efficiency, resulting in performance gain in the few-shot setup.

TABLE VI: Comparison results of RHS module and cross attention. CA represents the cross-attention. The best results are presented in bold.

Backbone	CUB-200-2011		Stanford Dogs	
	1-shot	5-shot	1-shot	5-shot
RSaD w/o RHS	80.28±0.87	91.18±0.46	72.03±0.87	85.42±0.52
RSaD w/ CA	81.58±0.81	91.53±0.40	72.47±0.84	86.50±0.49
RSaD w/ RHS	82.45±0.79	92.02±0.44	73.75±0.93	86.65±0.54

The DBI (Davies-Bouldin Index) is employed as a quantitative metric to assess the quality of model embedding space. It is defined as follows:

$$DBI = \frac{1}{n} \sum_{i=1}^n \max_{j \neq i} \left(\frac{\sigma_i + \sigma_j}{d(c_i, c_j)} \right), \quad (14)$$

where n represents the number of samples, c_i denotes the cluster of the i -th sub-class, σ_i signifies the average distance from the sample to the cluster in the i -th sub-class, and $d(c_i, c_j)$ represents the distance between different clusters. A smaller DBI value indicates a better clustering effect. As depicted in Table VII, applying RHS module can significantly improve the embedding space and the generalization performance in unseen subclasses.

To further confirm the efficacy of the RHS module, t-SNE [79] visualization is performed. According to the results shown in Fig. 6, two observations are highlighted. Firstly, introducing the RHS module can significantly reduce the intra-class distance, resulting in a tighter cluster of objects within the same category. This leads to an improved discriminative capability of the model. Secondly, applying the RHS module can effectively enhance the model’s generalization ability, resulting in consistently high performance in the novel set.

D. Discussion on Impact of KD

Considering the adoption of a symmetric structure and bidirectional distillation method in RSaD, it is reasonable to perform quantitative experiments to examine the effects of various distillation methods. Table VIII presents the comparative results of various distillation methods, incorporating both directional and structural asymmetry. The ”UD-KD” model

TABLE VII: Quantitative experiment on the effect of RHS module. The terms ”Base” and ”Novel” indicate testing in the Base set and Novel set, respectively. The best results are presented in bold.

Method	Backbone	Split	DBI ↓
RSaD w/o RHS	ResNet-12	Base	3.7632
RSaD w/ CA	ResNet-12	Base	3.0955
RSaD w/ RHS	ResNet-12	Base	2.9593 (↓ 0.8102)
RSaD w/o RHS	ResNet-12	Novel	3.3140
RSaD w/ CA	ResNet-12	Novel	2.9642
RSaD w/ RHS	ResNet-12	Novel	2.8374 (↓ 0.4766)

is the unidirectional knowledge distillation model. The loss function is defined as follows:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{cls} + \alpha * D_{KL}(p_I || p_S), \quad (15)$$

where \mathcal{L}_{cls} is the loss function calculated by the cross-entropy, and $D_{KL}(\cdot, \cdot)$ represents the KL divergence. Here, p_I denotes the logit distribution generated by the main branch, while p_S represents the corresponding saliency logit distribution generated by the auxiliary branch. The auxiliary branch in ”UD-KDP” is pre-trained, and its weights are frozen during training. The results demonstrate that RSaD outperforms the UD-KD model, suggesting that the bidirectional knowledge distillation strategy is better suited for distilling imperfect knowledge. Moreover, this study explores the performance of three different asymmetric branches on the CUB-200-2011 dataset: ResNet-12-Conv4 (strong-weak), ResNet-12-Visformer-tiny [80] (similar parameters but different structure), and ResNet-12-ResNet-18 (weak-strong). Employing a symmetric branch structure demonstrates superior performance compared to using an asymmetric branch structure.

TABLE VIII: The effect of different mutual information distillation approaches. The best results are presented in bold.

Method	Backbone	CUB-200-2011	
		5-way 1-shot	5-way 5-shot
UD-KD	ResNet-12/ResNet-12	80.55±0.81	89.57±0.48
UD-KDP	ResNet-12/ResNet-12	81.02±0.82	91.00±0.45
RSaD	ResNet-12/Conv4	81.70±0.81	91.60±0.40
RSaD	ResNet-12/Visformer	79.50±0.84	91.47±0.42
RSaD	ResNet-12/ResNet-18	79.72±0.82	90.88±0.47
RSaD	ResNet-12/ResNet-12	82.45±0.79	92.02±0.44

A parameter sensitivity analysis of the importance factor α is conducted to further explore the impact of saliency-aware supervision signal. According to the results shown in Table IX, two insights are gained. Firstly, The degree of saliency-aware supervision signal is significant. Excessive saliency-aware supervision signal causes the model to overly focus on fitting the foreground distribution instead of the label distribution, while insufficient mutual information results in the model degrading to baseline performance. Secondly, the degree of saliency-aware supervision signal varies across datasets. Considering the confidence of mutual information, less saliency-aware supervision signal may be more suitable for challenging datasets.

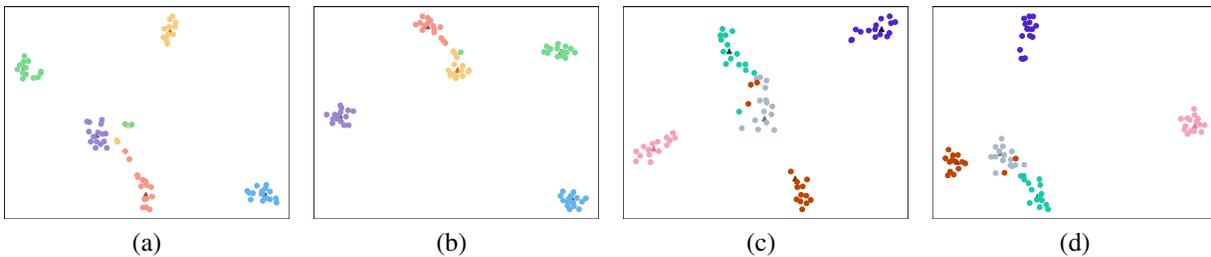


Fig. 6: t-SNE visualization results for 5-way 1-shot tasks of different models on various sets in Dataset CUB-200-2011. (a) RSaD w/o RHS module on base set, (b) RSaD w/ RHS module on base set, (c) RSaD w/o RHS module on novel set, (d) RSaD w/ RHS module on novel set. Triangles represent prototypes, circles represent queries, and different colors indicate distinct subclasses. The labels of the base set and the novel set are disjoint.

TABLE IX: Effect of the degree of the saliency-aware supervision. The best results are presented in bold.

α	Backbone	CUB-200-2011		Stanford Dogs	
		1-shot	5-shot	1-shot	5-shot
0.1	ResNet-12	78.95 \pm 0.83	89.12 \pm 0.50	71.78 \pm 0.91	85.32 \pm 0.55
1.0	ResNet-12	80.48 \pm 0.82	90.36 \pm 0.46	73.75\pm0.93	86.65\pm0.54
5.0	ResNet-12	82.45\pm0.79	92.02\pm0.44	72.65 \pm 0.92	86.58 \pm 0.55
10.0	ResNet-12	80.71 \pm 0.82	91.15 \pm 0.45	68.90 \pm 0.92	84.62 \pm 0.59

E. Model Complexity Analysis

The model complexity analysis is critical for model evaluation. This work presents a comparative analysis with four state-of-the-art open-source approaches to establish the superiority of the proposed approach. Two commonly used evaluation metrics, "Params.(M)" and "FLOPs(G)" are employed to analyze model complexity. Params.(M) refers to the number of parameters in the model, measured in millions. Meanwhile, FLOPs(G) denotes the number of arithmetic operations required to model feedforward, measured in billions.

As shown in Table X, the proposed approach is computationally inexpensive compared to state-of-the-art FG-FSL approaches such as BSFA and AGPF. There is a significant gap in FLOPs(G) despite having similar model parameters, with only half that of the other two approaches. Moreover, the proposed approach is inferior to the traditional FSL approach CAN in terms of model complexity due to the need for a detailed exploration of the relationship between different regions. Nevertheless, the proposed approach significantly improved over CAN in classification performance (see Table I). Furthermore, the proposed approach brings a negligible increase of 0.52 and 0.06 in Params.(M) and FLOPs(G), respectively. Three possible reasons account for this insignificant raise: 1) No spatial comparison in the proposed method during logit computation compared to the DN4 and BSFA. This article employs spatial descriptor aggregation followed by cost-effective channel comparison to improve computational efficiency. 2) No complex structure is introduced compared to AGPF. AGPF introduces a feature pyramid structure to capture the variations between sub-classes, which is computationally expensive, while the proposed method introduces a lightweight RHS module. 3) The proposed approach employs a dual-

branch network without an increase in model complexity. During the testing phase, only the main branch is preserved.

TABLE X: Comparison with state-of-the-art few-shot fine-grained visual recognition methods in model complexity. The best results are shown in bold.

Method	Backbone	Complexity	
		Params.(M)	FLOPs(G)
DN4 [7]	ResNet-12	12.42	67.39
CAN [70]	ResNet-12	8.04	12.75
AGPF [71]	ResNet-12	8.77	51.53
BSFA [44]	ResNet-12	8.04	50.64
Baseline	ResNet-12	8.00	24.80
RSaD	ResNet-12	8.52(\uparrow 0.52)	24.86(\uparrow 0.06)

VI. CONCLUSION

This article proposed RSaD for few-shot fine-grained visual recognition. RSaD leveraged saliency-aware guidance to improve model performance while maintaining computational efficiency. Specifically, RSaD integrated multiple saliency models to produce saliency maps of superior quality. Subsequently, RSaD introduced SaG, which leveraged the distillation technique on the salient region probability distribution to exploit explicit intrinsic relationships among sub-classes. RHS was designed to highlight deep descriptors relevant to the objects of interest and squeeze them into contextual embedding to ensure transferable representation. Extensive experiments demonstrated that the proposed model achieved comparable performance to the current state-of-the-art approach while maintaining low model complexity.

Recently, several methods [81]–[83] have been designed to segment objects in an image all at once. These efficient and promptable methods offer more robust saliency priors than saliency detection models. Future research in this article is not limited to few-shot fine-grained classification but will focus on developing a generalized few-shot visual recognition framework based on solid segmentation models. Furthermore, there have been significant breakthroughs in semi-supervised few-shot learning methods [84], [85] in recent years. Semi-supervised learning within FS-FGVR is a promising avenue worth investigating in future work.

REFERENCES

- [1] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2117–2125.
- [2] Y. Liufu, L. Jin, J. Xu, X. Xiao, and D. Fu, "Reformative noise-immune neural network for equality-constrained optimization applied to image target detection," *IEEE Transactions on Emerging Topics in Computing*, vol. 10, no. 2, pp. 973–984, 2021.
- [3] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=YicbFdNTTy>
- [4] Y. Ding, Z. Ma, S. Wen, J. Xie, D. Chang, Z. Si, M. Wu, and H. Ling, "AP-CNN: Weakly supervised attention pyramid convolutional neural network for fine-grained visual classification," *IEEE Transactions on Image Processing*, vol. 30, pp. 2826–2836, 2021.
- [5] C. Liu, H. Xie, Z. Zha, L. Yu, Z. Chen, and Y. Zhang, "Bidirectional attention-recognition model for fine-grained object classification," *IEEE Transactions on Multimedia*, vol. 22, no. 7, pp. 1785–1795, 2020.
- [6] J. He, J.-N. Chen, S. Liu, A. Kortylewski, C. Yang, Y. Bai, and C. Wang, "TransFG: A transformer architecture for fine-grained recognition," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 1, 2022, pp. 852–860.
- [7] 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 *CVPR*, 2019, pp. 7260–7268.
- [8] 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, 2022.
- [9] Y. Wu, B. Zhang, G. Yu, W. Zhang, B. Wang, T. Chen, and J. Fan, "Object-aware long-short-range spatial alignment for few-shot fine-grained image classification," in *ACM Multimedia*, 2021, pp. 107–115.
- [10] D. Wertheimer, L. Tang, and B. Hariharan, "Few-shot classification with feature map reconstruction networks," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 8012–8021.
- [11] H. Zhang, J. Zhang, and P. Koniusz, "Few-shot learning via saliency-guided hallucination of samples," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2770–2779.
- [12] C. Wang, S. Song, Q. Yang, X. Li, and G. Huang, "Fine-grained few shot learning with foreground object transformation," *Neurocomputing*, vol. 466, pp. 16–26, 2021.
- [13] L. Zhao, G. Liu, D. Guo, W. Li, and X. Fang, "Boosting few-shot visual recognition via saliency-guided complementary attention," *Neurocomputing*, vol. 507, pp. 412–427, 2022.
- [14] —, "Saliency-guided complementary attention for improved few-shot learning," in *2021 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 2021, pp. 1–6.
- [15] J. A. Deutsch and D. Deutsch, "Attention: Some theoretical considerations," *Psychological review*, vol. 70, no. 1, p. 80, 1963.
- [16] B. Sorscher, S. Ganguli, and H. Sompolinsky, "Neural representational geometry underlies few-shot concept learning," *Proceedings of the National Academy of Sciences*, vol. 119, no. 43, p. e2200800119, 2022.
- [17] S. Kong and C. Fowlkes, "Low-rank bilinear pooling for fine-grained classification," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 365–374.
- [18] 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 *CVPR*, 2018, pp. 1199–1208.
- [19] X. Yang, X. Nan, and B. Song, "D2N4: A discriminative deep nearest neighbor neural network for few-shot space target recognition," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 58, no. 5, pp. 3667–3676, 2020.
- [20] J. Xu, H. Le, M. Huang, S. Athar, and D. Samaras, "Variational feature disentangling for fine-grained few-shot classification," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 8812–8821.
- [21] S.-L. Xu, F. Zhang, X.-S. Wei, and J. Wang, "Dual attention networks for few-shot fine-grained recognition," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 3, 2022, pp. 2911–2919.
- [22] X.-S. Wei, Y. Shen, X. Sun, P. Wang, and Y. Peng, "Attribute-aware deep hashing with self-consistency for large-scale fine-grained image retrieval," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 11, pp. 13 904–13 920, 2023.
- [23] X.-S. Wei, Y.-Z. Song, O. M. Aodha, J. Wu, Y. Peng, J. Tang, J. Yang, and S. Belongie, "Fine-grained image analysis with deep learning: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 12, pp. 8927–8948, 2022.
- [24] K. Duan, D. Parikh, D. Crandall, and K. Grauman, "Discovering localized attributes for fine-grained recognition," in *2012 IEEE Conference on Computer Vision and Pattern Recognition*, 2012, pp. 3474–3481.
- [25] T. Berg and P. N. Belhumeur, "POOF: Part-based one-vs.-one features for fine-grained categorization, face verification, and attribute estimation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 955–962.
- [26] P.-H. Gosselin, N. Murray, H. Jégou, and F. Perronnin, "Revisiting the fisher vector for fine-grained classification," *Pattern recognition letters*, vol. 49, pp. 92–98, 2014.
- [27] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, pp. 91–110, 2004.
- [28] D. Chang, Y. Ding, J. Xie, A. K. Bhunia, X. Li, Z. Ma, M. Wu, J. Guo, and Y.-Z. Song, "The devil is in the channels: Mutual-channel loss for fine-grained image classification," *IEEE Transactions on Image Processing*, vol. 29, pp. 4683–4695, 2020.
- [29] T.-Y. Lin, A. RoyChowdhury, and S. Maji, "Bilinear cnn models for fine-grained visual recognition," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1449–1457.
- [30] S. Kong and C. Fowlkes, "Low-rank bilinear pooling for fine-grained classification," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 365–374.
- [31] N. Zhang, J. Donahue, R. Girshick, and T. Darrell, "Part-based R-CNNs for fine-grained category detection," in *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part I 13*. Springer, 2014, pp. 834–849.
- [32] X. Zhang, H. Xiong, W. Zhou, W. Lin, and Q. Tian, "Picking deep filter responses for fine-grained image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 1134–1142.
- [33] Z. Yang, T. Luo, D. Wang, Z. Hu, J. Gao, and L. Wang, "Learning to navigate for fine-grained classification," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 420–435.
- [34] Y. Han, X. Yang, T. Pu, and Z. Peng, "Fine-grained recognition for oriented ship against complex scenes in optical remote sensing images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–18, 2022.
- [35] C. Guo, Y. Lin, S. Chen, Z. Zeng, M. Shao, and S. Li, "From the whole to detail: Progressively sampling discriminative parts for fine-grained recognition," *Knowledge-Based Systems*, vol. 235, p. 107651, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0950705121009138>
- [36] Y. Hu, X. Jin, Y. Zhang, H. Hong, J. Zhang, Y. He, and H. Xue, "RAMS-Trans: Recurrent attention multi-scale transformer for fine-grained image recognition," in *Proceedings of the 29th ACM International Conference on Multimedia*, ser. MM '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 4239–4248. [Online]. Available: <https://doi.org/10.1145/3474085.3475561>
- [37] Y. Zhang, J. Cao, L. Zhang, X. Liu, Z. Wang, F. Ling, and W. Chen, "A free lunch from vit: Adaptive attention multi-scale fusion transformer for fine-grained visual recognition," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 3234–3238.
- [38] H. Sun, X. He, and Y. Peng, "SIM-Trans: Structure information modeling transformer for fine-grained visual categorization," in *Proceedings of the 30th ACM International Conference on Multimedia*, ser. MM '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 5853–5861. [Online]. Available: <https://doi.org/10.1145/3503161.3548308>
- [39] 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 Transactions on Image Processing*, vol. 28, no. 12, pp. 6116–6125, 2019.
- [40] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17. Red Hook, NY, USA: Curran Associates Inc., 2017, p. 4080–4090.

- [41] X. Li, J. Wu, Z. Sun, Z. Ma, J. Cao, and J. Xue, "BSNet: Bi-similarity network for few-shot fine-grained image classification," *IEEE Trans. Image Process.*, vol. 30, pp. 1318–1331, 2021.
- [42] R. T. Q. Chen, X. Li, R. Grosse, and D. Duvenaud, "Isolating sources of disentanglement in vaes," ser. NIPS'18. Red Hook, NY, USA: Curran Associates Inc., 2018, p. 2615–2625.
- [43] B. Zhang, J. Yuan, B. Li, T. Chen, J. Fan, and B. Shi, "Learning cross-image object semantic relation in transformer for few-shot fine-grained image classification," in *Proceedings of the 30th ACM International Conference on Multimedia*, 2022, pp. 2135–2144.
- [44] Z. Zha, H. Tang, Y. Sun, and J. Tang, "Boosting few-shot fine-grained recognition with background suppression and foreground alignment," *IEEE Transactions on Circuits and Systems for Video Technology*, 2023.
- [45] C. Song, Y. Huang, W. Ouyang, and L. Wang, "Mask-guided contrastive attention model for person re-identification," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 1179–1188.
- [46] C. Fu, J. Xu, F. Lin, F. Guo, T. Liu, and Z. Zhang, "Object saliency-aware dual regularized correlation filter for real-time aerial tracking," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 58, no. 12, pp. 8940–8951, 2020.
- [47] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015. [Online]. Available: <https://arxiv.org/abs/1503.02531>
- [48] S. Ahn, S. X. Hu, A. Damianou, N. D. Lawrence, and Z. Dai, "Variational information distillation for knowledge transfer," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 9163–9171.
- [49] D. Chen, J.-P. Mei, Y. Zhang, C. Wang, Z. Wang, Y. Feng, and C. Chen, "Cross-layer distillation with semantic calibration," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 8, 2021, pp. 7028–7036.
- [50] B. Zhao, Q. Cui, R. Song, Y. Qiu, and J. Liang, "Decoupled knowledge distillation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 11953–11962.
- [51] D. Chen, J.-P. Mei, H. Zhang, C. Wang, Y. Feng, and C. Chen, "Knowledge distillation with the reused teacher classifier," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 11933–11942.
- [52] S. Stanton, P. Izmailov, P. Kirichenko, A. A. Alemi, and A. G. Wilson, "Does knowledge distillation really work?" *Advances in Neural Information Processing Systems*, vol. 34, pp. 6906–6919, 2021.
- [53] Y. Zhang, T. Xiang, T. M. Hospedales, and H. Lu, "Deep mutual learning," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4320–4328.
- [54] K. Zhou, Y. Yang, A. Cavallaro, and T. Xiang, "Omni-scale feature learning for person re-identification," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 3702–3712.
- [55] H. Wei, L. Feng, X. Chen, and B. An, "Combating noisy labels by agreement: A joint training method with co-regularization," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 13726–13735.
- [56] Y. Pang, X. Zhao, L. Zhang, and H. Lu, "Multi-scale interactive network for salient object detection," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 9413–9422.
- [57] J. Ma, H. Xie, G. Han, S.-F. Chang, A. Galstyan, and W. Abd-Almageed, "Partner-assisted learning for few-shot image classification," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 10573–10582.
- [58] Z. Zhou, X. Qiu, J. Xie, J. Wu, and C. Zhang, "Binocular mutual learning for improving few-shot classification," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 8402–8411.
- [59] H.-J. Ye, L. Ming, D.-C. Zhan, and W.-L. Chao, "Few-shot learning with a strong teacher," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2022.
- [60] X. Qin, Z. Zhang, C. Huang, C. Gao, M. Dehghan, and M. Jagersand, "BASNet: Boundary-aware salient object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [61] X. Qin, Z. Zhang, C. Huang, M. Dehghan, O. R. Zaiane, and M. Jagersand, "U2-Net: Going deeper with nested u-structure for salient object detection," *Pattern recognition*, vol. 106, p. 107404, 2020.
- [62] L. Wang, H. Lu, Y. Wang, M. Feng, D. Wang, B. Yin, and X. Ruan, "Learning to detect salient objects with image-level supervision," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 136–145.
- [63] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie, "The caltech-ucsd birds-200-2011 dataset," California Institute of Technology, Tech. Rep. CNS-TR-2011-001, 2011.
- [64] A. Khosla, N. Jayadevaprakash, B. Yao, and F.-F. Li, "Novel dataset for fine-grained image categorization: Stanford dogs," in *Proc. CVPR workshop on fine-grained visual categorization (FGVC)*, vol. 2, no. 1. Citeseer, 2011.
- [65] J. Krause, M. Stark, J. Deng, and L. Fei-Fei, "3D object representations for fine-grained categorization," in *ICCV Workshops*, 2013, pp. 554–561.
- [66] W.-Y. Chen, Y.-C. Liu, Z. Kira, Y.-C. F. Wang, and J.-B. Huang, "A closer look at few-shot classification," in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=HkxLXnAcFQ>
- [67] Y. Wang, W.-L. Chao, K. Q. Weinberger, and L. van der Maaten, "SimpleShot: Revisiting nearest-neighbor classification for few-shot learning," *arXiv preprint arXiv:1911.04623*, 2019. [Online]. Available: <https://arxiv.org/abs/1911.04623>
- [68] Y. Chen, Z. Liu, H. Xu, T. Darrell, and X. Wang, "Meta-baseline: Exploring simple meta-learning for few-shot learning," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9062–9071.
- [69] H.-J. Ye, H. Hu, D.-C. Zhan, and F. Sha, "Few-shot learning via embedding adaptation with set-to-set functions," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 8808–8817.
- [70] R. Hou, H. Chang, B. Ma, S. Shan, and X. Chen, "Cross attention network for few-shot classification," in *NeurIPS*, 2019, pp. 4005–4016.
- [71] H. Tang, C. Yuan, Z. Li, and J. Tang, "Learning attention-guided pyramidal features for few-shot fine-grained recognition," *Pattern Recognit.*, vol. 130, p. 108792, 2022.
- [72] 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 Transactions on Multimedia*, vol. 23, pp. 1666–1680, 2020.
- [73] Y. Zhu, C. Liu, and S. Jiang, "Multi-attention meta learning for few-shot fine-grained image recognition," in *IJCAI*, 2020, pp. 1090–1096.
- [74] H. Zhang and P. Koniusz, "Power normalizing second-order similarity network for few-shot learning," in *2019 IEEE winter conference on applications of computer vision (WACV)*. IEEE, 2019, pp. 1185–1193.
- [75] H. Zhang, H. Li, and P. Koniusz, "Multi-level second-order few-shot learning," *IEEE Transactions on Multimedia*, 2022.
- [76] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 618–626.
- [77] L. Tang, B. Li, Y. Zhong, S. Ding, and M. Song, "Disentangled high quality salient object detection," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 3580–3590.
- [78] J. Zhang, Q. Liang, Q. Guo, J. Yang, Q. Zhang, and Y. Shi, "R2Net: Residual refinement network for salient object detection," *Image and Vision Computing*, vol. 120, p. 104423, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S026288562200052X>
- [79] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne." *Journal of machine learning research*, vol. 9, no. 11, 2008.
- [80] Z. Chen, L. Xie, J. Niu, X. Liu, L. Wei, and Q. Tian, "Visformer: The vision-friendly transformer," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 589–598.
- [81] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollar, and R. Girshick, "Segment anything," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023, pp. 4015–4026.
- [82] X. Zou, J. Yang, H. Zhang, F. Li, L. Li, J. Wang, L. Wang, J. Gao, and Y. J. Lee, "Segment everything everywhere all at once," in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [Online]. Available: <https://openreview.net/forum?id=UHBrWeFWIL>
- [83] W. Ji, J. Li, Q. Bi, W. Li, and L. Cheng, "Segment anything is not always perfect: An investigation of sam on different real-world applications," *arXiv preprint arXiv:2304.05750*, 2023. [Online]. Available: <https://arxiv.org/abs/2304.05750>
- [84] X.-S. Wei, H.-Y. Xu, F. Zhang, Y. Peng, and W. Zhou, "An embarrassingly simple approach to semi-supervised few-shot learning," in *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35. Curran Associates, Inc., 2022, pp. 14489–14500. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2022/file/5d3b57e06e3fc45f077eb5c9f28156d4-Paper-Conference.pdf

- [85] X.-S. Wei, H.-Y. Xu, Z. Yang, C.-L. Duan, and Y. Peng, "Negatives make a positive: An embarrassingly simple approach to semi-supervised few-shot learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–14, 2023.



Haiqi Liu received the B.S. degree in computer science and technology from South China University of Technology, Guangzhou, China, in 2020. He is currently pursuing the Ph.D. degree in computer science and technology with the School of Computer Science and Engineering.

His research interests mainly include Few shot Learning, Image Recognition and Affective Computing.



C. L. Philip Chen (S'88–M'88–SM'94–F'07) received the M.S. degree from the University of Michigan at Ann Arbor, Ann Arbor, MI, USA, in 1985 and the Ph.D. degree from the Purdue University in 1988, all in electrical and computer science.

He is the Chair Professor and Dean of the College of Computer Science and Engineering, South China University of Technology. He is the former Dean of the Faculty of Science and Technology. He is a Fellow of IEEE, AAAS, IAPR, CAA, and HKIE; a member of Academia Europaea (AE) and European

Academy of Sciences and Arts (EASA). He received IEEE Norbert Wiener Award in 2018 for his contribution in systems and cybernetics, and machine learnings. He is also a highly cited researcher by Clarivate Analytics in 2018, 2019, 2020, 2021 and 2022.

He was the Editor-in-Chief of the IEEE Transactions on Cybernetics (2020–2021) after he completed his term as the Editor-in-Chief of the IEEE Transactions on Systems, Man, and Cybernetics: Systems (2014–2019), followed by serving as the IEEE Systems, Man, and Cybernetics Society President from 2012 to 2013. Currently, he serves as a deputy director of CAAI Transactions on AI, an Associate Editor of the IEEE Transactions on AI, IEEE Trans on SMC: Systems, and IEEE Transactions on Fuzzy Systems, an Associate Editor of China Sciences: Information Sciences. He received Macau FDCT Natural Science Award three times and a First-rank Guangdong Province Scientific and Technology Advancement Award in 2019. His current research interests include cybernetics, computational intelligence, and systems.



Xinrong Gong received the B.S. degree in the School of Electronic and Optical Engineering from the Nanjing University of Science and Technology, Nanjing, China, in 2018. He is pursuing a Ph.D. at the School of Computer Science and Engineering at the South China University of Technology, in Guangzhou, China.

His research interests include Machine Learning, Broad Learning Systems, and Multimodal Emotion Recognition.



Tong Zhang (M'16–SM'24) received the B.S. degree in software engineering from Sun Yat-sen University, at Guangzhou, China, in 2009, and the M.S. degree in applied mathematics from University of Macau, at Macau, China, in 2011, and the Ph.D. degree in software engineering from the University of Macau, at Macau, China in 2016. Dr. Zhang currently is a professor and Associate Dean of the School of Computer Science and Engineering, South China University of Technology, China.

His research interests include affective computing, evolutionary computation, neural network, and other machine learning techniques and their applications. Prof. Zhang is the Associate Editor of the IEEE Transactions on Affective Computing, IEEE Transactions on Computational Social Systems, and Journal of Intelligent Manufacturing. He has been working in publication matters for many IEEE conferences.