







Eliminating Feature Ambiguity for Few-Shot Segmentation

Qianxiong Xu¹, Guosheng Lin^{1*}, Chen Change Loy¹, Cheng Long^{1*},
Ziyue Li², and Rui Zhao³

¹ S-Lab, Nanyang Technological University

² University of Cologne

³ SenseTime Research

{qianxiong.xu, gslin, ccloy, c.long}@ntu.edu.sg,
zlibn@wiso.uni-koeln.de, zhaorui@sensetime.com

Abstract. Recent advancements in few-shot segmentation (FSS) have exploited pixel-by-pixel matching between query and support features, typically based on cross attention, which selectively activate query foreground (FG) features that correspond to the same-class support FG features. However, due to the large receptive fields in deep layers of the backbone, the extracted query and support FG features are inevitably mingled with background (BG) features, impeding the FG-FG matching in cross attention. Hence, the query FG features are fused with less support FG features, *i.e.*, the support information is not well utilized. This paper presents a novel plug-in termed ambiguity elimination network (AENet), which can be plugged into any existing cross attention-based FSS methods. The main idea is to mine discriminative query FG regions to rectify the ambiguous FG features, increasing the proportion of FG information, so as to suppress the negative impacts of the doped BG features. In this way, the FG-FG matching is naturally enhanced. We plug AENet into three baselines CyCTR, SCCAN and HDMNet for evaluation, and their scores are improved by large margins, *e.g.*, the 1-shot performance of SCCAN can be improved by 3.0%+ on both PASCAL-5ⁱ and COCO-20ⁱ. The code is available at <https://github.com/Sam1224/AENet>.

Keywords: Discriminative prior mask · Discriminative query regions · Feature refinement

1 Introduction

Semantic segmentation is a fundamental task within computer vision, entailing the dense assignment of each pixel in an image to an appropriate class label [5, 6, 17]. This process is crucial for understanding the detailed composition of visual scenes. The advent and subsequent evolution of deep learning approaches have led to significant advancements in semantic segmentation [2, 23, 29, 51, 52]. However, most approaches relied heavily on precise pixel-level annotation of

* Co-corresponding authors

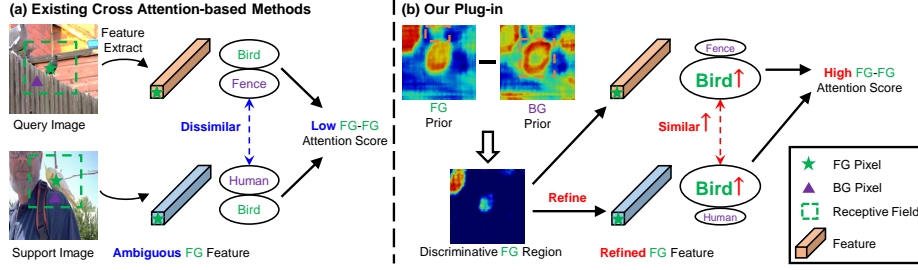


Fig. 1: Illustrations of (a) existing methods, and (b) the rationale of our plug-in. In (a), due to the large receptive fields during feature extraction, foreground (FG) pixels’ features are inevitably fused with background (BG) features (*e.g.*, the bird pixels also contain dissimilar BGs: fence and human), which hinders FG-FG matching in cross attention. In (b), we propose a plug-in to mine discriminative query FG regions, which exclude those regions that are similar to both support FG and BG features, for refining the ambiguous query and support FG features. As a result, the FG parts in the mingled FG features are increased, thus the FG-FG matching is naturally enhanced.

training datasets. Moreover, these models typically exhibit limitations in extending their recognition capabilities to novel, previously unseen classes.

To address these challenges, few-shot segmentation (FSS) has been introduced to segment query images containing arbitrary classes, with the help of a few support (image, mask) pairs sharing the same class. During training, FSS models would learn the class-agnostic pattern on some base classes, enabling it to identify query features resembling support FG features and classify them as query FG. Then, such pattern is directly applied to segment novel classes.

The effectiveness of FSS heavily relies on the skillful utilization of support samples, based on which the existing methods can be categorized into prototype-based [1, 4, 14, 18, 35, 40, 49] and cross attention-based methods [28, 37, 42, 45, 48, 50]. Prototype-based methods typically entail the extraction of support prototypes from support FG features, which are subsequently leveraged to segment the query image through feature comparison [40] or concatenation [35]. Nevertheless, the compression of features into prototypes would lead to potential information loss, as well as the disruption of the spatial structure of objects [44]. To address these issues, recent advancements in FSS have embraced cross attention [36] to fuse query features with the uncompressed support FG features.

Despite their success, existing cross attention-based methods overlook the *ineffective FG-FG matching* issue raised by *feature ambiguity*. In particular, it is a common practice [31, 40, 49] to forward query and support images to a pretrained backbone (*e.g.*, ResNet50 [8]) to extract their features. However, as illustrated in Fig. 1(a), deep convolution layers essentially have large receptive fields, so the extracted FG/BG pixels’ features are inevitably fused with other BG/FG features, especially for those pixels locating at the boundary area between FG and BG objects. Some evidences are provided in Fig. 1(b) in the form of prior masks [35]. Specifically, cosine similarity is measured between each query feature and sup-

port FG features to obtain the FG prior, with each value showing the probability of being a FG pixel. BG prior is obtained in the same way. It can be observed: (1) There are many wrongly activated query BG regions (in orange rectangle) in FG prior, because they have aggregated nearby FG features, thereby showing high similarity to the support FG; (2) In BG prior, support BG features can also match with query FG features, because the BG pixels on the border have also been integrated with support FG information. Take Fig. 1(a) as an example, where a query FG pixel contains FG (bird) and BG (fence) features, and support FG features include FG (bird) and BG (human). As query and support FG features are doped with different-class BG features, the similarity-based cross attention scores would be relatively smaller, which hinders query FG features from fusing more support FG features, thereby leading to ineffective FSS.

In this paper, we aim to design a plug-in to address the aforementioned issues. The key idea is to suppress those ambiguous query regions that are similar to both support FG and BG features (doped with some FG information). In this way, the remaining FG regions refer to the most discriminative query FG regions, receiving the least side-effects from the mingled BG features. Based on this idea, we propose a plug-in named **ambiguity elimination network (AENet)**, which includes: (1) **Prior generator (PG)**: We incorporate the idea into the learning-agnostic prior mask, and generate a pair of prior masks to figure out the approximate scope of query FG, while highlighting the most discriminative query FG regions, facilitating fast convergence [35]. The visualizations (in Fig. 1(b) and Fig. 5(b)) can directly show the effectiveness of the idea; (2) **Ambiguity eliminator (AE)**: The idea is further employed to rectify the query and support FG features, so as to naturally improve the FG-FG matching in cross attention, *i.e.*, to better utilize the support information. As shown in Fig. 1(b), the features of the discriminative query FG regions are fused with the query and support features for refinement. For a query FG pixel, its features would consequently contain more FG information, so the side-effects of the mingled BG features can be suppressed. In this paper, we validate the effectiveness of AENet on three cross attention-based FSS baselines, CyCTR [50], SCCAN [45] and HDMNet [28].

To our knowledge, we are the first to identify the negative impacts of feature ambiguity on FSS. We propose a simple yet effective idea to obtain the discriminative query FG regions. Then, we propose the plug-in network AENet to enhance the FG-FG matching for existing cross attention-based FSS methods. Extensive experiments are conducted on two public benchmarks PASCAL-5ⁱ and COCO-20ⁱ to validate the effectiveness of AENet. Notably, AENet is plugged into CyCTR, SCCAN and HDMNet, and can improve their performance by large margins, *e.g.*, the 1-shot performance of SCCAN can be boosted from 66.8% to 69.8%, and 46.3% to 49.4% on PASCAL-5ⁱ and COCO-20ⁱ, respectively.

2 Related Work

Few-shot segmentation. Different from semantic segmentation models that are trained and tested on the same set of classes, FSS is designed to segment

arbitrary classes with the help of a few labelled samples. FSS methods in recent literature can be broadly categorized into several groups.

Prototype-based methods [4, 16, 18, 22, 26, 31, 35, 38–40, 47] represent a prominent category within the realm of FSS, wherein support FG information is compressed into single or multiple prototypes, facilitating the segmentation of query images. These methodologies typically leverage techniques such as cosine similarity or feature concatenation to perform segmentation. In particular, OSLSM [31] makes pioneering contributions to the field by introducing the concept of FSS, laying the groundwork for subsequent research endeavors. PFENet [35] firstly derives a learning-agnostic query prior mask from the high-level query and support features to help coarsely locate the query FG objects and improve the convergence speed. Each value in the prior mask represents the feature similarity between the current query pixel to the support FG pixels. However, the existing method is heavily affected by noises [45], and there are many wrong responses on query BG areas. In this paper, we propose a prior generator (PG) to highlight the most discriminative query regions that can accurately locate query FG.

To prevent from the information loss and the structure disruption raised by prototypes, attention-based methods [10, 12, 37, 43, 45, 48, 50] build pixel-pixel matching between query and support features, and expect to learn similarity-based cross attentions to activate query FG features with the same-class support FG features. PGNet [48] builds up two graphs on the query and support features, then incorporate the attention mechanism with multi-scale graph reasoning to perform segmentation. Besides, CyCTR [50] devises a cycle-consistent attention to suppress the side-effects of harmful support features. Recently, SC-CAN [45] indicates the mismatch problem of query BG features to the support FG features when conducting cross attention, and propose a self-calibrated cross attention block to align query BG features with appropriate BG features for effective segmentation. Another lane of methods [9, 20, 27, 32, 44] regard the problem as semantic correspondence and use memory-expensive 4D attentions for the matching. Unfortunately, it would be ineffective for these methods to perform query-support FG-FG matching (via cross attention), because the extracted query and support FG features are likely to be mingled with dissimilar BG features, impeding the effective utilization of support information. To address it, we propose an ambiguity eliminator (AE) to purify the FG features and make the FG-FG matching effective.

Moreover, many methods extend the standard setting by using base class predictions [11, 14, 28, 34, 54], mining knowledge from unlabelled data [1] or more advanced pretrained backbone [13], or introducing extra textual information [46, 53]. To have better comparisons with recent advances, we have conducted experiments under both standard and BAM’s [14] settings in Sec. 5.3.

3 Problem Definition

Let $\mathcal{D}_{\text{train}}$ and $\mathcal{D}_{\text{test}}$ represent the training and testing sets, respectively, in the context of FSS. $\mathcal{D}_{\text{train}}$ encompasses a collection of base classes $\mathcal{C}_{\text{base}}$, while $\mathcal{D}_{\text{test}}$

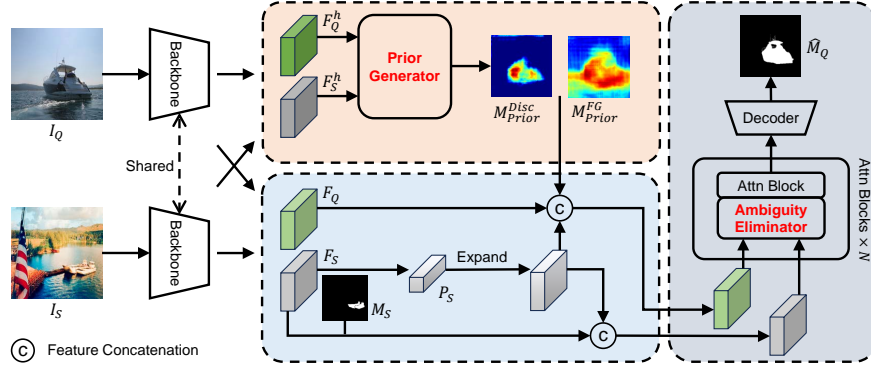


Fig. 2: Overview of ambiguity elimination network (AENet), which is designed for existing cross attention-based methods such as SCCAN [45]. Specifically, AENet includes prior generator (PG) and ambiguity eliminator (AE). (1) Learning-agnostic PG helps to accurately locate the query FG regions; (2) AE mines the most discriminative query FG regions, and then rectifies the query and support FG features to improve the FG-FG matching in cross attention.

encompasses another distinct set of novel classes $\mathcal{C}_{\text{novel}}$. FSS addresses a scenario where the sets of classes $\mathcal{C}_{\text{base}}$ and $\mathcal{C}_{\text{novel}}$ are disjoint, formally denoted as $\mathcal{C}_{\text{base}} \cap \mathcal{C}_{\text{novel}} = \emptyset$. Both $\mathcal{D}_{\text{train}}$ and $\mathcal{D}_{\text{test}}$ are composed of numerous episodes, which serve as the fundamental units of episodic training. In the context of a k -shot setting, each episode comprises a support set $\mathcal{S} = \{I_S^n, M_S^n\}_{n=1}^k$ and a query set $\mathcal{Q} = \{I_Q, M_Q\}$ specific to a particular class c . Here, I_S^n and M_S^n denote the n -th support image and its corresponding annotated binary mask, while I_Q and M_Q represent the query image and its associated mask. During the training phase, the model learns to segment the query image I_Q with guidance from the support set \mathcal{S} , focusing on classes from $\mathcal{C}_{\text{base}}$. Subsequently, this learned segmentation pattern is applied to $\mathcal{C}_{\text{novel}}$ during the testing phase.

4 Methodology

The standard FSS pipeline is described as follows: The query image I_Q and the support image I_S are forwarded to a frozen ImageNet-pretrained [30] backbone (*e.g.*, VGG16 [33] or ResNet50 [8]) for extracting their mid-level features $\{F_Q, F_S\}$ (from blocks 2 and 3) and high-level features $\{F_Q^h, F_S^h\}$ (from block 4). As explained in existing methods [35, 49], FSS models would witness severe overfitting phenomena if segmentation is carried out with the high-level features. In spite of this, PFENet [35] demonstrates that learning-agnostic prior masks can be derived from high-level features to coarsely locate the query FG objects and improve the convergence speed. Therefore, almost all recent advances [14, 28, 45] follow the same guideline to take the mid-level features for segmentation, and use the high-level features for generating prior masks. Then, support FG features are obtained with F_S and support mask M_S , which are fused with query

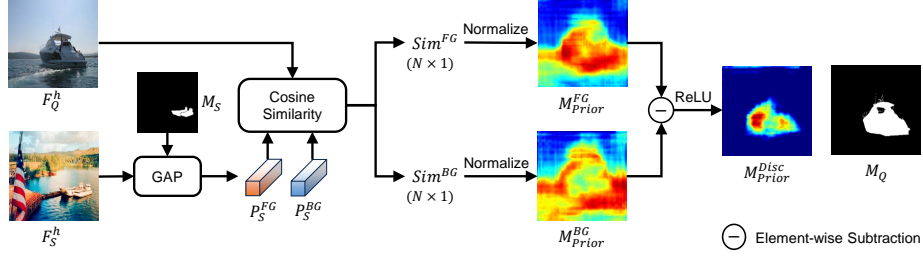


Fig. 3: Details of the learning-agnostic prior generator (PG) module. The discriminative prior mask can suppress those ambiguous query regions that are similar to both support FG and BG features.

features, either in a prototypical way or a cross attention manner, to activate those query FG features that share the same class. Finally, the enhanced query features are processed by a decoder to obtain the predicted query mask \hat{M}_Q .

To mitigate the *ineffective FG-FG matching* issue and *feature ambiguity* issue (as described in Sec. 1), we present the overview of our plug-in ambiguity elimination network (AENet) in Fig. 2, whose main components consist of the prior generator (PG) module (Sec. 4.1) and the ambiguity eliminator (AE) module (Sec. 4.2). AENet can be plugged into any existing cross attention-based FSS baselines such as CyCTR [50], SCCAN [45] and HDMNet [28]. Take SCCAN as an example, its pseudo mask aggregation (PMA) module is replaced with our proposed PG, and we insert one AE before each of its self-calibrated cross attention (SCCA) block, with other parts remaining untouched. Particularly, the learning-agnostic PG can not only approximately locate the query FG, but also extract the discriminative query FG regions. In this way, the FSS model can obtain some hints with no additional learnable parameters, and its converge speed is well improved. Besides, the AE module aims at using the rectified discriminative query FG regions to refine the query FG and support FG features, so as to suppress the negative impacts of their doped BG features. As a result, the FG-FG matching in cross attention is enhanced, leading to more effective FSS.

4.1 Prior Generator (PG)

Since PFENet [35], learning-agnostic prior masks are widely adopted in FSS models [28, 45, 50], to coarsely locate the query FG objects and improve the convergence speed. For each query pixel, most of the existing methods [21, 24, 35, 50] take its high-level features, calculate its cosine similarity scores with each support FG feature, and take the normalized maximum score to show how likely this query pixel belongs to FG class. Later, SCCAN [45] indicates the operation of taking a single maximum score is heavily affected by noises, and they propose to utilize all pairwise similarities for generating more robust prior masks.

Nevertheless, they overlook the fact that the high-level FG and BG features are inevitably fused with BG and FG features (*i.e.*, feature ambiguity), which

means query-support FG-FG may not consistently have large similarity scores, while query-support BG-FG can also mistakenly have high similarity. Hence, the generated prior masks would be less effective. We support our claim by providing some examples in Fig. 3, where M_{Prior}^{FG} and M_{Prior}^{BG} are the query prior masks calculated with support FG and BG features, respectively. We could observe that (1) there are many BG regions wrongly activated in M_{Prior}^{FG} , and (2) support BG can also have high responses on query FG in M_{Prior}^{BG} .

To tackle the aforementioned issue, we propose to rectify and extract the most discriminative query FG regions via a simple subtraction operation between M_{Prior}^{FG} and M_{Prior}^{BG} , whose essence is removing those ambiguous query regions that show high similarity to both support FG and BG features. As shown in Fig. 3, the discriminative prior mask M_{Prior}^{Disc} can remove the unexpected areas, suppressing the side-effects of the doped FG information in support BG features.

The details of PG are presented in Fig. 3, and we formally describe it as follows. Following existing methods, we take the high-level query features F_Q^h and support features F_S^h , as well as the downsampled support mask M_S as inputs. First of all, the support FG and BG prototypes P_S^{FG} and P_S^{BG} are obtained via:

$$P_S^{FG} = GAP(F_S^h, M_S), P_S^{BG} = GAP(F_S^h, 1 - M_S) \quad (1)$$

where $GAP(\cdot)$ denotes the global average pooling (GAP) operation. Then, two cosine similarity scores $Sim^{FG} \in [-1, 1]$ and $Sim^{BG} \in [-1, 1]$ are calculated between the flattened query features and the support prototypes. They are normalized and reshaped to obtain the prior masks $M_{Prior}^{FG} \in [0, 1]$ and $M_{Prior}^{BG} \in [0, 1]$, showing the query regions that are similar to support FG and BG features, respectively. Notably, our memory cost is $HW \times 1$ for each calculation, while that of existing methods [35, 45] is $HW \times HW$.

$$Sim^* = Cosine(F_Q^h, P_S^*) \quad (2)$$

$$M_{Prior}^* = Norm(Sim^*) \quad (3)$$

where superscript $*$ $\in \{FG, BG\}$, $Cosine(\cdot)$ means cosine similarity, $Norm(\cdot)$ is the min-max scaler to normalize the values into $[0, 1]$. Next, we perform a clipped subtraction to rectify and obtain the discriminative prior mask M_{Prior}^{Disc} .

$$M_{Prior}^{Disc} = ReLU(M_{Prior}^{FG} - M_{Prior}^{BG}) \quad (4)$$

where $ReLU(\cdot)$ is an operator to set the negative values as zeros. The negative values correspond to those query regions that are similar to support BG features, which are not helpful in FSS. Finally, we concatenate M_{Prior}^{FG} and M_{Prior}^{Disc} as the final prior masks, providing both the coarse location of query FG objects and the most discriminative query FG regions.

4.2 Ambiguity Eliminator (AE)

To tackle the *ineffective FG-FG matching* raised by *feature ambiguity*, we further design the AE module for existing cross attention-based FSS methods [28, 45,

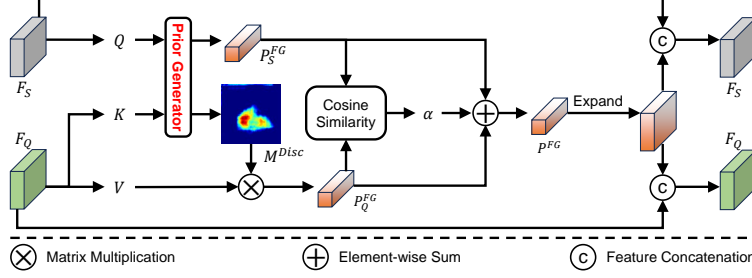


Fig. 4: Details of ambiguity eliminator (AE) module. AE mines discriminative query FG features to rectify the query and support features.

50]. As verified in Fig. 3, discriminative query FG regions can be easily and consistently extracted, whose features are the most discriminative ones that are less affected by *feature ambiguity*. Therefore, we use them to rectify the query and support FG features, so the proportion of pure FG information in a FG pixel’s mingled features can be naturally increased. Therefore, the FG-FG matching between query and support FG features is enhanced, and the query FG pixels can aggregate more support FG information.

As illustrated in Fig. 4, the mid-level query features F_Q and support features F_S are taken as the inputs. F_Q is processed by two linear layers to obtain the projections K and V , while F_S is projected to Q . Then, the Q and K are forwarded to the PG module to obtain the support FG prototype P_S^{FG} and the discriminative prior mask M^{Disc} , which is supervised by an auxiliary binary cross entropy (BCE) loss \mathcal{L}_{aux} . The procedure can be written as:

$$Q = Linear(F_S), K = Linear(F_Q), V = Linear(F_Q) \quad (5)$$

$$P_S^{FG}, M^{Disc} = PG(Q, K) \quad (6)$$

$$\mathcal{L}_{aux} = BCE(M^{Disc}, M_Q) \quad (7)$$

where $PG(\cdot)$ includes Eq. (1) to Eq. (4), \mathcal{L}_{aux} is utilized during training, M_Q is the labelled query mask. Then, a matrix multiplication is performed to extract and aggregate the discriminative query features into a prototype P_Q^{FG} . Next, the cosine similarity α between P_S^{FG} and P_Q^{FG} is measured, and re-scaled to $[0, 1]$. α is utilized to weighted fuse the prototypes to obtain P^{FG} , containing the support FG features and the most discriminative query FG features. After that, P^{FG} is expanded and concatenated with the input query and support features F_Q and F_S , and processed by a linear layer for feature refinement.

$$P_Q^{FG} = Softmax(M^{Disc}) \otimes V \quad (8)$$

$$\alpha = (Cosine(P_S^{FG}, P_Q^{FG}) + 1)/2 \quad (9)$$

$$P^{FG} = \alpha \cdot P_S^{FG} + (1 - \alpha) \cdot P_Q^{FG} \quad (10)$$

$$F_* = Linear(F_* || P^{FG}) \quad (11)$$

where \otimes is the matrix multiplication operation, subscript $*$ $\in \{Q, S\}$, \parallel denotes feature concatenation. Finally, we wrap AE with a Transformer block [36], and the refined query and support features are forwarded to existing cross attention blocks for feature matching and enhancement.

4.3 Overall Loss

The loss function consists of a main loss and an auxiliary loss. The former refers to the loss functions employed in the original baseline, while the auxiliary loss corresponds to Eq. (7). Take SCCAN as an example, its loss function is:

$$\mathcal{L} = \mathcal{L}_{main} + \lambda \cdot \mathcal{L}_{aux} \quad (12)$$

$$= Dice(\hat{M}_Q, M_Q) + \lambda \cdot \left(\frac{1}{N} \sum_{i=1}^N BCE(M_i^{Disc}, M_Q) \right) \quad (13)$$

where $Dice(\cdot)$ represents dice loss, \hat{M}_Q is the final prediction, λ is a hyperparameter, N is the number of attention blocks (*e.g.*, SCCAN has 8 SCCA blocks), M_i^{Disc} is the i -th discriminative mask obtained from the i -th AE module.

5 Experiments

5.1 Experiment Setup

Datasets. We assess the performance of our methodology on two widely used benchmark datasets, including PASCAL-5ⁱ [31] and COCO-20ⁱ [25]. PASCAL-5ⁱ comprises 20 distinct classes and is derived from PASCAL VOC 2012 [3], augmented with additional annotations from [7]. Conversely, COCO-20ⁱ is constructed from MSCOCO [19], presenting a more rigorous challenge with its 80 classes. Both PASCAL-5ⁱ and COCO-20ⁱ are partitioned into four folds for cross-validation purposes, with each fold containing either 5 (for PASCAL-5ⁱ) or 20 (for COCO-20ⁱ) classes. Within each fold, the training set encompasses the union of the other three folds, while the fold itself is reserved for testing. Furthermore, during testing, we randomly sample 1,000 episodes from PASCAL-5ⁱ and 4,000 episodes from COCO-20ⁱ, ensuring a comprehensive evaluation of our method’s efficacy across diverse scenarios.

Evaluation metrics. It is a common practice to use mean Intersection over Union (mIoU) and foreground-background IoU (FB-IoU) as the evaluation metrics [28, 35, 45]. Specifically, mIoU computes the average IoU scores across all FG classes within a fold, providing a consolidated measure of segmentation accuracy. Instead, FB-IoU treats all FG classes as a single FG class for measurement.

Implementation details. To validate the effectiveness of AENet, we plug it into three cross attention-based baselines: CyCTR [50], SCCAN [45] and HDM-Net [28]. In detail, their prior mask generation modules are replaced by PG (Sec. 4.1), and we insert an AE module (Sec. 4.2) before each of their cross attention blocks. All experiments are conducted with 4 NVIDIA V100 GPUs with

Table 1: Performance comparisons on PASCAL-5ⁱ in terms of mIoU and FB-IoU. “5ⁱ” shows the mIoU scores of 5 novel classes in fold i , “Mean” is the averaged mIoU score from all folds. The best results are highlighted in **bold**.

Backbone Method		1-shot						5-shot					
		5 ⁰	5 ¹	5 ²	5 ³	Mean	FB-IoU	5 ⁰	5 ¹	5 ²	5 ³	Mean	FB-IoU
VGG16	PFENet (TPAMI'20) [35]	56.9	68.2	54.4	52.5	58.0	72.0	59.0	69.1	54.8	52.9	59.0	72.3
	DACM (ECCV'22) [44]	61.8	67.8	61.4	56.3	61.8	75.5	66.1	70.6	65.8	60.2	65.7	77.8
	FECANet (TMM'23) [20]	66.5	68.9	63.6	58.3	64.3	76.2	68.6	70.8	66.7	60.7	66.7	77.6
	BAM (CVPR'22) [14]	63.2	70.8	66.1	57.5	64.4	77.3	67.4	73.1	70.6	64.0	68.8	81.1
	HDMNet (CVPR'23) [28]	64.8	71.4	67.7	56.4	65.1	-	68.1	73.1	71.8	64.0	69.3	-
	HDMNet + AENet	64.7	73.2	66.4	60.1	66.1 _{1.0} †	78.1	69.1	74.8	69.4	65.1	69.6 _{0.3} †	82.0
	SCCAN (ICCV'23) [45]	63.3	70.8	66.6	58.2	64.7	77.2	67.2	72.3	70.5	63.8	68.4	79.1
	SCCAN + AENet	66.3	73.3	68.5	58.4	66.6 _{1.9} †	79.0	70.8	75.1	72.2	64.2	70.6 _{2.2} †	81.8
ResNet50	PFENet (TPAMI'20)	61.7	69.5	55.4	56.3	60.8	73.3	63.1	70.7	55.8	57.9	61.9	73.9
	DACM (ECCV'22) [44]	66.5	72.6	62.2	61.3	65.7	77.8	72.4	73.7	69.1	68.4	70.9	81.3
	ABCNet (CVPR'23) [41]	68.8	73.4	62.3	59.5	66.0	76.0	71.7	74.2	65.4	67.0	69.6	80.0
	RiFeNet (AAAI'24) [1]	68.4	73.5	67.1	59.4	67.1	-	70.0	74.7	69.4	64.2	69.6	-
	FECANet (TMM'23) [20]	69.2	72.3	62.4	65.7	67.4	78.7	72.9	74.0	65.2	67.8	70.0	80.7
	BAM (CVPR'22) [14]	69.0	73.6	67.6	61.1	67.8	79.7	70.6	75.1	70.8	67.2	70.9	82.2
	MIANet (CVPR'23) [46]	68.5	75.8	67.5	63.2	68.7	79.5	70.2	77.4	70.0	68.8	71.6	82.2
	CyCTR (NIPS'21) [50]	67.8	72.8	58.0	58.0	64.2	-	71.1	73.2	60.5	57.5	65.6	-
	CyCTR + AENet	70.4	74.4	69.4	61.6	69.0 _{4.8} †	80.2	73.0	76.2	74.1	67.0	72.6 _{7.0} †	83.0
	SCCAN (ICCV'23) [45]	68.3	72.5	66.8	59.8	66.8	77.7	72.3	74.1	69.1	65.6	70.3	81.8
	SCCAN + AENet	72.2	75.5	68.5	63.1	69.8 _{3.0} †	80.8	74.2	76.5	74.8	70.6	74.1 _{3.8} †	84.5
	HDMNet (CVPR'23) [28]	71.0	75.4	68.9	62.1	69.4	-	71.3	76.2	71.3	68.5	71.8	-
	HDMNet + AENet	71.3	75.9	68.6	65.4	70.3 _{0.9} †	81.2	73.9	77.8	73.3	72.0	74.2 _{2.4} †	84.5

16GB onboard memory. For both datasets, we adopt the same augmentation strategy as [14, 35], and we follow the selected baselines to set their hyperparameters. In this paper, we perform evaluation under both standard [35] and BAM’s [14] settings. For the former, we adopt VGG16 [33] and ResNet50 [8] pretrained on ImageNet [30] as backbones, while for the latter, they are further fine-tuned for segmenting \mathcal{C}_{base} , which is more appropriate for FSS. Besides, we follow the best hyperparameter settings as the baselines (*e.g.*, number of attention blocks), while for AENet-related parameters, the weight λ (Eq. (12)) is set to 1, and is further studied in Sec. 5.3.

5.2 Quantitative Comparisons with State of the Arts

To well validate the effectiveness of AENet, we plug it into CyCTR [50], SCCAN [45] and HDMNet [28], and conduct experiments under both 1-shot and 5-shot settings, with VGG16 [33] and ResNet50 [8] as pretrained backbones.

PASCAL-5ⁱ. The quantitative results on PASCAL-5ⁱ are shown in Tab. 1, we could observe that our AENet plug-in consistently helps to improve the selected baselines by large margins. For example, by taking ResNet50 as the pretrained backbone, the mean mIoU score of SCCAN can be boosted from 66.8% to 69.8% on PASCAL-5ⁱ under 1-shot setting, and the improvement can be increased to 3.8% when 5 support pairs are provided. Similarly, FB-IoU can be improved by 2.3% and 2.7%, respectively. Notably, the 1-shot performance gain on CyCTR can be as high as 4.8%, showing the superiority of AENet. Moreover, SCCAN

Table 2: Performance comparisons on COCO-20ⁱ in terms of mIoU and FB-IoU. “20ⁱ” shows the mIoU scores of 20 novel classes in fold *i*, “Mean” is the averaged mIoU score from all folds. * means the reproduced results. The best results are highlighted in **bold**.

Backbone Method		1-shot					5-shot						
		5 ⁰	5 ¹	5 ²	5 ³	Mean	FB-IoU	5 ⁰	5 ¹	5 ²	5 ³	Mean	FB-IoU
VGG16	PFENet (TPAMI'20) [35]	35.4	38.1	36.8	34.7	36.3	63.3	38.2	42.5	41.8	38.9	40.4	65.0
	FECANet (TMM'23) [20]	34.1	37.5	35.8	34.1	35.4	65.5	39.7	43.6	42.9	39.7	41.5	67.7
	BAM (CVPR'22) [14]	39.0	47.0	46.4	41.6	43.5	71.1	47.0	52.6	48.6	49.1	49.3	73.3
	HDMNet* (CVPR'23) [28]	42.6	49.6	46.7	44.7	45.9	71.4	48.0	55.8	54.7	50.5	52.2	74.8
	HDMNet + AENet	41.9	49.5	49.2	45.0	46.4 _{0.5} ↑	71.8	48.3	56.7	53.4	50.0	52.1 _{0.1} ↓	75.8
	SCCAN (ICCV'23) [45]	38.3	46.5	43.0	41.5	42.3	66.9	43.4	52.5	54.5	47.3	49.4	71.8
	SCCAN + AENet	40.3	50.4	47.9	44.9	45.9 _{3.6} ↑	71.2	45.8	56.3	55.8	53.4	52.8 _{3.4} ↑	74.3
ResNet50	PFENet (TPAMI'20)	36.5	38.6	35.0	33.8	35.8	-	36.5	43.3	38.0	38.4	39.0	-
	DACM (ECCV'22) [44]	37.5	44.3	40.6	40.1	40.6	68.9	44.6	52.0	49.2	46.4	48.1	71.6
	ABCNet (CVPR'23) [41]	42.3	46.2	46.0	42.0	44.1	69.9	45.5	51.7	52.6	46.4	49.1	72.7
	FECANet (TMM'23) [20]	38.5	44.6	42.6	40.7	41.6	69.6	44.6	51.5	48.4	45.8	47.6	71.1
	RiFeNet (AAAI'24) [1]	39.1	47.2	44.6	45.4	44.1	-	44.3	52.4	49.3	48.4	48.6	-
	BAM (CVPR'22) [14]	43.4	50.6	47.5	43.4	46.2	-	49.3	54.2	51.6	49.6	51.2	-
	MIANet (CVPR'23) [46]	42.5	53.0	47.8	47.4	47.7	71.5	45.8	58.2	51.3	51.9	51.7	73.1
	CyCTR (NIPS'21) [50]	38.9	43.0	39.6	39.8	40.3	-	41.1	48.9	45.2	47.0	45.6	-
	CyCTR + AENet	42.2	53.1	47.4	45.3	47.0 _{6.7} ↑	71.8	49.0	59.0	50.6	51.1	52.4 _{6.8} ↑	75.0
	SCCAN (ICCV'23) [45]	40.4	49.7	49.6	45.6	46.3	69.9	47.2	57.2	59.2	52.1	53.9	74.2
SCCAN + AENet	43.1	56.0	50.3	48.4	49.4 _{3.1} ↑	73.6	51.7	61.9	57.9	55.3	56.7 _{2.8} ↑	76.5	
HDMNet* (CVPR'23) [28]	44.8	54.9	50.0	48.7	49.6	72.1	50.9	60.2	55.0	55.3	55.3	74.9	
HDMNet + AENet	45.4	57.1	52.6	50.0	51.3 _{1.7} ↑	74.4	52.7	62.6	56.8	56.1	57.1 _{1.8} ↑	78.5	

does not behave as well as the best baseline HDMNet, but the incorporation of AENet can help it to outperform HDMNet by considerable margins, *e.g.*, 66.6% vs. 65.1% (VGG16, 1-shot) and 74.1% vs. 71.8% (ResNet50, 5-shot). Besides, HDMNet can be improved by 2.4% (ResNet50, 5-shot) with AENet.

COCO-20ⁱ. COCO-20ⁱ appears to be a more challenging dataset, as the images usually contain small objects, multiple objects, and their BG are quite complex. Unfortunately, such characteristics would make the aforementioned issues much more severe, *e.g.*, with the same receptive field, small FG objects are more likely to be mingled with more BG information, compared with large FG objects. The results are displayed in Tab. 2, and note that the results of HDMNet are reproduced by us, because their validation data lists of COCO-20ⁱ are different from the uniformly deployed ones in other baselines. Similar to PASCAL-5ⁱ, our AENet can well boost the performance of CyCTR, SCCAN and HDMNet. In particular, AENet can help to improve the baseline better on COCO-20ⁱ than on PASCAL-5ⁱ, *e.g.*, 3.6% vs. 1.9% with SCCAN (VGG16, 1-shot), and 6.7% vs. 4.8% with CyCTR (Res50, 1-shot). We contribute it to the fact that the FG-FG matching in baselines become less effective in COCO-20ⁱ, while the proposed AENet can effectively mitigate this issue.

5.3 Ablation Study

In this section, extensive ablation studies are conducted to validate the effectiveness of AENet. Unless explicitly specified, the experiments are conducted on PASCAL-5ⁱ with ResNet50 as the pretrained backbone, under 1-shot setting.

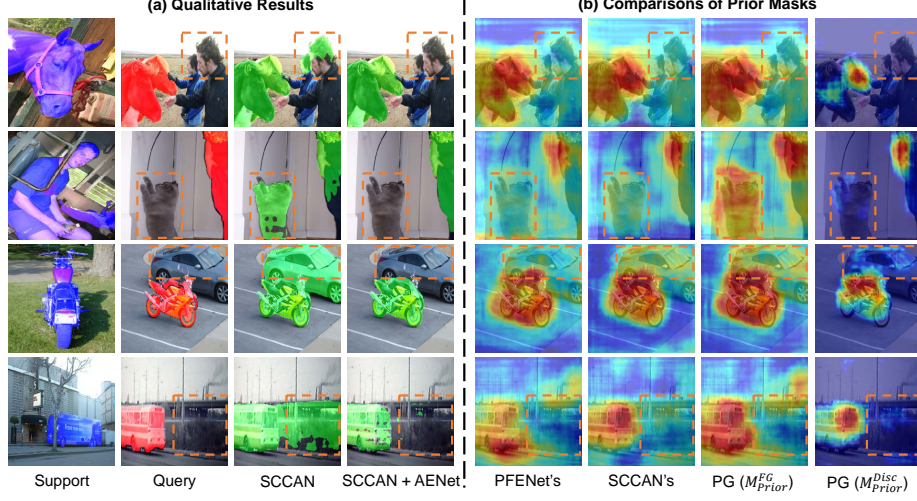


Fig. 5: Illustration of (a) qualitative results and (b) different training-agnostic prior masks. In (b), the proposed PG is compared with the existing prior masks from PFENet [35] and SCCAN [45]. We use some rectangles (in orange) to highlight some challenging areas, where existing prior masks wrongly activate them as FG and lead to wrong predictions, but our PG (M_{Prior}^{Disc}) can suppress them well.

Qualitative results and prior masks comparisons. To have a clearer understanding of the proposed modules to the final predictions, we jointly depict some qualitative results, as well as the visual comparisons among different prior masks in Fig. 5. Specifically, the proposed PG is compared with two existing prior mask generation methods from PFENet [35] and SCCAN [45]. Although they can roughly locate the query FG, they merely measure the similarities between each query pixel and all the support FG pixels to determine if the current query pixel is more likely to be FG or BG, regardless of the fact that the FG/BG features are mingled with BG/FG features (as explained in Sec. 1). In this way, some query BG pixels also contain query FG features, so there exist some similarities and they get wrongly activated. Instead, our PG not only measures the similarities with support FG (M_{Prior}^{FG}), but also with support BG (M_{Prior}^{BG}), then perform a subtraction operation to obtain the discriminative prior mask M_{Prior}^{Disc} . As shown in Fig. 5(b), M_{Prior}^{FG} suffer from the same problem as PFENet and SCCAN, but M_{Prior}^{Disc} can suppress the wrongly activated areas well. As a result, our plug-in AENet helps to make more accurate predictions.

Component-wise ablation study. BAM [14] is a controversial baseline that extends the standard FSS setting by: (1) They wrap the pretrained backbone with PSPNet [51], and fine-tune it for segmenting the base classes, which is more appropriate for FSS; (2) Then, the base class predictions are obtained and suppressed as BG regions during testing, which improves the accuracy. Despite of the controversy, most of the latest baselines [15, 28, 39, 46] follow BAM’s setting.

Table 3: Component-wise ablation study with SCCAN.

PG	AE	BAM	5 ⁰	5 ¹	5 ²	5 ³	Mean
			68.3	72.5	66.8	59.8	66.8
✓			69.7	73.1	66.9	61.5	67.8 _{1.0} ↑
	✓		69.3	72.8	67.4	62.1	67.9 _{1.1} ↑
✓	✓		70.1	73.4	66.9	63.1	68.3 _{1.5} ↑
✓		✓	71.4	75.1	68.0	62.4	69.2 _{2.4} ↑
	✓	✓	70.7	74.9	67.7	62.9	69.1 _{2.3} ↑
✓	✓	✓	72.2	75.5	68.5	63.1	69.8 _{3.0} ↑

Table 4: Detailed ablation study on PG and AE.

PG		AE	Eq. (6)	5 ⁰	5 ¹	5 ²	5 ³	Mean
M_{Prior}^{Disc}	M_{Prior}^{FG}							
				68.3	72.5	66.8	59.8	66.8
✓				70.7	73.2	65.7	60.4	67.5
✓	✓			69.7	73.1	66.9	61.5	67.8
		M_{Prior}^{FG}		69.0	72.2	66.5	59.8	66.9
		M_{Prior}^{Disc}		69.3	72.8	67.4	62.1	67.9

Hence, we experiment with both settings for better comparisons. It could be observed from Tab. 3 that the mean mIoU of SCCAN starts with 66.8%, and the score can be boosted to 67.8% and 67.9% after integrating PG and AE module, respectively. Then, when both PG and AE are utilized, the score is further increased to 68.3% (+1.5%), showing the effectiveness of the proposed AENet plug-in. With BAM’s ensemble, the final score can be as high as 69.8%.

Further discussion on PG. As described in Sec. 4.1, we generate three prior masks, including M_{Prior}^{FG} (measured with support FG features), M_{Prior}^{BG} (measured with support BG features), and M_{Prior}^{Disc} (clipped subtraction). Among them, M_{Prior}^{FG} is visually similar to the existing prior masks, while M_{Prior}^{Disc} appears to have better quality (as displayed in Fig. 5(b)). It can be observed from Tab. 4, when we replace SCCAN’s prior mask with M_{Prior}^{Disc} , the mean mIoU score can be boosted by 0.7%. If we further use M_{Prior}^{FG} , the final improvement can be 1.0%. We contribute the improvement to the different functions of these masks, *e.g.*, (1) M_{Prior}^{FG} is responsible for roughly locating the query FG regions. Although the generated prior masks also have high responses on BG regions, in most cases, the complete FG objects have already been included; and (2) for M_{Prior}^{Disc} , it only activates the most discriminative query FG regions that are dissimilar to the support BG features. Therefore, when they are both utilized, the model would firstly be aware of the regions it should focus on, then it can take the discriminative regions as anchor, and find other similar parts in the query image where there would be no intra-class differences [4, 26].

Detailed ablation on AE. We further conduct experiments to prove that the “subtraction” operation is critically important. As illustrated in the last two rows of Tab. 4, we use M_{Prior}^{Disc} and M_{Prior}^{FG} to distinguish the cases where there is a subtraction operation (Eq. (4)) or not. According to the table, we can observe that merely taking support FG information into consideration can hardly make any improvement (66.8% vs. 66.9%). Instead, if M_{Prior}^{Disc} is alternatively utilized, the improvement can be as high as 1.1%. Then, we explain this phenomenon as follows: (1) The query FG pixels’ features are doped with *base class-specific* BG features, because different classes tend to have their own sets of BG classes. Therefore, when the auxiliary loss (Eq. (7)) is applied for regularization, the model would be likely to learn *base class-specific* operations to detach the min-

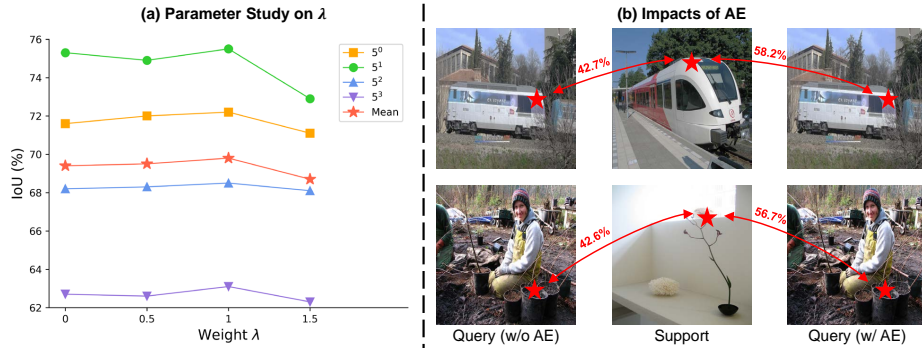


Fig. 6: Illustration of (a) parameter study on weight λ , and (b) the impacts of AE. In (b), \star represent FG pixels, the values are measured by cosine similarity.

gled BG features from query FG features. Thus, the model gets biased, and the learned pattern cannot be safely applied to the novel classes; (2) Instead, the subtraction operation (Eq. (4)) can provide the model with some *class-agnostic* guidance, so the learned pattern can be relatively uniform for all the classes.

Parameter study on loss weight λ . We set the loss weight λ (in Eq. (12)) as $\{0, 0.5, 1, 1.5\}$ to explore their effects, and show the results in Fig. 6(a). It could be observed: (1) Even if we do not apply additional supervision signals on the generated discriminative mask in AE, the mean mIoU score can already reach 69%+, showing the effectiveness of the discriminative mask; (2) When $\lambda = 1$, the best performance is achieved; (3) When $\lambda > 1$, the performance starts decreasing. Therefore, we set $\lambda = 1$ by default in the paper.

Impacts of AE. AE is designed to mitigate the *feature ambiguity* issue, so as to improve the FG-FG matching. To show the impacts of AE, we draw two examples in Fig. 6(b), where the features before and after the AE module are taken to measure the query-support similarity. As shown in the figure, AE module consistently helps to improve the FG-FG similarity between query and support FG pixels. In this way, query FG pixels can fuse more FG features from the support samples, *i.e.*, the support information is well utilized for more effective FSS.

6 Conclusion

In this paper, we identify the negative impacts of feature ambiguity to the cross attention modules in FSS. To alleviate the issue, we design a plug-in ambiguity elimination network (AENet) which includes a prior generator (PG) and an ambiguity eliminator (AE) module. Learning-agnostic PG is responsible for roughly locating the query FG objects, and highlighting the most discriminative query FG regions. AE utilizes the discriminative query FG features to rectify the features, so as to enhance the cross attentions. We plug AENet to three FSS baselines, and can improve their performance by large margins. Extensive experiments have been conducted to validate the effectiveness of AENet.

Acknowledgement. This study is supported under the RIE2020 Industry Alignment Fund - Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

References

1. Bao, X., Qin, J., Sun, S., Zheng, Y., Wang, X.: Relevant intrinsic feature enhancement network for few-shot semantic segmentation. arXiv preprint arXiv:2312.06474 (2023)
2. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence* **40**(4), 834–848 (2017)
3. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. *International journal of computer vision* **88**(2), 303–338 (2010)
4. Fan, Q., Pei, W., Tai, Y.W., Tang, C.K.: Self-support few-shot semantic segmentation. In: *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIX*. pp. 701–719. Springer (2022)
5. Garcia-Garcia, A., Orts-Escolano, S., Oprea, S., Villena-Martinez, V., Garcia-Rodriguez, J.: A review on deep learning techniques applied to semantic segmentation. arXiv preprint arXiv:1704.06857 (2017)
6. Guo, Y., Liu, Y., Georgiou, T., Lew, M.S.: A review of semantic segmentation using deep neural networks. *International journal of multimedia information retrieval* **7**, 87–93 (2018)
7. Hariharan, B., Arbeláez, P., Girshick, R., Malik, J.: Simultaneous detection and segmentation. In: *European conference on computer vision*. pp. 297–312. Springer (2014)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
9. Hong, S., Cho, S., Nam, J., Lin, S., Kim, S.: Cost aggregation with 4d convolutional swin transformer for few-shot segmentation. In: *European Conference on Computer Vision*. pp. 108–126. Springer (2022)
10. Hu, T., Yang, P., Zhang, C., Yu, G., Mu, Y., Snoek, C.G.: Attention-based multi-context guiding for few-shot semantic segmentation. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 33, pp. 8441–8448 (2019)
11. Iqbal, E., Safarov, S., Bang, S.: Msanet: Multi-similarity and attention guidance for boosting few-shot segmentation. arXiv preprint arXiv:2206.09667 (2022)
12. Jiao, S., Zhang, G., Navasardyan, S., Chen, L., Zhao, Y., Wei, Y., Shi, H.: Mask matching transformer for few-shot segmentation. arXiv preprint arXiv:2301.01208 (2022)
13. Kang, D., Koniusz, P., Cho, M., Murray, N.: Distilling self-supervised vision transformers for weakly-supervised few-shot classification & segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 19627–19638 (2023)
14. Lang, C., Cheng, G., Tu, B., Han, J.: Learning what not to segment: A new perspective on few-shot segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8057–8067 (2022)

15. Lang, C., Cheng, G., Tu, B., Li, C., Han, J.: Base and meta: A new perspective on few-shot segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023)
16. Lang, C., Tu, B., Cheng, G., Han, J.: Beyond the prototype: Divide-and-conquer proxies for few-shot segmentation. *arXiv preprint arXiv:2204.09903* (2022)
17. Lateef, F., Ruichek, Y.: Survey on semantic segmentation using deep learning techniques. *Neurocomputing* **338**, 321–348 (2019)
18. Li, G., Jampani, V., Sevilla-Lara, L., Sun, D., Kim, J., Kim, J.: Adaptive prototype learning and allocation for few-shot segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8334–8343 (2021)
19. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *European conference on computer vision*. pp. 740–755. Springer (2014)
20. Liu, H., Peng, P., Chen, T., Wang, Q., Yao, Y., Hua, X.S.: Fecanet: Boosting few-shot semantic segmentation with feature-enhanced context-aware network. *IEEE Transactions on Multimedia* (2023)
21. Liu, J., Bao, Y., Xie, G.S., Xiong, H., Sonke, J.J., Gavves, E.: Dynamic prototype convolution network for few-shot semantic segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11553–11562 (2022)
22. Liu, Y., Liu, N., Cao, Q., Yao, X., Han, J., Shao, L.: Learning non-target knowledge for few-shot semantic segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11573–11582 (2022)
23. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3431–3440 (2015)
24. Luo, X., Tian, Z., Zhang, T., Yu, B., Tang, Y.Y., Jia, J.: Pfenet++: Boosting few-shot semantic segmentation with the noise-filtered context-aware prior mask. *arXiv preprint arXiv:2109.13788* (2021)
25. Nguyen, K., Todorovic, S.: Feature weighting and boosting for few-shot segmentation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 622–631 (2019)
26. Okazawa, A.: Interclass prototype relation for few-shot segmentation. In: *European Conference on Computer Vision*. pp. 362–378. Springer (2022)
27. Park, S., Lee, S., Hyun, S., Seong, H.S., Heo, J.P.: Task-disruptive background suppression for few-shot segmentation. *arXiv preprint arXiv:2312.15894* (2023)
28. Peng, B., Tian, Z., Wu, X., Wang, C., Liu, S., Su, J., Jia, J.: Hierarchical dense correlation distillation for few-shot segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 23641–23651 (2023)
29. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* 18. pp. 234–241. Springer (2015)
30. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International journal of computer vision* **115**(3), 211–252 (2015)
31. Shaban, A., Bansal, S., Liu, Z., Essa, I., Boots, B.: One-shot learning for semantic segmentation. *arXiv preprint arXiv:1709.03410* (2017)

32. Shi, X., Wei, D., Zhang, Y., Lu, D., Ning, M., Chen, J., Ma, K., Zheng, Y.: Dense cross-query-and-support attention weighted mask aggregation for few-shot segmentation. In: European Conference on Computer Vision. pp. 151–168. Springer (2022)
33. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
34. Sun, Y., Chen, Q., He, X., Wang, J., Feng, H., Han, J., Ding, E., Cheng, J., Li, Z., Wang, J.: Singular value fine-tuning: Few-shot segmentation requires few-parameters fine-tuning. arXiv preprint arXiv:2206.06122 (2022)
35. Tian, Z., Zhao, H., Shu, M., Yang, Z., Li, R., Jia, J.: Prior guided feature enrichment network for few-shot segmentation. IEEE transactions on pattern analysis and machine intelligence (2020)
36. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)
37. Wang, H., Zhang, X., Hu, Y., Yang, Y., Cao, X., Zhen, X.: Few-shot semantic segmentation with democratic attention networks. In: European Conference on Computer Vision. pp. 730–746. Springer (2020)
38. Wang, J., Li, J., Chen, C., Zhang, Y., Shen, H., Zhang, T.: Adaptive fss: a novel few-shot segmentation framework via prototype enhancement. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 5463–5471 (2024)
39. Wang, J., Li, J., Chen, C., Zhang, Y., Shen, H., Zhang, T.: Adaptive fss: A novel few-shot segmentation framework via prototype enhancement. arXiv preprint arXiv:2312.15731 (2023)
40. Wang, K., Liew, J.H., Zou, Y., Zhou, D., Feng, J.: Panet: Few-shot image semantic segmentation with prototype alignment. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9197–9206 (2019)
41. Wang, Y., Sun, R., Zhang, T.: Rethinking the correlation in few-shot segmentation: A buoys view. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7183–7192 (2023)
42. Wang, Y., Sun, R., Zhang, Z., Zhang, T.: Adaptive agent transformer for few-shot segmentation. In: European Conference on Computer Vision. pp. 36–52. Springer (2022)
43. Xie, G.S., Liu, J., Xiong, H., Shao, L.: Scale-aware graph neural network for few-shot semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5475–5484 (2021)
44. Xiong, Z., Li, H., Zhu, X.X.: Doubly deformable aggregation of covariance matrices for few-shot segmentation. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XX. pp. 133–150. Springer (2022)
45. Xu, Q., Zhao, W., Lin, G., Long, C.: Self-calibrated cross attention network for few-shot segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 655–665 (2023)
46. Yang, Y., Chen, Q., Feng, Y., Huang, T.: Mianet: Aggregating unbiased instance and general information for few-shot semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7131–7140 (2023)
47. Zhang, B., Xiao, J., Qin, T.: Self-guided and cross-guided learning for few-shot segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8312–8321 (2021)

48. Zhang, C., Lin, G., Liu, F., Guo, J., Wu, Q., Yao, R.: Pyramid graph networks with connection attentions for region-based one-shot semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9587–9595 (2019)
49. Zhang, C., Lin, G., Liu, F., Yao, R., Shen, C.: Canet: Class-agnostic segmentation networks with iterative refinement and attentive few-shot learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5217–5226 (2019)
50. Zhang, G., Kang, G., Yang, Y., Wei, Y.: Few-shot segmentation via cycle-consistent transformer. *Advances in Neural Information Processing Systems* **34**, 21984–21996 (2021)
51. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2881–2890 (2017)
52. Zhou, T., Wang, W., Konukoglu, E., Van Gool, L.: Rethinking semantic segmentation: A prototype view. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2582–2593 (2022)
53. Zhu, L., Chen, T., Ji, D., Ye, J., Liu, J.: Llafs: When large-language models meet few-shot segmentation. *arXiv preprint arXiv:2311.16926* (2023)
54. Zhu, L., Chen, T., Yin, J., See, S., Liu, J.: Addressing background context bias in few-shot segmentation through iterative modulation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3370–3379 (2024)