

# Cross-Domain Few-Shot Segmentation via Multi-view Progressive Adaptation

Jiahao Nie<sup>1,2\*</sup> Guanqiao Fu<sup>2\*</sup> Wenbin An<sup>3</sup> Yap-Peng Tan<sup>4,2</sup> Alex C. Kot<sup>5,4,2</sup> Shijian Lu<sup>2†</sup>

<sup>1</sup>Interdisciplinary Graduate Programme, Nanyang Technological University

<sup>2</sup>Nanyang Technological University <sup>3</sup>Xi'an Jiaotong University <sup>4</sup>VinUniversity <sup>5</sup>SMBU

jiahao007@e.ntu.edu.sg

shijian.lu@ntu.edu.sg

## Abstract

*Cross-Domain Few-Shot Segmentation aims to segment categories in data-scarce domains conditioned on a few exemplars. Typical methods first establish few-shot capability in a large-scale source domain and then adapt it to target domains. However, due to the limited quantity and diversity of target samples, existing methods still exhibit constrained performance. Moreover, the source-trained model's initially weak few-shot capability in target domains, coupled with substantial domain gaps, severely hinders the effective utilization of target samples and further impedes adaptation. To this end, we propose Multi-view Progressive Adaptation, which progressively adapts few-shot capability to target domains from both data and strategy perspectives. (i) From the data perspective, we introduce Hybrid Progressive Augmentation, which progressively generates more diverse and complex views through cumulative strong augmentations, thereby creating increasingly challenging learning scenarios. (ii) From the strategy perspective, we design Dual-chain Multi-view Prediction, which fully leverages these progressively complex views through sequential and parallel learning paths under extensive supervision. By jointly enforcing prediction consistency across diverse and complex views, MPA achieves both robust and accurate adaptation to target domains. Extensive experiments demonstrate that MPA effectively adapts few-shot capability to target domains, outperforming state-of-the-art methods by a large margin (+7.0%). Our code is available at <https://github.com/niejiahao1998/MPA>.*

## 1. Introduction

Leveraging the paradigm of meta-learning [5, 21, 25, 26, 62, 73], Few-Shot Segmentation (FSS) has achieved great success by learning and adapting few-shot capability from base

\*Equal contribution

†Corresponding author

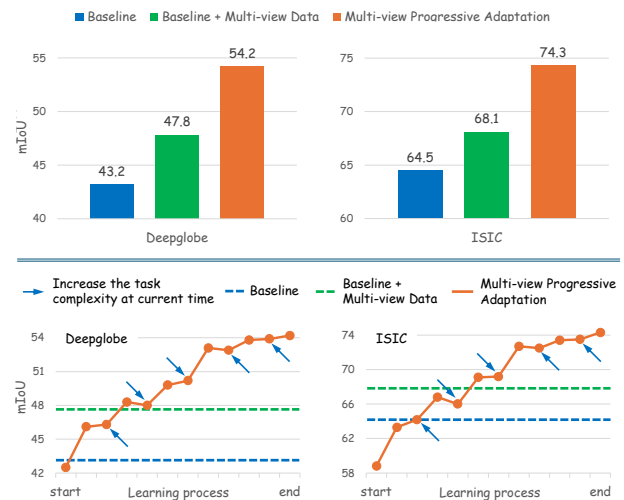


Figure 1. **Up:** Simply incorporating multiple augmented views from the accessible target samples increases the sample available for establishing few-shot capability but yields only marginal gains, as the large domain gap limits effective utilization of heavily perturbed views. In contrast, our proposed Multi-View Progressive Adaptation (MPA) significantly improves the performance. **Down:** MPA adopts a progressive strategy to address the challenges posed by large domain gaps. Specifically, it starts with an easy task and progressively increases task complexity as the model becomes more capable during adaptation. This design enables a smooth adaptation of source-trained model and effectively establishes few-shot capability in the target domains.

categories to novel categories [10, 13, 38, 43, 50, 52, 61]. FSS typically requires abundant training samples from base categories to establish the capability for segmenting query images conditioned on a few provided support images. However, acquiring abundant base-category samples is often challenging in various data-scarce domains such as medical images [2, 7] and satellite images [8]. Enabling FSS in data-scarce domains is a critical and valuable endeavor, driven by its broad range of real-world applications.

Cross-Domain Few-Shot Segmentation (CD-FSS) [19, 31, 40, 51, 53] has been explored to tackle this challenge by

leveraging a large-scale source domain to obtain the abundant training samples of base categories. CD-FSS methods typically adopt two learning stages: (i) meta-training over a large-scale source domain (such as Pascal VOC [11]) with abundant base-category samples for establishing few-shot capability; and (ii) adapting this capability to each data-scarce target domain with limited accessible exemplars. However, this paradigm faces two critical challenges: (i) only extremely limited samples in target domains is available for adaptation; and (ii) substantial domain gaps between the source and target domains hinder the effective adaptation of the few-shot capability. Consequently, the baseline method [31] yields limited performance in target domains, as shown in Fig. 1 (Up).

Motivated by the observation that CD-FSS methods typically achieve higher performance in multi-shot settings than in the 1-shot setup, we attempt to tackle the aforementioned challenges by augmenting more views from the accessible samples for adaptation. As illustrated in Fig. 1 (Up), incorporating multiple views yields performance gains but still remains suboptimal. This limitation stems from the large domain gaps, which hinder the establishment of few-shot capability under heavily perturbed augmentations during the early adaptation stage, when the source-trained model exhibits weak few-shot capability in target domains. Consequently, the effective utilization of augmented views is restricted, leading to limited performance. These observations highlight the necessity of designing an appropriate adaptation strategy to fully exploit augmented views.

Inspired by the success of progressive strategies in domain generalization [34, 46, 58, 75], we propose Multi-view Progressive Adaptation (MPA) to gradually tackle CD-FSS challenges. Specifically, MPA progressively increases task complexity and diversity to establish few-shot capability from both the data and strategy perspectives. **First**, from the data perspective, we design Hybrid Progressive Augmentation (HPA) that progressively increases the complexity and diversity of the augmented views as adaptation proceeds and the model becomes more capable: (i) generating more challenging query views through the accumulation of multiple strong augmentation strategies; (ii) increasing the number of augmented views to establish support–query correspondences under more diverse and challenging situations. **Second**, from the strategy perspective, we develop Dual-chain Multi-view Prediction (DMP) to effectively exploit the progressively complex augmented views, which learns through two complementary prediction chains: (i) a sequential chain, which progressively extends support–query correspondences [40] across all augmented views, where prediction errors are propagated and accumulated in later views; (ii) a parallel chain, which performs multiple “support-to-query” predictions across the augmented views, where diverse predictions lead to var-

ied errors. DMP applies supervision to regularize these accumulated and diverse errors, thereby facilitating the establishment of few-shot capability. Thanks to these designs, MPA progressively brings performance gains with each increment in task complexity and successfully establishes few-shot capability in target domains by the end of adaptation (Fig. 1 (Down)). Extensive experiments along with in-depth analyses demonstrate the superiority of MPA, which effectively addresses the challenges of domain gaps and limited data accessibility in CD-FSS.

The contributions of this work can be summarized in three aspects: **First**, We identify two key constraints in CD-FSS task: (i) the accessible target samples for adaptation is limited in both quantity and diversity; and (ii) the source-trained model exhibits weak few-shot capability in target domains due to substantial large domain gaps. **Second**, we propose Multi-view Progressive Adaptation (MPA), which progressively augments views with increasing complexity and diversity from accessible samples and establishes few-shot capability in target domains from both data and strategy perspectives. **Third**, extensive experiments demonstrate that MPA achieves strong performances over state-of-the-art methods (+7.0%) and enhances learning efficiency.

## 2. Related Work

**Few-Shot Segmentation (FSS)** [6, 37, 39, 45, 48, 63–65, 68, 72, 76], a topic that has been extensively studied, focuses on segmenting novel categories using only a few support images. Existing methods can generally be categorized into two types. Prototype-based methods [10, 13, 29, 30, 32, 35, 50, 52, 56] segment masks by computing similarities between all query features and support prototypes. In contrast, affinity-based methods [36, 38, 43, 72] establish dense correspondence between query and support, heavily relying on rich contextual information. While these approaches are well-established, their robustness under cross-domain setups remains suboptimal [40, 59].

**Cross-Domain Few-Shot Segmentation (CD-FSS)** tackles the FSS setup applied in data-scarce target domains [1, 3, 4, 15, 31, 36, 54, 57, 60], aiming to adapt models trained on a large-scale source domain to diverse data-scarce target domains. CD-FSS presents significant challenges due to the following factors: (i) during the adaptation stage, the limited target samples increases the risk of overfitting; and (ii) the source and target domains exhibit substantial domain gaps. Early works explore this challenge through approaches from dynamic adaptation refinement [12] and knowledge-transfer [22] aspects, but these methods exhibit limited generalization capabilities. Recently, efforts have been made to adapt source-trained models by disentangling the feature frequency [53] and by fine-tuning specific structures rather than the entire model to improve the robustness and mitigate overfitting [14, 51]. Furthermore, some recent

Table 1. The mIoU (%) of views with more complex augmentation operations is consistently lower than views with simpler augmentations, indicating that the task becomes more difficult.

Augmentation	Deepglobe	ISIC
Flip	<b>53.1</b>	<b>71.1</b>
Flip+hue variation	47.9 <sub>↓</sub>	69.8 <sub>↓</sub>
Flip+hue variation+brightness change	44.5 <sub>↓↓</sub>	64.4 <sub>↓↓</sub>

works [19, 41, 44, 67, 74] leverage the strong generalization capability of SAM [27], demonstrating improved performance. However, ABCDFSS [20] argues that training on the source domain may hinder the development of few-shot capability in the target domain by introducing additional domain gaps. Therefore, this work places greater emphasis on the adaptation stage and explores how to leverage the limited accessible samples to adapt the source-trained model and establish few-shot capability in target domains.

### 3. Preliminary Study

#### 3.1. Problem Formulation

Few-Shot Segmentation aims to segment novel categories using category-agnostic knowledge learned from abundant base-category samples. However, obtaining sufficient training samples is infeasible in data-scarce domains, leading to a two-stage Cross-Domain Few-Shot Segmentation (CD-FSS) paradigm: It first meta-trains the backbone model with a large-scale source domain  $\mathcal{D}_{\text{source}}$  to establish few-shot capability, and then adapts the few-shot capability to each data-scarce target domain  $\mathcal{D}_{\text{target}}$  separately. More details are in the supplementary materials.

Following the meta-learning framework [31], we utilize an episodic training strategy. For the  $K$ -shot setting, each episode consists of a support set  $S = \{(I_s^i, M_s^i)\}_{i=1}^K$  and a query set  $Q = \{(I_q, M_q)\}$ , containing samples from the same category. Here,  $I_s$  and  $I_q$  represent support and query images, while  $M_s$  and  $M_q$  denote their ground-truth masks. The framework consists of two stages: (i) training the backbone model in  $\mathcal{D}_{\text{source}}$  with both  $S$  and  $Q$  sets; and (ii) fine-tuning the trained model to  $\mathcal{D}_{\text{target}}$  with  $S$  only.

#### 3.2. Comparison of Difficulty across Multiple Views

We increase the task difficulty through two designs: stronger data augmentations and sequential predictions across multiple views. The verification is provided below.

**Difficulty definition.** We define difficulty from a quantitative perspective. Specifically, we perform “support-to-query” inference with different query images, where lower segmentation performance indicates higher task difficulty.

**Stronger data augmentations.** We hypothesize that cumulatively applying more complex augmentations introduces stronger perturbations, thereby increasing the difficulty of predicting accurate masks on the augmented views. For example, an easy view involves only a horizontal flip, a harder

Table 2. The mIoU (%) of later views in the sequential prediction chain is consistently lower than that of earlier views, indicating that errors are propagated and accumulated.

Dataset	1 <sup>st</sup> view	3 <sup>rd</sup> view	6 <sup>th</sup> view
Deepglobe	<b>53.1</b>	52.0 <sub>↓</sub>	49.6 <sub>↓↓</sub>
ISIC	<b>71.1</b>	70.1 <sub>↓</sub>	68.5 <sub>↓↓</sub>

view combines flipping with hue variation, and adding brightness adjustment further increases the difficulty. We experimentally validate this assumption in Tab. 1.

**Sequential prediction across multiple views.** We assume that errors accumulate and propagate across sequential predictions. Specifically, the prediction for the first query is guided by the support image, while subsequent queries are conditioned on the predictions of their preceding ones. This assumption is experimentally validated in Tab. 2, showing that sequentially leveraging multiple augmented views provides a promising way to increase task difficulty.

### 4. Method

To adapt the source-trained model and progressively establish few-shot capability in data-scarce target domains [58], we propose Multi-view Progressive Adaptation using multiple augmented views from accessible samples [47] with well-designed strategies, as illustrated in Fig. 2. The proposed framework consists of two major designs: (i) Hybrid Progressive Augmentation (HPA): progressively augment more views with increasing complexity and diversity as the adaptation proceeds; and (ii) Dual-chain Multi-view Prediction (DMP): establish few-shot capability through both sequential and parallel chains with augmented views. The basic pipeline can be formulated as follows: Given an input support image and its corresponding mask  $\{(I_s, M_s)\}$ , we generate  $N$  augmented query images with masks, denoted as  $\{(I_{q_i}, M_{q_i})\}_{i=1}^N$ , through adaptive augmentation operations upon progressive criteria. Next, all support and query images are processed through a weight-shared encoder to extract features  $\{F_s, \{F_{q_i}\}_{i=1}^N\}$ . Then, the support feature  $F_s$  and its mask  $M_s$  are processed using masked average pooling to generate the support prototype  $P_s$ . Finally,  $P_s$ ,  $F_s$ , and  $\{F_{q_i}\}_{i=1}^N$  are utilized to progressively establish few-shot capability in data-scarce target domains. We elaborate our designs as follows: HPA and DMP are detailed in Sec. 4.1 and Sec. 4.2, respectively. Furthermore, the loss functions are explained in Sec. 4.3. Finally, extension to  $K$ -shot setting is in Sec. 4.4.

#### 4.1. Hybrid Progressive Augmentation

In the target domains, only one support image  $I_s$  and corresponding ground-truth mask  $M_s$  are accessible under the 1-shot setting, thus  $N$  query images  $\{I_{q_i}\}_{i=1}^N$  with corresponding labels  $\{M_{q_i}\}_{i=1}^N$  are derived from support image via augmentation operations:

$$\{(I_{q_i}, M_{q_i})\}_{i=1}^N = \text{AUG}(I_s), \text{AUG}(M_s). \quad (1)$$

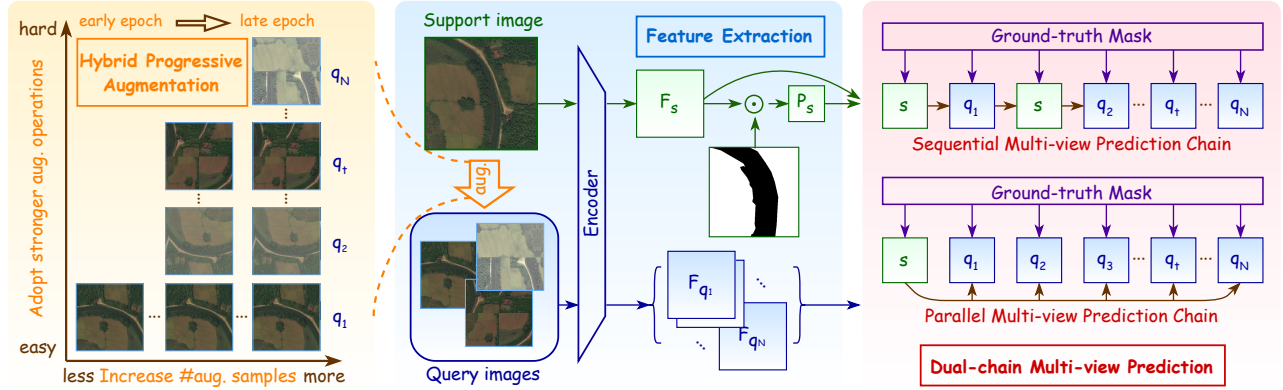


Figure 2. The framework of the proposed Multi-view Progressive Adaptation (MPA). MPA starts with Hybrid Progressive Augmentation (HPA) as highlighted in the yellow box, which introduces two strategies to progressively increase the complexity and diversity of augmented query images. Leveraging the augmented views, MPA establishes few-shot capabilities in target domains via Dual-chain Multi-view Prediction (DMP), incorporating sequential and parallel chains as highlighted in the red box. Note  $s$  and  $q_i$  denote support and  $i^{th}$  query images,  $F_s$  and  $F_{q_i}$  denote the support and the  $i^{th}$  query features as extracted by the encoder, and  $P_s$  denotes the support prototype.

The progressive strategy refers to “learning from easy to more complex conditions”, which facilitates gradual establishment of few-shot capability in target domains that exhibit large gaps from the source domain [58]. We propose Hybrid Progressive Augmentation (HPA) to generate more complex query views as the adaptation proceeds and the model becomes more capable, as illustrated in Fig. 2 (yellow box). Specifically, we increase the task complexity through two designs:

**Augment more challenging views.** As the adaptation proceeds, we progressively apply stronger augmentation operations. At the beginning, we apply simple augmentations such as flipping to generate the first query image  $I_{q_1}$ . Subsequently, increasingly complex transformations are introduced to generate subsequent query images, ensuring that augmentations across views become progressively more complex through cumulative operations. For example,  $I_{q_1}$  uses only flipping,  $I_{q_2}$  further adds brightness adjustment, ..., and  $I_{q_6}$  combines all previous operations and grid shuffle together (refer to Fig. 2 (yellow box)).

**Increase the number of views.** Augmenting additional views and establishing support-query correspondence among all of them implicitly increases the challenge, as prediction errors accumulate and propagate to subsequent views (refer to Tab. 2).  $N$  is set to 1 at the beginning, only requiring the model to predict accurate mask of  $I_{q_1}$  based on  $I_s$ . As the adaptation proceeds,  $N$  is adaptively increased, allowing for the augmentation of more query images. Subsequently, the model is expected to accurately predict across all  $N$  augmented query images  $\{I_{q_i}\}_{i=1}^N$ .

## 4.2. Dual-chain Multi-view Prediction

Even though we obtain  $N$  query images, establishing few-shot capability remains challenging, as an effective strategy for leveraging these augmented views is essential. Consequently, we design a dual-chain strategy to comprehensively

exploit these views: a sequential multi-view prediction chain and a parallel one (refer to Fig. 2 (red box)).

**Sequential multi-view predictions and supervisions.** It is straightforward to propose a chain that includes  $I_s$  and all  $\{I_{q_i}\}_{i=1}^N$ , requiring the model to predict accurate masks across all of them sequentially. Based on the success of [14, 40], we incorporate SSP [13] to enhance prediction accuracy. Support prototype  $P_s$  and prototype of the first query  $P_{q_1}^{seq}$  are derived from support feature  $F_s$ , support mask  $M_s$ , and query feature  $F_{q_1}$ <sup>1</sup>:

$$P_s = MAP(F_s, M_s), \quad P_{q_1}^{seq} = SSP(F_{q_1}, P_s), \quad (2)$$

where  $MAP$  is masked average pooling operation. More details of SSP are in the supplementary materials.

Next, we predict support mask  $\hat{M}_s$  and mask  $\hat{M}_{q_1}^{seq}$  for the first query image:

$$\hat{M}_s = \sigma(\cos(F_s, P_s)), \quad \hat{M}_{q_1}^{seq} = \sigma(\cos(F_{q_1}, P_{q_1}^{seq})), \quad (3)$$

where  $\cos$  is cosine similarity, and  $\sigma$  represents softmax function. Concurrently, support base loss  $\mathcal{L}_{bs}$  and sequential query loss  $\mathcal{L}_{q_1}^{seq}$  for the first query image are computed:

$$\mathcal{L}_{bs} = \mathcal{L}_{bce}(\hat{M}_s, M_s), \quad \mathcal{L}_{q_1}^{seq} = \mathcal{L}_{bce}(\hat{M}_{q_1}^{seq}, M_{q_1}), \quad (4)$$

where  $\mathcal{L}_{bce}$  is the binary cross entropy loss. Drawing inspiration from the effectiveness of bi-directional predictions in alleviating overfitting [40, 72], we also incorporate reverse prediction, where the roles of the support and query images are exchanged, with the query image guiding the segmentation of the support:

$$P_s^{seqq_1} = SSP(F_s, P_{q_1}^{seq}), \quad \hat{M}_s^{seqq_1} = \sigma(\text{Sim}(F_s, P_s^{seqq_1})), \\ \mathcal{L}_s^{seqq_1} = \mathcal{L}_{bce}(\hat{M}_s^{seqq_1}, M_s), \quad (5)$$

<sup>1</sup>We use the superscript “seq” (abbr. for sequential) to indicate that the variables or results are obtained within the sequential prediction chain. Similarly, the superscript “par” (abbr. for parallel) signifies that variables or results are from the parallel chain.

where  $P_s^{seq_1}$  is another support prototype predicted from  $I_{q_1}$ <sup>2</sup>.  $\mathcal{L}_s^{seq_1}$  provides additional regularization, aiding in establishing few-shot capability with limited data. With additional query images introduced by HPA, the above predictions can be seamlessly extended into a sequential chain.

Specifically, for the  $j^{th}$  ( $1 \leq j \leq N$ ) query image  $I_{q_j}$ , we obtain  $\mathcal{L}_s^{seq}$ , and  $\mathcal{L}_s^{seq_j}$  from  $P_s^{seq(j-1)}$  ( $P_s^{seq_0} = P_s$  in Eq. 2), with details in the supplementary materials. The results from Eqs. 2, 3 and 5 correspond to the case when  $j = 1$ . Ideally, a well-trained model should make accurate predictions on all queries sequentially [40]. Otherwise, errors in earlier predictions propagate, affecting subsequent results, while the corresponding loss provides additional regularization for target domains.

**Parallel multi-view predictions and supervisions.** To directly mimic ‘‘support-to-query’’ predictions during inference [20, 31, 40] and further establish correspondences between the accessible support image  $I_s$  and query images  $\{I_{q_i}\}_{i=1}^N$ , we propose a parallel multi-view prediction chain that operates with the sequential chain simultaneously. Specifically, it leverages  $I_s$  to segment  $\{I_{q_i}\}_{i=1}^N$  in parallel ( $1 \leq i \leq N$ ) and includes the reverse predictions. Then, ground-truth masks provide supervision signals:

$$\begin{aligned} P_{q_i}^{par} &= SSP(F_{q_i}, P_s), \quad \hat{M}_{q_i}^{par} = \sigma(\cos(F_{q_i}, P_{q_i}^{par})), \\ P_s^{par_i} &= SSP(F_s, P_{q_i}^{par}), \quad \hat{M}_s^{par_i} = \sigma(\cos(F_s, P_s^{par_i})), \\ \mathcal{L}_{q_i}^{par} &= \mathcal{L}_{bce}(\hat{M}_{q_i}^{par}, M_{q_i}), \quad \mathcal{L}_s^{par_i} = \mathcal{L}_{bce}(\hat{M}_s^{par_i}, M_s). \end{aligned} \quad (6)$$

### 4.3. Loss Function

**Sequential loss.** We would like to clarify that  $\mathcal{L}_s^{seq}$  and  $\mathcal{L}_s^{seq_1}$  are identical to  $\mathcal{L}_{q_1}^{par}$  and  $\mathcal{L}_s^{par_1}$ , respectively. Therefore, our proposed sequential loss starts from index 2. The supervision from the sequential prediction chain can be summarized as follows:

$$\mathcal{L}^{seq} = \sum_{i=2}^N (\mathcal{L}_{q_i}^{seq} + \mathcal{L}_s^{seq_i}). \quad (7)$$

**Parallel loss.** Similarly, the parallel supervision consists of two parts, parallel support loss and parallel query loss:

$$\mathcal{L}_s^{par} = \sum_{i=1}^N \mathcal{L}_s^{par_i}, \quad \mathcal{L}_q^{par} = \sum_{i=1}^N \mathcal{L}_{q_i}^{par}. \quad (8)$$

**Total loss.** The total loss combines and balances the previous loss terms  $\mathcal{L}_{bs}$ ,  $\mathcal{L}_s^{seq}$ ,  $\mathcal{L}_s^{par}$ , and  $\mathcal{L}_q^{par}$ :

$$\begin{aligned} \mathcal{L} &= \lambda_{bs} \times \mathcal{L}_{bs} + \lambda^{seq} \times \mathcal{L}^{seq} \\ &+ \lambda_s^{par} \times \mathcal{L}_s^{par} + \lambda_q^{par} \times \mathcal{L}_q^{par}, \end{aligned} \quad (9)$$

<sup>2</sup>A pseudo prototype of the support image,  $P_s^{seq_1}$ , can be predicted from the predicted pseudo prototype  $P_{q_1}^{seq}$  of the first query image.  $P_s^{seq_1}$  differs from  $P_s$ , which is obtained from the ground-truth support mask  $M_s$ . We include  $i$  in the superscript to indicate that this support prototype is derived from the predicted result of the  $i^{th}$  query image.

where  $\lambda_{bs} = 0.2$ ,  $\lambda^{seq} = 0.1$ ,  $\lambda_s^{par} = 0.4$ , and  $\lambda_q^{par} = 1$ . The determination of the values of these parameters is discussed in the supplementary materials. The overall MPA algorithm is presented step by step in the supplementary materials.

### 4.4. Extension to $K$ -shot Setting

In the extension to the  $K$ -shot ( $K > 1$ ) setting,  $K$  support images with their corresponding masks, denoted as  $S = \{(I_s^i, M_s^i)\}_{i=1}^K$ , are provided for adaptation. All augmented queries  $\{(I_{q_i}, M_{q_i})\}_{i=1}^N$  are derived from a randomly selected support image-mask pair  $\{(I_s^i, M_s^i)\}$ . MPA can be efficiently extended to this setting as follows. As described in Sec. 4.2, the initial prediction utilizes the averaged support prototype  $\bar{P}_s = \frac{1}{K} \sum_{i=1}^K P_s^i$  to predict  $\hat{M}_q$ . In the reversed prediction procedure, we use pseudo prototype  $P_{q_i}$  of  $i^{th}$  query image to predict each  $\hat{M}_s^i$  in parallel.

## 5. Experiment

### 5.1. Datasets

We conduct extensive experiments over five data-scarce target datasets, covering satellite images [8], medical screenings [2, 7, 24, 55], tiny objects [33], and underwater scenes [23]. Fig. 3 shows sample images from these datasets. **Deepglobe** [8] is a satellite image dataset containing 6 terrain categories, such as urban, agriculture, rangeland, forest, water, and barren. For CD-FSS, we follow PATNet [31] to split images into smaller pieces. **ISIC2018** [7, 55] comprises medical images of skin lesions. It captures three types of skin lesions. **Chest X-Ray** [2, 24] is collected for Tuberculosis test. Its grayscale images enhance the diversity of evaluation. **FSS-1000** [33] comprises natural images of everyday objects. It poses significant challenges due to its scarce samples and tiny objects in images. **SUIM** [23] consists of 7 categories of underwater objects, including fish, plants, divers, robots, ruins, and rocks.

### 5.2. Implementation Details

We adopt the ResNet-50 [18] pre-trained on ImageNet [9] as the backbone model. Consistent with previous works [13, 40], we discard the last stage and last ReLU for better generalization. The model is implemented with PyTorch [42]. Following PATNet [31], we resize all images to  $400 \times 400$  to reduce memory consumption and accelerate the learning. The learning rate is set as  $5e-4$  for all datasets. For data augmentation in Sec. 4.1, we adopt a set of transformations in PyTorch [42], including horizontal-flip, vertical-flip, 90-degree rotation, brightness-variation, and hue-variation. We adopt an adaptive criterion to dynamically increase the task complexity during progressive adaptation. Specifically, when the performance stagnates for three consecutive epochs, we regard it as performance saturation and intro-

Table 3. Quantitative comparisons between the proposed MPA and existing methods over four widely adopted data-scarce domains (in mIoU (%)). The best mIoU number under each setup is highlighted by **bold** font.

Backbone	Method	Deepglobe		ISIC		Chest X-Ray		FSS-1000		Average	
		1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
Vgg-16	AMP [49]	37.6	40.6	28.4	30.4	51.2	53.0	57.2	59.2	43.6	45.8
	PATNet [31]	28.7	34.8	33.1	45.8	57.8	60.6	71.6	76.2	47.8	54.4
Res-50	PGNet [70]	10.7	12.4	21.9	21.3	34.0	28.0	62.4	62.7	32.2	31.1
	PANet [56]	36.6	45.3	25.3	34.0	57.8	69.3	69.2	71.7	47.2	55.1
	CaNet [71]	22.3	23.1	25.2	28.2	28.4	28.6	70.7	72.0	36.6	38.0
	RPMMs [66]	13.0	13.5	18.0	20.0	30.1	30.8	65.1	67.1	31.6	32.9
	PFENet [52]	16.9	18.0	23.5	23.8	27.2	27.6	70.9	70.5	34.6	35.0
	RePRI [1]	25.0	27.4	23.3	26.2	65.1	65.5	71.0	74.2	46.1	48.3
	HSNet [38]	29.7	35.1	31.2	35.1	51.9	54.4	77.5	81.0	47.6	51.4
	PATNet [31]	37.9	43.0	41.2	53.6	66.6	70.2	78.6	81.2	56.1	62.0
	SSP [13]	41.3	54.2	48.6	65.4	72.6	73.0	77.0	79.4	60.0	68.0
	DR-Adapter [51]	41.3	50.1	40.8	48.9	82.4	82.3	79.1	80.4	60.9	65.4
	IFA [40]	50.6	58.8	66.3	69.8	74.0	74.6	80.1	<b>82.4</b>	67.8	71.4
	ABCFSS [20]	42.6	49.0	45.7	53.3	79.8	81.4	74.6	76.2	60.7	65.0
	<b>MPA (w/ Source-Training)</b>	<b>54.2</b>	<b>60.8</b>	<b>74.3</b>	<b>74.4</b>	<b>89.1</b>	<b>91.0</b>	<b>81.4</b>	81.4	<b>74.8</b>	<b>76.9</b>
	<b>MPA (w/o Source-Training)</b>	<b>53.1</b>	<b>59.4</b>	<b>71.1</b>	<b>71.3</b>	<b>89.0</b>	<b>90.6</b>	<b>80.2</b>	80.8	<b>73.4</b>	<b>75.5</b>

Table 4. Quantitative comparisons between the proposed MPA and existing methods over SUIM (in mIoU (%)). The best mIoU number under each setup is highlighted by **bold** font.

Method	1-shot	5-shot
HSNet [38]	28.8	-
SCL [69]	31.8	-
HDMNet [43]	23.4	30.9
RemDiff [57]	34.7	-
RestNet [22]	25.2	-
PATNet [31]	32.1	40.2
PxMch [3]	34.8	-
ABCFSS [20]	35.1	41.3
<b>MPA (w Source-Training)</b>	<b>55.5</b>	<b>62.0</b>
<b>MPA (w/o Source-Training)</b>	<b>54.2</b>	<b>61.1</b>

duce an additional, more challenging query view. This new view cumulatively incorporates all previously used augmentation operations along with an additional, more complex operation. The mIoU [38] serves as the evaluation metric.

### 5.3. Comparison with State-of-the-art Methods

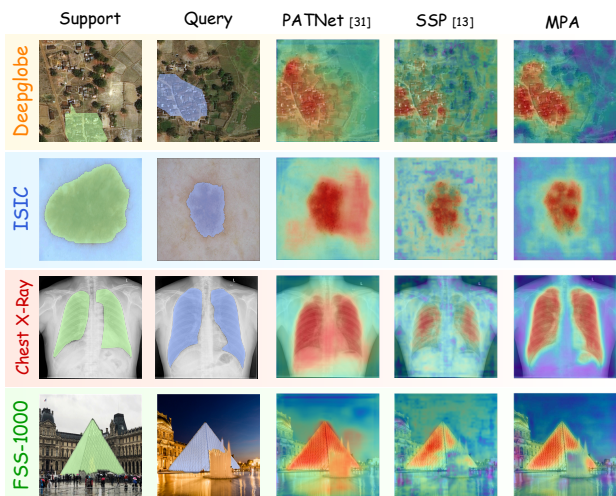
**Effectiveness.** Tabs. 3 and 4 compare MPA with the state-of-the-art methods, showing that MPA consistently achieves superior performance by a substantial margin. In particular, it surpasses IFA [40] by 7.0% (1-shot) and 5.5% (5-shot) in mIoU. For the five target benchmarks, Deepglobe [18] is unique due to its aerial view and complex background. MPA can handle such a complex segmentation task well and improve the mIoU by 3.6% under the 1-shot setup. The medical images in ISIC [7, 55] and Chest X-Ray [2, 24] have a clean background, and their target regions occupy a large portion of the image. MPA also segments such images well. FSS-1000 [33] is more challenging due to substantial variations across images in hundreds of categories. Though

Table 5. Quantitative comparisons between the proposed MPA and SAM-based methods (in mIoU (%)). The best mIoU number under each setup is highlighted by **bold** font.

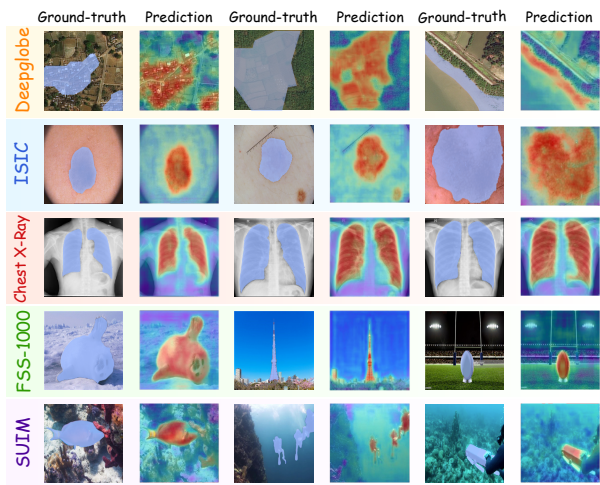
Method	Deepglobe	ISIC	FSS-1000	Avg.
APSeg [19]	35.9	45.4	79.7	53.7
TAVP [67]	46.1	54.9	79.1	60.0
PerSAM [74]	31.4	23.9	71.2	42.2
<b>MPA (w/o Source)</b>	<b>53.1</b>	<b>71.1</b>	<b>80.2</b>	<b>68.1</b>

existing methods benefit from a small domain gap with respect to the source domain (Pascal VOC [11]), MPA still excels under the 1-shot setup. For the very different underwater images in SUIM [23], MPA achieves new state-of-the-art performance as well. The extensive experiments over these different benchmarks demonstrate the superior robustness and generalization capability of the proposed MPA. As shown in Tab. 5, MPA also outperforms SAM-based methods in terms of both effectiveness and model size, highlighting its overall superiority. In addition, we perform qualitative comparisons to verify the superiority of the MPA over multiple CD-FSS benchmarks. As Fig. 3(a) shows, MPA in the Column 5 achieves clearly better segmentation as compared with the baseline SSP [13] in the Column 4. Fig. 3(b) and Figure in the supplementary materials show more visualizations, further demonstrating the great segmentation results of the proposed MPA.

**Discussion.** Previous methods typically rely on source-domain training to establish few-shot capability. However, as discussed earlier, MPA is designed to progressively establish few-shot capability with augmented query views in target domains. To further validate its effectiveness, we conduct experiments without source-training, directly adapting the backbone model to target domains with MPA. Remarkably, this single-stage adaptation achieves perfor-



(a) Comparison with CD-FSS Methods



(b) Qualitative Illustrations of MPA

Figure 3. Qualitative illustrations over five data-scarce domains, including Deepglobe, ISIC, Chest X-Ray, FSS-1000, and SUIM from up to down. Segmentation comparisons with state-of-the-art methods are in (a). For each pair of support image (ground truth highlighted by green) and query image (ground truth highlighted by blue), Column 3-5 show the corresponding segmentation heatmaps by PATNet [31], SSP [13], and our MPA. Segmentation heatmaps of MPA over samples from five domains are shown in (b). Best viewed in color.

Table 6. Ablation studies for technical designs of MPA. Our proposed HPA and DMP designs both show significant effectiveness.

Incorporated design	Deepglobe	ISIC
Baseline	42.1	42.2
+ HPA	47.8	61.2
+ <b>DMP + HPA</b>	<b>53.1</b>	<b>71.1</b>

mance comparable to the conventional two-stage pipeline, attaining 73.4% mIoU in the 1-shot setting and 75.5% in the 5-shot setting on average. Compared with IFA [40], MPA without source training improves mIoU by 5.6% (1-shot) and 4.1% (5-shot), while delivering substantial average gains of 12.7% (1-shot) and 10.5% (5-shot) over the source-free ABCDFSS [20]. These results confirm that most performance gains are derived from the adaptation stage, validating that designing suitable adaptation strategies with limited accessible data is the key to CD-FSS. In addition, removing the source-training stage reduces total training time by approximately 80%, substantially improving efficiency. Consequently, all subsequent experiments are reported under the source-free setting by default, with a detailed discussion of the efficiency benefits in Sec. 5.5.

#### 5.4. Ablation Studies

**Technical designs ablation.** We conduct ablation experiments over our designed Hybrid Progressive Augmentation (HPA) and Dual-chain Multi-view Prediction (DMP) techniques. We use SSP [13] with the pre-trained ResNet-50 [18] as the baseline. First, we incorporate HPA (Sec. 4.1) by progressively increasing the complexity of augmented views, while limiting the number of views during training. As shown in Tab. 6, this yields performance gains of 5.7%




Figure 4. Ablation study on different adaptation strategies.

and 19.0% on the two datasets, respectively. Then, we incorporate HPA and DMP, and observe additional performance gains of 5.3% and 9.9% on the two datasets. Notably, DMP serves as the cornerstone of the MPA framework, effectively building diverse prediction scenarios and achieving larger improvements by exhaustively exploiting the multiple views generated by HPA. These two designs complement each other to progressively establish few-shot capability. Overall, the results underscore the importance of augmenting more views and designing appropriate strategies for adaptation.

**Adaptation strategies.** We benchmark MPA against several widely adopted adaptation strategies from prior studies. For a fair comparison, all strategies are evaluated under both source-trained and direct adaptation setups. Specifically, Task-Adaptive Fine-Tuning (TAFT) [31] works at the meta-testing stage, which leverages Kullback-Leibler divergence loss to minimize the distance between the category prototypes of the segmented query image and the support set. Copy-Paste Fine-Tuning (CPFT) [16] mixes the foreground region of a target image with the background region of a random source image as the support set, forcing source-to-target adaptation as inspired by [17]. IFA [40] establishes support-query correspondences during the fine-

Table 7. Effectiveness of progressive strategies. Both explicit (generating more challenging query views) and implicit (increasing the number of augmented views) progressive augmentation benefit the final performance.

Setups	mIoU
always 1 augmented query	50.5
<b>implicit progressive strategy</b>	<b>52.0</b>
always simple augmentation	51.3
<b>explicit progressive strategy</b>	<b>52.4</b>
<b>combine both strategies</b>	<b>53.1</b>

tuning stage. It augments a single view of the given support image only and is susceptible to overfitting. As Fig. 4 shows, MPA delivers consistently strong performance and remains competitive across both setups. These results highlight that MPA does not depend heavily on source-domain training and can reliably establish few-shot capability in the target domain using only limited accessible data.

**Progressive strategies.** We examine the effectiveness of explicit and implicit progressive strategies as in Tab. 7. The implicit progressive strategy increases the number of augmented queries as performance improves, which inherently accumulates prediction errors and leads to inaccuracy in the subsequent queries [40] (in the sequential chain of DMP). Nevertheless, such inaccuracy facilitates the evaluation of the generalization capability [40], and more discussion about this implicit design is the supplementary materials. As Tab. 7 shows, the implicit progressive improves the mIoU by around 1.4%. The explicit progressive strategy generates more challenging queries progressively via stronger augmentations, *e.g.*, from simple flipping to complex color jittering. As Tab. 7 shows, the explicit progressive strategy brings in around 1.1% improvement in mIoU. More detail of the involved hyper-parameters is discussed in the supplementary materials.

## 5.5. Analysis

**Analysis of DMP.** We perform an in-depth analysis of the DMP design. For the parallel chain, the augmented query views create multiple independent learning routes, effectively increasing the data available for mimicking the “support-to-query” prediction. This leads to clear performance gains by increasing the amount of adaptation data. For the sequential chain, errors can accumulate and propagate across multiple prediction turns. As training progresses, later views undergo stronger augmentations, resulting in increasingly perturbed predictions that deviate from the ground truth (see Tab. 2). To tackle these errors, we apply dense supervision to every sequentially augmented prediction, encouraging the model to learn representations robust to a broad spectrum of perturbations. This design improves the model’s generalization potential. Qualitative results in the supplementary material further support the motivation of the sequential prediction chain.

Table 8. The mIoU (%) under fine-tuning with different augmentation strategies. Our proposed cumulative augmentation strategy achieves the best performance.

Augmentation	Deepglobe	ISIC
Simple augmentation	51.3	67.9
Augmentation replacement	51.9	68.5
<b>Ours</b>	<b>53.1</b>	<b>71.1</b>

Table 9. MPA significantly reduces the training time (min) and improves the mIoU performance (%).

<i>Adaptation with Source-Training</i>				
Adaptation strategy	Deepglobe		ISIC	
	time	mIoU	time	mIoU
PATNet [31]	538	37.9	532	41.2
IFA [40]	555	50.6	551	66.3
<i>Adaptation without Source-Training</i>				
<b>MPA</b>	<b>98</b>	<b>53.1</b>	<b>95</b>	<b>71.1</b>

**Analysis of HPA.** We conduct adaptation under three setups: (1) using only simple augmentation, (2) replacing simple augmentation with a single, more complex augmentation per view, and (3) our augmentation accumulation strategy. The results in Tab. 8 show that our adopted augmentation accumulation strategy achieves the best results.

**Efficiency.** We benchmark the training time of MPA (w/o source-training) with several representative methods. As Tab. 9 shows, MPA saves over 80% time as compared with the state-of-the-art training-required method IFA but achieves clearly better segmentation results<sup>3</sup>. Additionally, data augmentation in MPA is performed only once at the beginning, which introduces negligible overhead. Further, feature extraction of all augmented views in MPA runs in parallel without incurring extra computational cost, and all experiments can be performed on GTX 2080Ti GPUs.

## 6. Conclusion

In this paper, we tackle the challenges of Cross-Domain Few-Shot Segmentation, where large domain gaps and scarce target-domain data hinder the establishment of few-shot capability. We proposed Multi-View Progressive Adaptation (MPA), which gradually builds few-shot capability from both the data and strategy perspectives. By introducing Hybrid Progressive Augmentation to gradually generate more complex and diverse views, and Dual-chain Multi-view Prediction to exhaustively exploit them, MPA effectively establish few-shot capability in target domains. Extensive experiments demonstrate its superiority across multiple benchmarks. In future work, we plan to extend the progressive adaptation framework to broader cross-domain tasks and explore its potential in real-world applications with more complex domain shifts.

<sup>3</sup>MPA performs adaptation with one image under 1-shot setting, after which it is evaluated on the entire test set.

## Acknowledgement

This study is supported by the Interdisciplinary Graduate Programme, Nanyang Technological University, and the Ministry of Education Singapore, under the Tier-2 project scheme with a project number MOE-T2EP20123-0003.

## References

- [1] Malik Boudiaf, Hoel Kervadec, Ziko Imtiaz Masud, Pablo Piantanida, Ismail Ben Ayed, and Jose Dolz. Few-shot segmentation without meta-learning: A good transductive inference is all you need? In *CVPR*, 2021. 2, 6
- [2] Sema Candemir, Stefan Jaeger, Kannappan Palaniappan, Jonathan P Musco, Rahul K Singh, Zhiyun Xue, Alexandros Karargyris, Sameer Antani, George Thoma, and Clement J McDonald. Lung segmentation in chest radiographs using anatomical atlases with nonrigid registration. *TMI*, 2013. 1, 5, 6
- [3] Hao Chen, Yonghan Dong, Zheming Lu, Yunlong Yu, and Jungong Han. Pixel matching network for cross-domain few-shot segmentation. In *WACV*, 2024. 2, 6
- [4] Jiayi Chen, Rong Quan, and Jie Qin. Cross-domain few-shot semantic segmentation via doubly matching transformation. In *IJCAI*, 2024. 2
- [5] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *ICLR*, 2019. 1
- [6] Gong Cheng, Chunbo Lang, and Junwei Han. Holistic prototype activation for few-shot segmentation. *TPAMI*, 2022. 2
- [7] Noel Codella, Veronica Rotemberg, Philipp Tschandl, M Emre Celebi, Stephen Dusza, David Gutman, Brian Helba, Aadi Kalloo, Konstantinos Liopyris, Michael Marchetti, et al. Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic). *arXiv preprint arXiv:1902.03368*, 2019. 1, 5, 6
- [8] Ilke Demir, Krzysztof Koperski, David Lindenbaum, Guan Pang, Jing Huang, Saikat Basu, Forest Hughes, Devis Tuia, and Ramesh Raskar. Deepglobe 2018: A challenge to parse the earth through satellite images. In *CVPRW*, 2018. 1, 5
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 5
- [10] Nanqing Dong and Eric P Xing. Few-shot semantic segmentation with prototype learning. In *BMVC*, 2018. 1, 2
- [11] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010. 2, 6, 1
- [12] Haoran Fan, Qi Fan, Maurice Pagnucco, and Yang Song. Darnet: Bridging domain gaps in cross-domain few-shot segmentation with dynamic adaptation. *arXiv preprint arXiv:2312.04813*, 2023. 2
- [13] Qi Fan, Wenjie Pei, Yu-Wing Tai, and Chi-Keung Tang. Self-support few-shot semantic segmentation. In *ECCV*, 2022. 1, 2, 4, 5, 6, 7
- [14] Qi Fan, Kaiqi Liu, Nian Liu, Hisham Cholakkal, Rao Muhammad Anwer, Wenbin Li, and Yang Gao. Adapting in-domain few-shot segmentation to new domains without retraining. *arXiv preprint arXiv:2504.21414*, 2025. 2, 4
- [15] Yuqian Fu, Yu Wang, Yixuan Pan, Lian Huai, Xingyu Qiu, Zeyu Shangguan, Tong Liu, Yanwei Fu, Luc Van Gool, and Xingqun Jiang. Cross-domain few-shot object detection via enhanced open-set object detector. In *ECCV*, 2024. 2
- [16] Yipeng Gao, Lingxiao Yang, Yunmu Huang, Song Xie, Shiyong Li, and Wei-Shi Zheng. Acrofofod: An adaptive method for cross-domain few-shot object detection. In *ECCV*, 2022. 7
- [17] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *CVPR*, 2021. 7
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5, 6, 7
- [19] Weizhao He, Yang Zhang, Wei Zhuo, Linlin Shen, Jiaqi Yang, Songhe Deng, and Liang Sun. Apseg: Auto-prompt network for cross-domain few-shot semantic segmentation. In *CVPR*, 2024. 1, 3, 6
- [20] Jonas Herzog. Adapt before comparison: A new perspective on cross-domain few-shot segmentation. In *CVPR*, 2024. 3, 5, 6, 7
- [21] Ruibing Hou, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Cross attention network for few-shot classification. In *NeurIPS*, 2019. 1
- [22] Xinyang Huang, Chuang Zhu, and Wenkai Chen. Restnet: Boosting cross-domain few-shot segmentation with residual transformation network. In *BMVC*, 2023. 2, 6
- [23] Md Jahidul Islam, Chelsey Edge, Yuyang Xiao, Peigen Luo, Muntaqim Mehtaz, Christopher Morse, Sadman Sakib Enan, and Junaed Sattar. Semantic segmentation of underwater imagery: Dataset and benchmark. In *IROS*, 2020. 5, 6
- [24] Stefan Jaeger, Alexandros Karargyris, Sema Candemir, Les Folio, Jenifer Siegelman, Fiona Callaghan, Zhiyun Xue, Kannappan Palaniappan, Rahul K Singh, Sameer Antani, et al. Automatic tuberculosis screening using chest radiographs. *TMI*, 2013. 5, 6
- [25] Bingyi Kang, Zhuang Liu, Xin Wang, Fisher Yu, Jiashi Feng, and Trevor Darrell. Few-shot object detection via feature reweighting. In *ICCV*, 2019. 1
- [26] Dahyun Kang, Heeseung Kwon, Juhong Min, and Minsu Cho. Relational embedding for few-shot classification. In *ICCV*, 2021. 1
- [27] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *ICCV*, 2023. 3
- [28] Kurt Koffka. *Principles of Gestalt psychology*. 2013. 1, 2
- [29] Chunbo Lang, Gong Cheng, Binfei Tu, and Junwei Han. Learning what not to segment: A new perspective on few-shot segmentation. In *CVPR*, 2022. 2

- [30] Chunbo Lang, Gong Cheng, Binfei Tu, Chao Li, and Junwei Han. Base and meta: A new perspective on few-shot segmentation. *TPAMI*, 2023. 2
- [31] Shuo Lei, Xuchao Zhang, Jianfeng He, Fanglan Chen, Bowen Du, and Chang-Tien Lu. Cross-domain few-shot semantic segmentation. In *ECCV*, 2022. 1, 2, 3, 5, 6, 7, 8
- [32] Gen Li, Varun Jampani, Laura Sevilla-Lara, Deqing Sun, Jonghyun Kim, and Joongkyu Kim. Adaptive prototype learning and allocation for few-shot segmentation. In *CVPR*, 2021. 2
- [33] Xiang Li, Tianhan Wei, Yau Pun Chen, Yu-Wing Tai, and Chi-Keung Tang. Fss-1000: A 1000-class dataset for few-shot segmentation. In *CVPR*, 2020. 5, 6
- [34] Qing Lian, Fengmao Lv, Lixin Duan, and Boqing Gong. Constructing self-motivated pyramid curriculums for cross-domain semantic segmentation: A non-adversarial approach. In *ICCV*, 2019. 2
- [35] Yuanwei Liu, Nian Liu, Xiwen Yao, and Junwei Han. Intermediate prototype mining transformer for few-shot semantic segmentation. In *NeurIPS*, 2022. 2
- [36] Zhihe Lu, Sen He, Xiatian Zhu, Li Zhang, Yi-Zhe Song, and Tao Xiang. Simpler is better: Few-shot semantic segmentation with classifier weight transformer. In *ICCV*, 2021. 2
- [37] Xiaoliu Luo, Zhuotao Tian, Taiping Zhang, Bei Yu, Yuan Yan Tang, and Jiaya Jia. Pfenet++: Boosting few-shot semantic segmentation with the noise-filtered context-aware prior mask. *TPAMI*, 2023. 2
- [38] Juhong Min, Dahyun Kang, and Minsu Cho. Hypercorrelation squeeze for few-shot segmentation. In *ICCV*, 2021. 1, 2, 6
- [39] Seonghyeon Moon, Samuel S Sohn, Honglu Zhou, Sejong Yoon, Vladimir Pavlovic, Muhammad Haris Khan, and Mubbasir Kapadia. Msi: Maximize support-set information for few-shot segmentation. In *ICCV*, 2023. 2
- [40] Jiahao Nie, Yun Xing, Gongjie Zhang, Pei Yan, Aoran Xiao, Yap-Peng Tan, Alex C Kot, and Shijian Lu. Cross-domain few-shot segmentation via iterative support-query correspondence mining. In *CVPR*, 2024. 1, 2, 4, 5, 6, 7, 8, 3
- [41] Jiahao Nie, Yun Xing, Wenbin An, Qingsong Zhao, Jiawei Shao, Yap-Peng Tan, Alex C Kot, Shijian Lu, and Xuelong Li. Boosting sam for cross-domain few-shot segmentation via conditional point sparsification. *arXiv preprint arXiv:2602.05218*, 2026. 3
- [42] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 5
- [43] Bohao Peng, Zhuotao Tian, Xiaoyang Wu, Chengyao Wang, Shu Liu, Jingyong Su, and Jiaya Jia. Hierarchical dense correlation distillation for few-shot segmentation. In *CVPR*, 2023. 1, 2, 6
- [44] Shi-Feng Peng, Guolei Sun, Yong Li, Hongsong Wang, and Guo-Sen Xie. Sam-aware graph prompt reasoning network for cross-domain few-shot segmentation. In *AAAI*, 2025. 3
- [45] Ri-Zhao Qiu, Yu-Xiong Wang, and Kris Hauser. Aligndiff: aligning diffusion models for general few-shot segmentation. In *ECCV*, 2024. 2
- [46] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Guided curriculum model adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. In *ICCV*, 2019. 2
- [47] Sarinda Samarasinghe, Mamshad Nayeem Rizve, Navid Kardan, and Mubarak Shah. Cdfls-v: Cross-domain few-shot learning for videos. In *ICCV*, 2023. 3
- [48] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots. One-shot learning for semantic segmentation. In *BMVC*, 2017. 2
- [49] Mennatullah Siam, Boris N Oreshkin, and Martin Jagersand. Amp: Adaptive masked proxies for few-shot segmentation. In *ICCV*, 2019. 6
- [50] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, 2017. 1, 2
- [51] Jiapeng Su, Qi Fan, Wenjie Pei, Guangming Lu, and Fanglin Chen. Domain-rectifying adapter for cross-domain few-shot segmentation. In *CVPR*, 2024. 1, 2, 6
- [52] Zhuotao Tian, Hengshuang Zhao, Michelle Shu, Zhicheng Yang, Ruiyu Li, and Jiaya Jia. Prior guided feature enrichment network for few-shot segmentation. *TPAMI*, 2020. 1, 2, 6
- [53] Jintao Tong, Yixiong Zou, Yuhua Li, and Ruixuan Li. Lightweight frequency masker for cross-domain few-shot semantic segmentation. In *NeurIPS*, 2024. 1, 2
- [54] Jintao Tong, Yixiong Zou, Guangyao Chen, Yuhua Li, and Ruixuan Li. Self-disentanglement and re-composition for cross-domain few-shot segmentation. *arXiv preprint arXiv:2506.02677*, 2025. 2
- [55] Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific Data*, 2018. 5, 6
- [56] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *ICCV*, 2019. 2, 6
- [57] Wenjian Wang, Lijuan Duan, Yuxi Wang, Qing En, Junsong Fan, and Zhaoxiang Zhang. Remember the difference: Cross-domain few-shot semantic segmentation via meta-memory transfer. In *CVPR*, 2022. 2, 6
- [58] Xin Wang, Yudong Chen, and Wenwu Zhu. A survey on curriculum learning. *TPAMI*, 2021. 2, 3, 4
- [59] Yuan Wang, Rui Sun, Zhe Zhang, and Tianzhu Zhang. Adaptive agent transformer for few-shot segmentation. In *ECCV*, 2022. 2
- [60] Jiamin Wu, Xin Liu, Xiaotian Yin, Tianzhu Zhang, and Yongdong Zhang. Task-adaptive prompted transformer for cross-domain few-shot learning. In *AAAI*, 2024. 2
- [61] Aoran Xiao, Weihao Xuan, Heli Qi, Yun Xing, Ruijie Ren, Xiaoqin Zhang, Ling Shao, and Shijian Lu. Cat-sam: Conditional tuning for few-shot adaptation of segment anything model. In *ECCV*, 2024. 1
- [62] Yang Xiao, Vincent Lepetit, and Renaud Marlet. Few-shot object detection and viewpoint estimation for objects in the wild. *TPAMI*, 2022. 1
- [63] Qianxiong Xu, Wenting Zhao, Guosheng Lin, and Cheng Long. Self-calibrated cross attention network for few-shot segmentation. In *ICCV*, 2023. 2

- [64] Qianxiong Xu, Guosheng Lin, Chen Change Loy, Cheng Long, Ziyue Li, and Rui Zhao. Eliminating feature ambiguity for few-shot segmentation. In *ECCV*, 2024.
- [65] Qianxiong Xu, Xuanyi Liu, Lanyun Zhu, Guosheng Lin, Cheng Long, Ziyue Li, and Rui Zhao. Hybrid mamba for few-shot segmentation. In *NeurIPS*, 2024. 2
- [66] Boyu Yang, Chang Liu, Bohao Li, Jianbin Jiao, and Qixiang Ye. Prototype mixture models for few-shot semantic segmentation. In *ECCV*, 2020. 6
- [67] Jiaqi Yang, Ye Huang, Xiangjian He, Linlin Shen, and Guoping Qiu. Tavp: Task-adaptive visual prompt for cross-domain few-shot segmentation. *arXiv preprint arXiv:2409.05393*, 2024. 3, 6
- [68] Yong Yang, Qiong Chen, Yuan Feng, and Tianlin Huang. Mianet: Aggregating unbiased instance and general information for few-shot semantic segmentation. In *CVPR*, 2023. 2
- [69] Bingfeng Zhang, Jimin Xiao, and Terry Qin. Self-guided and cross-guided learning for few-shot segmentation. In *CVPR*, 2021. 6
- [70] Chi Zhang, Guosheng Lin, Fayao Liu, Jiushuang Guo, Qingyao Wu, and Rui Yao. Pyramid graph networks with connection attentions for region-based one-shot semantic segmentation. In *ICCV*, 2019. 6
- [71] Chi Zhang, Guosheng Lin, Fayao Liu, Rui Yao, and Chunhua Shen. Canet: Class-agnostic segmentation networks with iterative refinement and attentive few-shot learning. In *CVPR*, 2019. 6
- [72] Gengwei Zhang, Guoliang Kang, Yi Yang, and Yunchao Wei. Few-shot segmentation via cycle-consistent transformer. In *NeurIPS*, 2021. 2, 4, 1
- [73] Gongjie Zhang, Zhipeng Luo, Kaiwen Cui, Shijian Lu, and Eric P Xing. Meta-detr: Image-level few-shot detection with inter-class correlation exploitation. *TPAMI*, 2022. 1
- [74] Renrui Zhang, Zhengkai Jiang, Ziyu Guo, Shilin Yan, Junting Pan, Xianzheng Ma, Hao Dong, Peng Gao, and Hongsheng Li. Personalize segment anything model with one shot. In *ICLR*, 2023. 3, 6
- [75] Yang Zhang, Philip David, Hassan Foroosh, and Boqing Gong. A curriculum domain adaptation approach to the semantic segmentation of urban scenes. *TPAMI*, 2019. 2
- [76] Lanyun Zhu, Tianrun Chen, Jianxiong Yin, Simon See, and Jun Liu. Addressing background context bias in few-shot segmentation through iterative modulation. In *CVPR*, 2024. 2