

Cross-Domain Few-Shot Segmentation via Iterative Support-Query Correspondence Mining

Jiahao Nie^{1,2*} Yun Xing^{2*} Gongjie Zhang³ Pei Yan^{4,2} Aoran Xiao²
Yap-Peng Tan² Alex C. Kot² Shijian Lu^{2†}

¹Interdisciplinary Graduate Programme, Nanyang Technological University

²Nanyang Technological University ³Black Sesame Technologies

⁴Huazhong University of Science and Technology

{jiahao007,xing0047}@e.ntu.edu.sg gjz@ieee.org yanpei@hust.edu.cn

{aoran.xiao,eyptan,eackot,shijian.lu}@ntu.edu.sg

Abstract

Cross-Domain Few-Shot Segmentation (CD-FSS) poses the challenge of segmenting novel categories from a distinct domain using only limited exemplars. In this paper, we undertake a comprehensive study of CD-FSS and uncover two crucial insights: (i) the necessity of a fine-tuning stage to effectively transfer the learned meta-knowledge across domains, and (ii) the overfitting risk during the naïve fine-tuning due to the scarcity of novel category examples. With these insights, we propose a novel cross-domain fine-tuning strategy that addresses the challenging CD-FSS tasks. We first design Bi-directional Few-shot Prediction (BFP), which establishes support-query correspondence in a bi-directional manner, crafting augmented supervision to reduce the overfitting risk. Then we further extend BFP into Iterative Few-shot Adaptor (IFA), which is a recursive framework to capture the support-query correspondence iteratively, targeting maximal exploitation of supervisory signals from the sparse novel category samples. Extensive empirical evaluations show that our method significantly outperforms the state-of-the-arts (+7.8%), which verifies that IFA tackles the cross-domain challenges and mitigates the overfitting simultaneously. The code is available at: <https://github.com/niejiahao1998/IFA>.

1. Introduction

Few-Shot Segmentation (FSS) aims to segment novel categories based on very limited support exemplars, typically by transferring category-agnostic knowledge learned from abundant base categories to novel categories [10, 13, 40,




Figure 1. We investigate two types of category correspondence, **left:** *Support-to-Query (S2Q)* and **right:** *Support-to-Query-to-Support (S2Q2S)* under four experimental setups (a-d). (a) In-domain performances without fine-tuning (FT) set oracle baselines for Cross-Domain Few-Shot Segmentation (CD-FSS). (b) Cross-domain results without fine-tuning suffer from severe performance drops, which verifies the necessity of bridging domain gap for CD-FSS. (c) Cross-domain setups with naïve fine-tuning only bring small performance gains, which is attributed to the overfitting risk of CD-FSS fine-tuning. Notably, there also underlies rich unexplored category correspondence in *S2Q2S*. (d) Cross-domain setup with our proposed Iterative Few-Shot Adaptor (IFA) achieve significant performance gains. IFA comprehensively exploits maximum information content in the given data by capturing both *S2Q* and *S2Q2S* category correspondence during fine-tuning.

43, 46, 47]. Similar to other few-shot tasks [6, 24, 27, 28, 57, 63, 65], FSS is generally addressed with meta-learning [1, 6, 37, 50, 54] and learns generalizable category correspondence from constructed support-query pairs.

Current state-of-the-art FSS approaches [13, 43] have demonstrated impressive performance when novel categories fall in the same domain with base categories, yet still suffer from severe performance drop when a clear domain gap is present [33] (*e.g.*, base categories from Pascal VOC [11], and novel categories from Deepglobe [8]). To

*Equal contribution

†Corresponding author

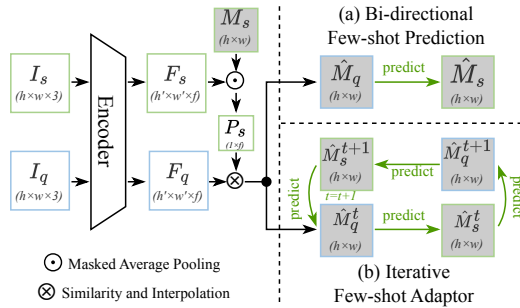


Figure 2. Illustration of our designs for CD-FSS: (a) Bi-directional Few-shot Prediction, and (b) Iterative Few-shot Adaptor. I_s and I_q denote support and query images respectively. F_s and F_q denote the corresponding support and query features as extracted by the *Encoder*. M_s denotes the support mask, and P_s denotes the generated support prototype.

bridge this gap, the task of Cross-Domain Few-Shot Segmentation (CD-FSS) has been explored for generalizing the ability of FSS beyond its own domain [33, 54]. Nevertheless, directly applying a meta-learned model to target domain is incapable of mitigating this challenge [33], largely because the category correspondence learned under meta-learning still biases toward the source domain.

We conduct a series of in-domain and cross-domain FSS experiments to verify our hypothesis and explore solutions¹. Specifically, we investigate two types of category correspondence, *Support-to-Query (S2Q)* and *Support-to-Query-to-Support (S2Q2S)*, to measure the generalization capability under different setups. Concretely, *S2Q* uses support exemplars to segment query images, while *S2Q2S* uses *S2Q* outcomes to predict back on support images. We introduce category correspondence of *S2Q2S* because it is not directly optimized in meta-learning and has a lower risk of overfitting. As shown in Fig. 1(a), in-domain performances without fine-tuning set oracle performances for cross-domain setups. Notably, cross-domain performances (Fig. 1(b)) are lower than the corresponding in-domain performances (Fig. 1(a)), verifying our hypothesis that the learned category correspondence has a clear bias toward the source domain. To generalize the category correspondence beyond its own domain, it is straightforward to introduce a fine-tuning stage [18, 33, 52]. However, naïve fine-tuning on novel category exemplars (Fig. 1(c)) only brings small performance gains. We attribute this phenomenon to overfitting caused by the limited accessible samples from the target domain, which is a core challenge in CD-FSS.

To address the overfitting challenge, we propose a novel fine-tuning strategy for CD-FSS. We first design Bi-directional Few-shot Prediction (BFP), which captures both *S2Q* and *S2Q2S* category correspondence simultaneously

¹Experiments in Fig. 1 are conducted by meta-learning on the same amount of base data and generalizing to the same novel categories (in-domain: Deepglobe, cross-domain: Pascal VOC to Deepglobe).

during meta-learning (Fig. 2(a)). This design leverages support masks as additional supervision to mitigate the overfitting risk. Building upon BFP, we further design Iterative Few-shot Adaptor (IFA), as illustrated in Fig. 2(b). Specifically, IFA iteratively conducts BFP and constructs supervision signals for predictions in every iteration. Hence, IFA comprehensively exploits the supervision from few-shot target exemplars, thereby mining extensive support-query correspondence during fine-tuning. Extensive experiments on four CD-FSS benchmarks show the effectiveness of our designs, especially in mitigating the overfitting challenge.

Our contributions can be summarized in four major aspects:

- We conduct a comprehensive study on the CD-FSS challenge, verifying the necessity of a fine-tuning stage and the overfitting risk during the naïve fine-tuning.
- We design a Bi-directional Few-shot Prediction (BFP) module to establish support-query correspondence, which leverages extensive supervision signals to mitigate overfitting risk during fine-tuning.
- We extend BFP to Iterative Few-shot Adaptor (IFA) in a recursive framework, fully exploiting the supervision signals from limited samples with iterative support-query correspondence mining.
- Our method tackles the cross-domain and overfitting challenges simultaneously and surpasses the state-of-the-art methods by large margins.

2. Related Work

Few-Shot Segmentation (FSS) [10, 13, 43, 44, 47, 53, 64] performs segmentation for novel categories with only a few annotations, which has been studied extensively. Most existing works can be categorized into two types. Prototype-based methods [10, 31, 32, 34, 36, 46, 47, 53] perform segmentation by similarities between all query features and support prototypes. In contrast to prototype-based methods, affinity-based ones [37, 40, 43, 64] mine dense correspondence between query and support features, which rely on rich contextual information. Although the aforementioned methods are well-established, their robustness under cross-domain setups is under-examined. In Cross-Domain Few-Shot Segmentation (CD-FSS) setups, we highlight that: (i) existing prototype-based methods yield unsatisfactory performance because learned category correspondence is hardly generalized to target domains; (ii) existing affinity-based methods are also unsuitable, as affected by much irrelevant information [55] when performing fine-tuning with limited target data. Different from prior attempts, our design demonstrates distinct superiority in CD-FSS, leveraging an effective iterative fine-tuning strategy.

Domain Adaptive Segmentation (DAS) is a paradigm to mitigate costly annotation and domain gap issues. Existing DA methods mitigate these issues by: (i) employ-

ing a discriminator to alleviate differences between different domains at output-level [5, 16, 48, 51] or feature-level [21, 23, 38, 66]; (ii) re-training a model learned from source domain with pseudo labels derived from target domain predictions [17, 39, 58, 67]; or (iii) generating source-style target data and reducing domain gap at input-level [22, 56]. Recently, FSS in cross-domain setup has been explored [1, 37, 54]. Lei *et al.* [33] first propose the Cross-Domain Few-Shot Segmentation (CD-FSS) framework to transfer the trained models to different low-resource domains. Then Fan *et al.* [12] and Huang *et al.* [25] design approaches from test-time fine-tuning and knowledge-transfer aspects, respectively. Motivated by these pioneer works, Chen *et al.* [4] propose a universal method to solve in-domain and cross-domain FSS tasks together. CD-FSS is different from DAS in two folds: (i) when adapting to the target domain, unlabeled data are abundant in DAS but only limited support data is given in CD-FSS, posing risks to overfitting; (ii) DAS assumes source and target label spaces are the same, while the labels spaces of source and target domains in CD-FSS are disjoint. Distinct from the above efforts, we focus on designing an effective fine-tuning strategy for few-shot challenge with limited accessible data.

3. Problem Formulation

Cross-Domain Few-Shot Segmentation (CD-FSS) transfers meta-learned capability of segmenting novel categories to new target domains with only a few annotated support images. Please refer to supplementary materials for more description of CD-FSS. By definition, the model is trained on the source domain \mathcal{D}_{source} and is evaluated on the target domain \mathcal{D}_{target} . Let $\{\mathcal{X}_s, \mathcal{Y}_s\}$ and $\{\mathcal{X}_t, \mathcal{Y}_t\}$ denote the sets in \mathcal{D}_{source} and \mathcal{D}_{target} respectively, where \mathcal{X} denotes the data distribution and \mathcal{Y} denotes the label space. The data distribution of source and target domains are different, and there is no overlap between the source and target label spaces, *i.e.*, $\mathcal{X}_s \neq \mathcal{X}_t$, $\mathcal{Y}_s \cap \mathcal{Y}_t = \emptyset$.

Following the meta-learning in [33], we adopt the episode training strategy. Specifically, each episode is constructed by a support set $S = \{(I_s^i, M_s^i)\}_{i=1}^K$ and a query set $Q = \{(I_q, M_q)\}$ within the same category, where I_s, I_q denote the support and query images, and M_s, M_q denote their masks. The framework can be divided into three steps: (i) training the model in \mathcal{D}_{source} with both M_s and M_q ; (ii) fine-tuning the trained model to \mathcal{D}_{target} with only M_s ; and (iii) testing the adapted model in \mathcal{D}_{target} .

4. Method

To transfer the capability of segmenting novel categories to target domains, we propose a method that mines support-query correspondence iteratively, as illustrated in Fig. 3. The proposed method consists of two major steps: (i) train-

ing models with Bi-directional Few-shot Prediction (BFP) on the source domain; and (ii) fine-tuning trained models to target domains with Iterative Few-shot Adaptator (IFA). The basic pipeline can be formulated as follows: The input support and query images $\{I_s, I_q\}$ are fed into a weight-shared encoder to extract features $\{F_s, F_q\}$. Then the support feature F_s and its mask M_s are processed by the masked average pooling to generate support prototype P_s . Finally, using P_s, F_s , and F_q to predict query masks \hat{M}_q . We elaborate our designs as follows: we first revisit SSP method [13] in Sec. 4.1, which is used directly in our method. Then, the BFP and IFA are presented in Sec. 4.2 and Sec. 4.3, respectively. Finally, we explain how to extend our design into K -shot setting in Sec. 4.4.

4.1. Revisiting of SSP

Motivated by the simple Gestalt principle [30] that pixels belonging to the same object are more similar than those from different objects, thus the given support may not be a good reference for predicting query mask. The target datasets of Cross-Domain Few-Shot Segmentation (CD-FSS) comply with the Gestalt principle, which is verified in supplementary materials. Fan *et al.* propose a Self-Support Prototype (SSP) module to alleviate this problem [13]. Firstly, SSP generates support prototype P_s from support feature F_s and mask M_s :

$$P_s = MAP(F_s, M_s), \quad (1)$$

where MAP is masked average pooling operation. Different from traditional prototypical learning [10] with direct matching between the support prototype and query image, SSP takes a two-step matching. Concretely, SSP uses a support prototype to find the most similar region in the query image first and then takes such region self-matching in query image to predict the mask. Our method (1-shot) is illustrated in Fig. 3, and we adopt SSP module to predict the query prototype \hat{P}_q from support prototype P_s and query feature F_q :

$$\hat{P}_q = SSP(F_q, P_s). \quad (2)$$

Compared with previous methods, SSP yields a more representative prototype for query [13], which is important for our design. More implementation details of applying SSP for our method are shown in supplementary materials.

4.2. Bi-directional Few-shot Prediction

Considering that only a few support annotations are accessible when fine-tuning to target domains, a well-designed strategy for establishing support-query correspondence is essential. The previous uni-directional *Support-to-Query* ($S2Q$) prediction is presented below. Predicted support

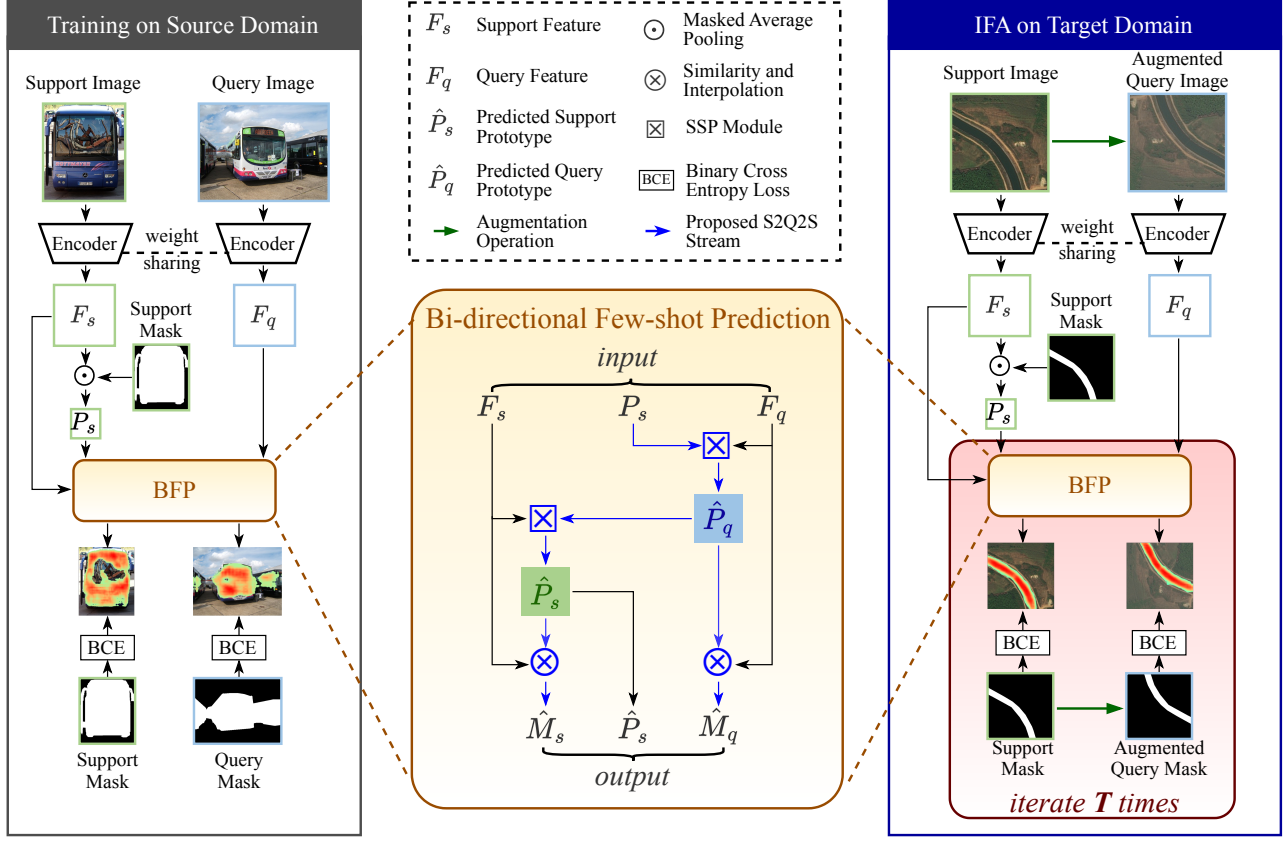


Figure 3. Overall architecture of the proposed Iterative Few-shot Adaptor (IFA), which is composed of two essential steps: training on the source domain, and fine-tuning over the target domain. In the training stage, we only adopt the Bi-directional Few-shot Prediction (BFP) (illustrated in yellow box), which is the fundamental unit of IFA. BFP is composed of both $S2Q$ and $S2Q2S$ streams together with supervision signals from both sides (blue arrows). In the fine-tuning stage where the target exemplars are extremely scarce, IFA is designed to iterate BFP T times, recursively mining the support-query correspondence (illustrated in red box). To show the predictions clearly, we only visualize the region where confidence is higher than 0.5.

mask \hat{M}_s and query mask \hat{M}_q can be obtained by:

$$\begin{aligned}\hat{M}_s &= \text{softmax}(\text{Sim}(F_s, P_s)), \\ \hat{M}_q &= \text{softmax}(\text{Sim}(F_s, \hat{P}_q)),\end{aligned}\quad (3)$$

where Sim is cosine similarity. Then support base loss \mathcal{L}_{bs} and query base loss \mathcal{L}_{bq} are adopted:

$$\begin{aligned}\mathcal{L}_{bs} &= \text{BCE}(\hat{M}_s, M_s), \\ \mathcal{L}_{bq} &= \text{BCE}(\hat{M}_q, M_q),\end{aligned}\quad (4)$$

where BCE is the binary cross entropy loss. Because accessible target exemplars are limited, the naïve meta-learning paradigm easily leads to overfitting. Consequently, it is important to introduce extra information via given data.

Therefore, we propose another *Support-to-Query-Support (S2Q2S)* stream to conduct more predictions:

$$\hat{P}_{s'} = \text{SSP}(F_s, \hat{P}_q),\quad (5)$$

where $\hat{P}_{s'}$ is another support prototype predicted from query. Then, we can introduce corresponding loss $\mathcal{L}_{s'}$ via

support ground-truth:

$$\begin{aligned}\hat{M}_{s'} &= \text{softmax}(\text{Sim}(F_s, \hat{P}_{s'})), \\ \mathcal{L}_{s'} &= \text{BCE}(\hat{M}_{s'}, M_s),\end{aligned}\quad (6)$$

combining \mathcal{L}_b and $\mathcal{L}_{s'}$ helps establishing more robust support-query correspondence.

Our proposed method is supervised by both support and query masks, which are accessible when training models on the source domain. However, query labels are unavailable when fine-tuning to target domains. Consequently, we derive query image I_q with corresponding label M_q from support image I_s and mask M_s :

$$I_q, M_q = \text{AUG}(I_s), \text{AUG}(M_s),\quad (7)$$

where AUG is an augmentation operation, which is described in Sec. 5.2. It is worth noting that we adopt the same transformation operations on I_s and M_s , to ensure that generated I_q and M_q are matched.

4.3. Iterative Few-shot Adaptor

Since I_q and M_q are augmented from I_s and M_s , the region-of-interest should be same in I_s and I_q . If the model is trained well, \hat{M}_q should be accurately predicted from P_s , and $\hat{M}_{s'}$ should also be securely predicted from \hat{P}_q . Otherwise, iteratively conducting predictions may lead to variant results, which measure the generalization capability of learned category correspondence. To further tackle the challenge brought by data scarcity, conducting BFP iteratively on the given data helps the model learn better support-query correspondence. Assuming we totally have T times bi-directional predictions, procedures at the iteration j are:

$$\begin{aligned}\hat{P}_q^{j+1} &= SSP(F_q, \hat{P}_s^j), \\ \hat{P}_s^{j+1} &= SSP(F_s, \hat{P}_q^{j+1}), \\ \hat{M}_q^{j+1} &= softmax(Sim(F_s, \hat{P}_q^{j+1})), \\ \hat{M}_{s'}^{j+1} &= softmax(Sim(F_s, \hat{P}_s^{j+1})).\end{aligned}\quad (8)$$

\hat{P}_q , $\hat{P}_{s'}$, \hat{M}_q and $\hat{M}_{s'}$ in Sec. 4.2 can be treated as \hat{P}_q^1 , \hat{P}_s^1 , \hat{M}_q^1 and $\hat{M}_{s'}^1$ when $j = 0$.

In the meantime, prior predicted errors also accumulate and lead to inaccurate outcomes after T times iterations. To tackle this problem, we introduce supervision signals \mathcal{L}_q^{j+1} and \mathcal{L}_s^{j+1} in every iteration:

$$\begin{aligned}\mathcal{L}_q^{j+1} &= BCE(\hat{M}_q^{j+1}, M_q), \\ \mathcal{L}_s^{j+1} &= BCE(\hat{M}_{s'}^{j+1}, M_{s'}),\end{aligned}\quad (9)$$

\mathcal{L}_{bq} , $\mathcal{L}_{s'}$ in Sec. 4.2 can be treated as \mathcal{L}_q^1 and \mathcal{L}_s^1 when $j = 0$, respectively. It is worth noting that we only conduct iterative predictions in the fine-tuning step to mine category correspondence of target domain. Once the model is learned well, we assume that support prototype \hat{P}_s should be representative, thus we only directly adopt support-to-query prediction in the testing stage.

The total loss \mathcal{L} includes all aforementioned losses with different weights:

$$\begin{aligned}\mathcal{L} &= \lambda_{bs} \times \mathcal{L}_{bs} + \lambda_{bq} \times \mathcal{L}_{bq} + \lambda_{s'} \times \mathcal{L}_{s'} \\ &+ \lambda_i \times \sum_{j=1}^{T-1} (\mathcal{L}_q^{j+1} + \mathcal{L}_s^{j+1}),\end{aligned}\quad (10)$$

where $\lambda_{bs} = 0.2$, $\lambda_{bq} = 1.0$, $\lambda_{s'} = 0.4$, and $\lambda_i = 0.1$. The ablation studies on hyper-parameters are shown in supplementary materials. The overall IFA flow in target domains is presented in Algo. 1.

4.4. Extension to K -shot Setting

In extension to K -shot ($K > 1$) setting, K support images with the corresponding masks $S = \{(I_s^i, M_s^i)\}_{i=1}^K$ are given

Algorithm 1 Pipeline of our proposed Iterative Few-shot Adaptor (1-shot setup).

Require: Support image I_s , Query image I_q , Ground-truth support mask M_s , Ground-truth query mask M_q , Source-trained model \mathcal{G}_s , and Determined iterative times T of predictions.

Ensure: Adapted target model \mathcal{G}_t .

```

1:  $\mathcal{G}_t = \mathcal{G}_s$ 
2: while not reach the maximum epoch do
3:   initialize:
4:      $j = 0, \mathcal{L} = 0$ .
5:   /* Derive fundamental information: */
6:   Extract support feature  $F_s$  via  $\mathcal{G}_t$  and  $I_s$ .
7:   Extract query feature  $F_q$  via  $\mathcal{G}_t$  and  $I_q$ .
8:   Calculate  $P_s$  by Eqn.1.
9:   Predict  $\hat{M}_s$  by Eqn.3.
10:  Calculate  $\mathcal{L}_{bs}$  by Eqn.4 and add to  $\mathcal{L}$ .
11:  /* Start iterative predictions: */
12:   $\hat{P}_s^0 = P_s$ .
13:  while  $j < T$  do
14:    Calculate  $\hat{P}_q^{j+1}$  and  $\hat{P}_s^{j+1}$  by Eqn.8.
15:    Predict  $\hat{M}_q^{j+1}$  and  $\hat{M}_{s'}^{j+1}$  by Eqn.8.
16:    Calculate  $\mathcal{L}_q^{j+1}$  and  $\mathcal{L}_s^{j+1}$  by Eqn.9 and add to  $\mathcal{L}$ .
17:     $j = j + 1$ .
18:  end while
19:  /* Model parameter updating: */
20:  Optimize and update target model  $\mathcal{G}_t$ .
21: end while

```

for fine-tuning. We derive (I_q, M_q) from one randomly picked pair (I_s^i, M_s^i) . IFA can be quickly and easily extended to the new setting as follows.

As indicated in Sec. 4.2, $S2Q$ prediction becomes using averaged $\bar{P}_s = \frac{1}{K} \sum_{i=1}^K P_s^i$ to predict \hat{P}_q . Accordingly, we use \hat{P}_q to predict each $\hat{P}_{s'}^i$, one by one in the $S2Q2S$ prediction procedure. In the iterative design in Sec. 4.3, IFA repeats the above steps T times for K -shot setting.

5. Experiment

5.1. Benchmark

Following the settings in [33], we utilize PASCAL VOC 2012 dataset [11] with SBD augmentation [19] as the source domain for training, which contains 20 daily object categories. Subsequently, we fine-tune the trained model to four target domains, encompassing satellite images, medical screenings, and tiny daily objects. Deepglobe [8] is a satellite image dataset with 7 categories: areas of urban, agriculture, rangeland, forest, water, barren, and unknown. For Cross-Domain Few-Shot Segmentation (CD-FSS), we inherit the same pre-processing in [33] to filter out unknown category and cut images into small pieces. ISIC2018 [7, 49]

is an RGB medical image dataset of skin lesions, containing 3 different types of lesions. To augment the diversity of the target domains, we also incorporate a black-white Chest X-Ray medical screening dataset [2, 26], which is collected for Tuberculosis. FSS-1000 [35] is an everyday object dataset. However, it presents considerable challenges because it contains scarce and tiny objects not present in the source domain. Sample images and their corresponding masks from four target datasets are shown in Fig. 4.




Figure 4. Examples of images and their corresponding ground-truth masks from four target domain datasets, encompassing a diverse range from satellite images and medical screenings to minuscule everyday objects.

5.2. Implementation Details

We adopt the popular ResNet [20] pre-trained on ImageNet [9] as the backbone. Following baseline SSP [13], we discard the last backbone stage and the last ReLU for better generalization. The model is implemented in PyTorch [42]. During training on the source domain, our model is optimized with 0.9 momentum and an initial learning rate of $1e-3$. Consistent with PATNet [33], we resize both support and query images to 400×400 , for reducing the memory consumption and speeding up the learning process. In the fine-tuning step, the learning rate is set as $5e-4$ for Deepglobe, ISIC, and FSS-1000, and $1e-5$ for Chest X-Ray. For each dataset, a total of 40 epochs are optimized, with 20 epochs dedicated to training and the remaining for fine-tuning. We randomly adopt a set of transforms function in PyTorch [42], including horizontal-flip, vertical-flip, rotation by 90 degrees, brightness-variation, and hue-variation, for the augmentation operation in Sec. 4.2.

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

In Tab. 1, we report comparisons of our proposed Iterative Few-shot Adaptor (IFA) with state-of-the-art methods. SSP [13] is initially proposed for FSS, we reproduce it

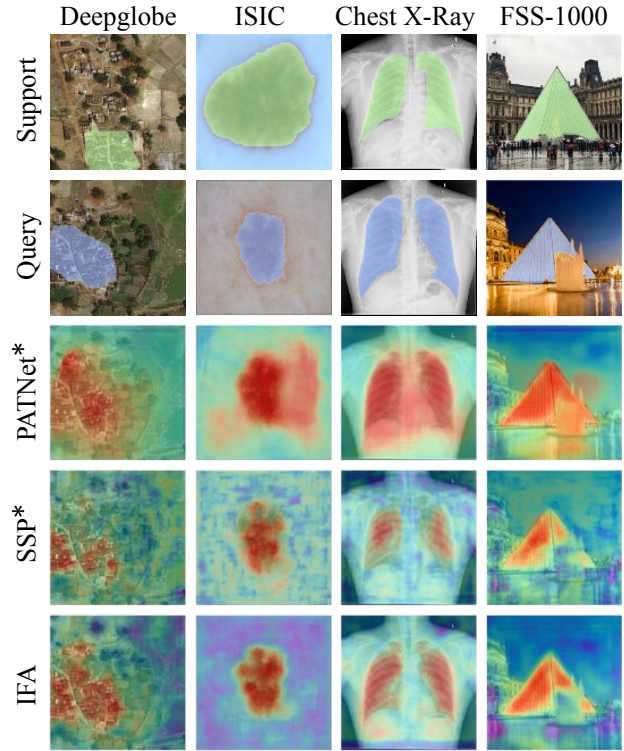


Figure 5. Qualitative results of the samples in four target datasets. From left to right, each column shows examples from Deepglobe, ISIC, Chest X-Ray, and FSS-1000. From up to down, each row shows the examples of support images with ground-truth masks (green), query images with ground-truth masks (blue), PATNet results, Our baseline (SSP) results, and Our IFA results. * represents the model reproduced by ourselves. Best viewed in color.

for CD-FSS. The mIoU (%) is used as the evaluation metric [40]. We construct our ResNet [20] backbone following most compared methods. Our method surpasses state-of-the-arts consistently by large margins. Particularly on the ISIC dataset, our IFA achieves 17.7% (1-shot) and 4.4% (5-shot) of mIoU improvements over the previous best method.

Deepglobe presents unique challenges as every pixel in the image belongs to a well-defined category. Nevertheless, only some specific pixels are region-of-interest in the CD-FSS task, while others are treated as background. The superior performance (+9.3% on 1-shot) of the Deepglobe dataset demonstrates that our IFA excels at complex scene datasets. ISIC and Chest X-Ray are two medical screening datasets, where the background is clean and the target region occupies a large portion of the entire image. The outstanding performances affirm that our IFA effectively establishes the category correspondence of medical domains. The difficulty of FSS-1000 arises from the significant variations between 1000 categories. Despite previous methods performing well due to the relatively small domain gap between the source domain (Pascal VOC) and FSS-1000, our IFA still outperforms priors.

Source Domain: Pascal VOC 2012 → Target Domain: Below											
Methods	Backbone	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
AMP [45]	Vgg-16	37.6	40.6	28.4	30.4	51.2	53.0	57.2	59.2	43.6	45.8
PATNet [33]		28.7	34.8	33.1	45.8	57.8	60.6	71.6	76.2	47.8	54.4
PGNet [61]	Res-50	10.7	12.4	21.9	21.3	34.0	28.0	62.4	62.7	32.2	31.1
PANet [53]		36.6	45.3	25.3	34.0	57.8	69.3	69.2	71.7	47.2	55.1
CaNet [62]		22.3	23.1	25.2	28.2	28.4	28.6	70.7	72.0	36.6	38.0
RPMMs [59]		13.0	13.5	18.0	20.0	30.1	30.8	65.1	67.1	31.6	32.9
PFENet [47]		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 [40]		29.7	35.1	31.2	35.1	51.9	54.4	77.5	81.0	47.6	51.4
PATNet [33]		37.9	43.0	41.2	53.6	66.6	70.2	<u>78.6</u>	<u>81.2</u>	56.1	62.0
SSP* [13]		<u>41.3</u>	<u>54.2</u>	<u>48.6</u>	<u>65.4</u>	<u>72.6</u>	<u>73.0</u>	77.0	79.4	<u>60.0</u>	<u>68.0</u>
IFA $T=3$		50.6	58.8	66.3	69.8	74.0	74.6	80.1	82.4	67.8	71.4

Table 1. Quantitative CD-FSS results under the mIoU (%) evaluation metric. The best and second best results are highlighted with **bold** and underline, respectively. * means the results are reproduced by ourselves.

It is worth noting that our IFA achieves better results on 5-shot setups compared with 1-shot, indicating that our iterative design leverages more information when provided with more support images. To better analyze and understand the superiority of IFA, we visualize the segmentation results, as shown in Fig. 5. Our results (5th row) are significantly better compared to the previous best methods (3rd and 4th row). We also visualize more qualitative results of IFA in supplementary materials.

5.4. Ablation Studies

Component-wise ablation. We conduct ablation studies of our designed components, as shown in Tab. 2. The baseline is our reproduced SSP [13]. To verify our assumption that the *Support-to-Query-to-Support* (S2Q2S) stream mines more support-query correspondence, we first only introduce Bi-directional Few-shot Prediction (BFP) on the source domain and apply the fine-tuning without query supervision on target domains. The performance improvements on Deepglobe(+2.8%) and ISIC(+13.7%) verify the effectiveness of our bi-directional design. After deriving query from annotated support in target domains and repeating BFP three times to formulate our IFA design, we achieve further performance gains of 6.5% and 4.0% on two datasets, respectively. Qualitative results of SSP baseline (4th row) and IFA (5th row) are shown in Fig. 5.

Fine-tuning strategy ablation. We also compare our proposed IFA strategy with other fine-tuning strategies used in relevant tasks, as shown in Tab. 3. All the strategies are performed on the same model (trained via Baseline+BFP)

Pascal VOC 2012 → Below			
Train on source	Adapt to target	Deepglobe	ISIC
Baseline	None	41.3	48.6
Baseline+BFP	w/o query label	44.1	62.3
Baseline+BFP	IFA $T=3$	50.6	66.3

Table 2. Ablation studies for components of our method.

to ensure fairness. Task-Adaptive Fine-Tuning (TAFT) [33] is encapsulated in the meta-testing stage, which is different from our independent fine-tuning design. Specifically, TAFT introduces Kullback-Leibler divergence loss to mitigate the distance between the category prototype of the segmented query image and that of the support set. Copy-Paste Fine-Tuning (CPFT) [14] is employed in few-shot object detection task. Concretely, CPFT mixes the image with the foreground region of the target domain and the background region from a random image in the source domain as the support set. It compels to adapt knowledge learned from source to target domains, which is inspired from [15]. From the experimental results, we find that our IFA is the most effective fine-tuning strategy for CD-FSS.

5.5. Analysis

Effectiveness of iteration design. To demonstrate the effectiveness of iterating BFP multiple times, we compare the results of different times of iteration, as shown in Tab. 4. We

Pascal VOC 2012 → Below		
Fine-tuning strategy	Deepglobe	ISIC
TAFT [33]	43.2	64.3
CPFT [14]	44.6	54.5
IFA $T=3$	50.6	66.3

Table 3. Ablation study on different fine-tuning strategies. We conduct the same training process on the source domain, and adopt different fine-tuning strategies on the target domains.

observe that iterating BFP leads to improved performance, and performing 3 times strikes the best balance between performance and computational cost. We assume that iterating too many times may introduce superfluous information for optimization. Therefore, performing more than 3 iterations yields performance saturation.

Pascal VOC 2012 → Deepglobe									
T	1	2	3	4	5	6	7	8	9
mIoU	46.6	47.2	50.6	50.6	50.8	50.7	50.8	50.9	50.8

Table 4. Quantitative comparison results of different iteration times of BFP. Some visualized results and corresponding analysis are shown in supplementary materials.

Comparison with Segment Anything Model (SAM). We compare the results of our IFA with SAM [29] in two medical image segmentation datasets (*i.e.*, ISIC and Chest X-Ray), as shown in Tab. 5. The results demonstrate substantial improvements from our proposed method over SAM, despite our method relying on a common meta-learning paradigm. It is worth noting that our model employed in this experiment has much fewer parameters and training data than those of the SAM.

Methods	Backbone	ISIC	Chest X-Ray
SAM [29] (<i>zero-shot</i>)	ViT	36.1	27.8
IFA $T=3$ (<i>1-shot</i>)	Res-50	66.3	74.0

Table 5. Comparison with Segment Anything Model (SAM) [29] in two medical image segmentation tasks. The results of IFA are under Res-50 backbone with 1-shot setup. The results of SAM are borrowed from [3].

Results under Domain-Shift Few-Shot Segmentation (DS-FSS). To demonstrate the robustness of our method across different datasets, we also evaluate IFA on Domain-Shift Few-Shot Segmentation (DS-FSS) tasks following the recent work of [1]. Referring to the definition of CD-FSS, DS-FSS can be regarded as a special condition of CD-FSS, where the source domain is COCO-20i [41] while the target domain is Pascal-5i [44] without overlapped categories. DS-FSS is slightly simpler than four tasks in [33], because the

Source Domain: COCO-20i → Target Domain: Pascal-5i			
Methods	Backbone	1-shot	5-shot
		Mean	Mean
RPMs [59]	Res-50	49.6	53.8
PFENet [47]		60.8	61.9
RePRI [1]		57.0	67.9
ASGNet [34]		57.4	66.6
HSNet [40]		61.5	68.4
CWT [37]		59.4	66.5
Meta-Memory [54]		<u>65.6</u>	<u>70.1</u>
IFA $T=3$		71.0	80.9
SCL [60]	Res-101	59.4	60.3
HSNet [40]		63.2	70.0
Meta-Memory [54]		<u>66.6</u>	<u>72.7</u>
IFA $T=3$		79.6	83.4

Table 6. Quantitative Domain-Shift Few-Shot Segmentation results under the mIoU (%) evaluation metric. The best and second best results are highlighted with **bold** and underline, respectively. The concrete performances of each fold are shown in supplementary materials.

categories in the source domain and target domain are both daily objects. Experiment results shown in Tab. A3 also surpass the previous methods by significant margins (+5.4%).

6. Conclusion

This paper presents an in-depth exploration of Cross-Domain Few-Shot Segmentation (CD-FSS) tasks, highlighting the critical role of fine-tuning in transferring FSS capabilities across various domains. We also uncover the challenge of overfitting due to limited data availability. With these insights, we propose a novel cross-domain fine-tuning strategy. We first design Bi-directional Few-shot Prediction (BFP) that establishes extensive category correspondence to mitigate the overfitting risk. Then we extend BFP to Iterative Few-shot Adaptor (IFA), which recursively mines support-query correspondence by maximally exploiting supervisory signals from limited data. Extensive experiments show that our designs tackle both cross-domain and overfitting challenges simultaneously and outperform state-of-the-arts by large margins. We hope that our work inspires further research in developing more effective algorithms and exploring additional facets of few-shot tasks.

Acknowledgement

This study is supported by the Interdisciplinary Graduate Programme, Nanyang Technological University, and the Ministry of Education Singapore, under the Tier-1 scheme with project number RG18/22.

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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13979–13988, 2021. [1](#), [3](#), [7](#), [8](#)
- [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. *IEEE Transactions on Medical Imaging*, 33(2):577–590, 2013. [6](#)
- [3] Hao Chen, Yonghan Dong, Zheming Lu, Yunlong Yu, Yingming Li, Jungong Han, and Zhongfei Zhang. Dense affinity matching for few-shot segmentation. *arXiv preprint arXiv:2307.08434*, 2023. [8](#)
- [4] Hao Chen, Yonghan Dong, Zheming Lu, Yunlong Yu, and Jungong Han. Pixel matching network for cross-domain few-shot segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 978–987, 2024. [3](#)
- [5] Minghao Chen, Hongyang Xue, and Deng Cai. Domain adaptation for semantic segmentation with maximum squares loss. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2090–2099, 2019. [3](#)
- [6] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *International Conference on Learning Representations*, 2019. [1](#)
- [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. [5](#), [1](#)
- [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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 172–181, 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 248–255. Ieee, 2009. [6](#)
- [10] Nanqing Dong and Eric P Xing. Few-shot semantic segmentation with prototype learning. In *Proceedings of the British Machine Vision Conference*, 2018. [1](#), [2](#), [3](#)
- [11] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88:303–338, 2010. [1](#), [5](#)
- [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. [3](#)
- [13] Qi Fan, Wenjie Pei, Yu-Wing Tai, and Chi-Keung Tang. Self-support few-shot semantic segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 701–719. Springer, 2022. [1](#), [2](#), [3](#), [6](#), [7](#)
- [14] Yipeng Gao, Lingxiao Yang, Yunmu Huang, Song Xie, Shiyong Li, and Wei-Shi Zheng. AcrofoD: An adaptive method for cross-domain few-shot object detection. In *Proceedings of the European Conference on Computer Vision*, pages 673–690. Springer, 2022. [7](#), [8](#)
- [15] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2918–2928, 2021. [7](#)
- [16] Dayan Guan, Jiaying Huang, Shijian Lu, and Aoran Xiao. Scale variance minimization for unsupervised domain adaptation in image segmentation. *Pattern Recognition*, 112: 107764, 2021. [3](#)
- [17] Dayan Guan, Jiaying Huang, Aoran Xiao, and Shijian Lu. Domain adaptive video segmentation via temporal consistency regularization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8053–8064, 2021. [3](#)
- [18] Yunhui Guo, Noel C Codella, Leonid Karlinsky, James V Codella, John R Smith, Kate Saenko, Tajana Rosing, and Rogerio Feris. A broader study of cross-domain few-shot learning. In *Proceedings of the European Conference on Computer Vision*, pages 124–141. Springer, 2020. [2](#)
- [19] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 991–998. IEEE, 2011. [5](#)
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. [6](#)
- [21] Judy Hoffman, Dequan Wang, Fisher Yu, and Trevor Darrell. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649*, 2016. [3](#)
- [22] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International Conference on Machine Learning*, pages 1989–1998. Pmlr, 2018. [3](#)
- [23] Weixiang Hong, Zhenzhen Wang, Ming Yang, and Junsong Yuan. Conditional generative adversarial network for structured domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1335–1344, 2018. [3](#)
- [24] Ruibing Hou, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Cross attention network for few-shot clas-

- sification. *Advances in Neural Information Processing Systems*, 32, 2019. **1**
- [25] Xinyang Huang, Chuang Zhu, and Wenkai Chen. Restnet: Boosting cross-domain few-shot segmentation with residual transformation network. *arXiv preprint arXiv:2308.13469*, 2023. **3**
- [26] 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. *IEEE Transactions on Medical Imaging*, 33(2):233–245, 2013. **6**
- [27] Bingyi Kang, Zhuang Liu, Xin Wang, Fisher Yu, Jiashi Feng, and Trevor Darrell. Few-shot object detection via feature reweighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8420–8429, 2019. **1**
- [28] Dahyun Kang, Heeseung Kwon, Juhong Min, and Minsu Cho. Relational embedding for few-shot classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8822–8833, 2021. **1**
- [29] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. **8**
- [30] Kurt Koffka. *Principles of Gestalt psychology*. Routledge, 2013. **3, 2**
- [31] Chunbo Lang, Gong Cheng, Binfei Tu, and Junwei Han. 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*, pages 8057–8067, 2022. **2**
- [32] Chunbo Lang, Gong Cheng, Binfei Tu, Chao Li, and Junwei Han. Base and meta: A new perspective on few-shot segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. **2**
- [33] Shuo Lei, Xuchao Zhang, Jianfeng He, Fanglan Chen, Bowen Du, and Chang-Tien Lu. Cross-domain few-shot semantic segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 73–90. Springer, 2022. **1, 2, 3, 5, 6, 7, 8**
- [34] Gen Li, Varun Jampani, Laura Sevilla-Lara, Deqing Sun, Jonghyun Kim, and Joongkyu Kim. Adaptive prototype learning and allocation for few-shot segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8334–8343, 2021. **2, 8, 3**
- [35] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2869–2878, 2020. **6, 2**
- [36] Yuanwei Liu, Nian Liu, Xiwen Yao, and Junwei Han. Intermediate prototype mining transformer for few-shot semantic segmentation. *Advances in Neural Information Processing Systems*, 35:38020–38031, 2022. **2**
- [37] Zhihe Lu, Sen He, Xi Tian Zhu, Li Zhang, Yi-Zhe Song, and Tao Xiang. Simpler is better: Few-shot semantic segmentation with classifier weight transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8741–8750, 2021. **1, 2, 3, 8**
- [38] Yawei Luo, Ping Liu, Tao Guan, Junqing Yu, and Yi Yang. Significance-aware information bottleneck for domain adaptive semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6778–6787, 2019. **3**
- [39] Zhipeng Luo, Zhongang Cai, Changqing Zhou, Gongjie Zhang, Haiyu Zhao, Shuai Yi, Shijian Lu, Hongsheng Li, Shanghang Zhang, and Ziwei Liu. Unsupervised domain adaptive 3d detection with multi-level consistency. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8866–8875, 2021. **3**
- [40] Juhong Min, Dahyun Kang, and Minsu Cho. Hypercorrelation squeeze for few-shot segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6941–6952, 2021. **1, 2, 6, 7, 8, 3**
- [41] Khoi Nguyen and Sinisa Todorovic. Feature weighting and boosting for few-shot segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 622–631, 2019. **8, 1**
- [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. **6**
- [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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23641–23651, 2023. **1, 2**
- [44] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots. One-shot learning for semantic segmentation. 2017. **2, 8, 1**
- [45] Mennatullah Siam, Boris N Oreshkin, and Martin Jagersand. Amp: Adaptive masked proxies for few-shot segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5249–5258, 2019. **7**
- [46] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in Neural Information Processing Systems*, 30, 2017. **1, 2**
- [47] Zhuotao Tian, Hengshuang Zhao, Michelle Shu, Zhicheng Yang, Ruiyu Li, and Jiaya Jia. Prior guided feature enrichment network for few-shot segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(2):1050–1065, 2020. **1, 2, 7, 8, 3**
- [48] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schuster, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7472–7481, 2018. **3**
- [49] 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*, 5(1):1–9, 2018. **5, 1**
- [50] Hung-Yu Tseng, Hsin-Ying Lee, Jia-Bin Huang, and Ming-Hsuan Yang. Cross-domain few-shot classification via

- learned feature-wise transformation. In *International Conference on Learning Representations*, 2019. [1](#)
- [51] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2517–2526, 2019. [3](#)
- [52] Heng Wang, Tan Yue, Xiang Ye, Zihang He, Bohan Li, and Yong Li. Revisit finetuning strategy for few-shot learning to transfer the embeddings. In *International Conference on Learning Representations*, 2022. [2](#)
- [53] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9197–9206, 2019. [2](#), [7](#)
- [54] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7065–7074, 2022. [1](#), [2](#), [3](#), [8](#)
- [55] Yuan Wang, Rui Sun, Zhe Zhang, and Tianzhu Zhang. Adaptive agent transformer for few-shot segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 36–52. Springer, 2022. [2](#)
- [56] Aoran Xiao, Jiaying Huang, Dayan Guan, Kaiwen Cui, Shijian Lu, and Ling Shao. Polarmix: A general data augmentation technique for lidar point clouds. *Advances in Neural Information Processing Systems*, 35:11035–11048, 2022. [3](#)
- [57] Yang Xiao, Vincent Lepetit, and Renaud Marlet. Few-shot object detection and viewpoint estimation for objects in the wild. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):3090–3106, 2022. [1](#)
- [58] Yun Xing, Dayan Guan, Jiaying Huang, and Shijian Lu. Domain adaptive video segmentation via temporal pseudo supervision. In *Proceedings of the European Conference on Computer Vision*, pages 621–639. Springer, 2022. [3](#)
- [59] Boyu Yang, Chang Liu, Bohao Li, Jianbin Jiao, and Qixiang Ye. Prototype mixture models for few-shot semantic segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 763–778. Springer, 2020. [7](#), [8](#), [3](#)
- [60] Bingfeng Zhang, Jimin Xiao, and Terry Qin. Self-guided and cross-guided learning for few-shot segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8312–8321, 2021. [8](#), [3](#)
- [61] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9587–9595, 2019. [7](#)
- [62] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5217–5226, 2019. [7](#)
- [63] Gongjie Zhang, Kaiwen Cui, Rongliang Wu, Shijian Lu, and Yonghong Tian. Pnpdet: Efficient few-shot detection without forgetting via plug-and-play sub-networks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3823–3832, 2021. [1](#)
- [64] Gengwei Zhang, Guoliang Kang, Yi Yang, and Yunchao Wei. Few-shot segmentation via cycle-consistent transformer. *Advances in Neural Information Processing Systems*, 34:21984–21996, 2021. [2](#), [1](#)
- [65] Gongjie Zhang, Zhipeng Luo, Kaiwen Cui, Shijian Lu, and Eric P Xing. Meta-detr: Image-level few-shot detection with inter-class correlation exploitation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. [1](#)
- [66] Jingyi Zhang, Jiaying Huang, Zhipeng Luo, Gongjie Zhang, Xiaoqin Zhang, and Shijian Lu. Da-detr: Domain adaptive detection transformer with information fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23787–23798, 2023. [3](#)
- [67] Yang Zou, Zhiding Yu, BVK Kumar, and Jinsong Wang. Un-supervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European Conference on Computer Vision*, pages 289–305, 2018. [3](#)