

Bridging Granularity Gaps: Hierarchical Semantic Learning for Cross-domain Few-shot Segmentation

Sujun Sun^{1,2}, Haowen Gu¹, Cheng Xie¹, Yanxu Ren¹, Mingwu Ren^{1,2}, Haofeng Zhang^{1,2,*}

¹School of Computer Science and Engineering, Nanjing University of Science and Technology, China

²State Key Laboratory of Intelligent Manufacturing of Advanced Construction Machinery, China

{egg, guhaowen, city945, 324106010241, renmingwu, zhanghf}@njust.edu.cn

Abstract

Cross-domain Few-shot Segmentation (CD-FSS) aims to segment novel classes from target domains that are not involved in training and have significantly different data distributions from the source domain, using only a few annotated samples, and recent years have witnessed significant progress on this task. However, existing CD-FSS methods primarily focus on style gaps between source and target domains while ignoring segmentation granularity gaps, resulting in insufficient semantic discriminability for novel classes in target domains. Therefore, we propose a Hierarchical Semantic Learning (HSL) framework to tackle this problem. Specifically, we introduce a Dual Style Randomization (DSR) module and a Hierarchical Semantic Mining (HSM) module to learn hierarchical semantic features, thereby enhancing the model’s ability to recognize semantics at varying granularities. DSR simulates target domain data with diverse foreground-background style differences and overall style variations through foreground and global style randomization respectively, while HSM leverages multi-scale superpixels to guide the model to mine intra-class consistency and inter-class distinction at different granularities. Additionally, we also propose a Prototype Confidence-modulated Thresholding (PCMT) module to mitigate segmentation ambiguity when foreground and background are excessively similar. Extensive experiments are conducted on four popular target domain datasets, and the results demonstrate that our method achieves state-of-the-art performance.

Introduction

In recent years, deep learning has achieved remarkable progress in numerous computer vision tasks (Hu et al. 2023; Wang et al. 2024; Bai et al. 2024). However, this success is typically predicated on large-scale annotated datasets (Yu et al. 2020; Liao, Xie, and Geiger 2022; Huang, Zhao, and Huang 2019), whose construction is time-consuming and labor-intensive, particularly for dense prediction tasks like semantic segmentation. To overcome data scarcity, Few-shot Segmentation (FSS) has been proposed and widely studied (Peng et al. 2023; Sun et al. 2024; Xu et al. 2024a), and it learns generalizable category correspondence via meta-learning on base classes with abundant annotated sam-

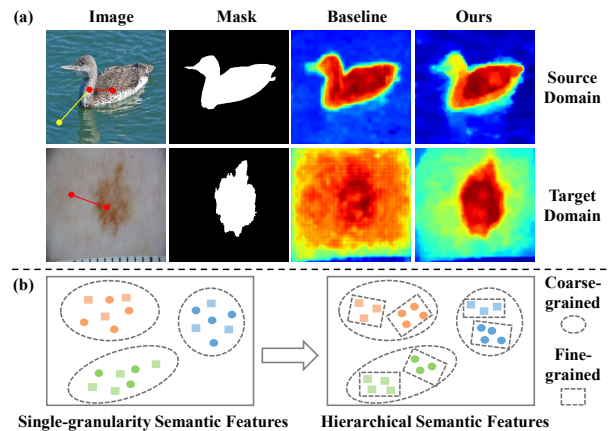


Figure 1: Motivation of the proposed method. (a) Segmentation granularity gaps between the source and target domains, *e.g.*, the foreground-background differences in the target domain are more similar to the finer-grained differences within the foreground of the source domain. The model trained on the source domain only focuses on single-granularity differences between training classes and background, with insufficient ability to distinguish the foreground and background in the target domain. (b) We aim to extract hierarchical semantic features to adapt to target domain data with different segmentation granularities.

ples, thereby enabling the segmentation of novel classes using only a few annotated samples. Nevertheless, existing FSS methods assume base classes and novel classes originate from the same domain. In cross-domain scenarios, these methods suffer significant performance degradation due to domain gaps between base classes and novel classes, severely limiting their practical applications.

To tackle both few-shot and cross-domain challenges simultaneously, Cross-domain Few-shot Segmentation (CD-FSS) (Lei et al. 2022) has recently gained wide attention. Existing CD-FSS methods typically mitigate overfitting to source domain by applying style perturbation to features (Su et al. 2024; Liu et al. 2025), transforming features into a domain-agnostic spaces (Lei et al. 2022; He et al. 2024), or fine-tuning model parameters using limited target domain

*Corresponding Author.

annotated data (Nie et al. 2024; Peng et al. 2025; Herzog 2024). Since fine-tuning with target domain images significantly increases computational costs during testing and easily disrupts cross-view semantic consistency, we focus on methods without target domain fine-tuning. However, we observe that these methods primarily address inter-domain style variations but ignore a core issue: segmentation granularity gaps exist between source and target domains, limiting their performance improvement.

Specifically, as shown in Fig. 1(a), the foreground (bird) and background (other classes) in the source domain image exhibit distinct differences, which we refer to as a coarse-grained semantic partition. In contrast, the distinction between the bird’s neck and wings, or between its brown and gray feathers, is an example of a fine-grained semantic partition. Since the model is trained only on source domain with the aim of reducing intra-class differences and increase inter-class differences for training classes, it tends to extract homogeneous foreground features and treats the entire foreground object as a single entity, ignoring fine-grained differences within the foreground. The foreground (skin lesion area) and background (normal skin area) in the target domain image show less contrast, resembling the fine-grained differences of bird feather colors. Consequently, the model trained on the source domain struggles to effectively capture the finer-grained foreground and background semantic partitions in target domains. To solve this problem, an intuitive idea is to extract hierarchical semantic features that exhibit intra-class consistency and inter-class distinction at different granularities, as shown in Fig. 1(b), thereby ensuring the model better adapt to granularity gaps.

Based on the above analysis, we propose a Hierarchical Semantic Learning (HSL) framework for CD-FSS, which consists of three key modules. First, we propose a Dual Style Randomization (DSR) module that sequentially performs foreground style randomization and global style randomization on training data. Specifically, foreground style randomization modulates the foreground style by using randomly selected local region styles, enhancing the model’s ability to capture varying degrees of inter-class differences without destroying the semantic content of images. Global style randomization modulates the overall image style via random convolution, adapting the model to diverse inter-domain style variations. To further enhance the hierarchical partitioning ability of original features, we propose a Hierarchical Semantic Mining (HSM) module. Since multi-scale superpixel segmentation masks naturally partition an image into local regions approximating semantic regions at different granularities, HSM leverages these masks to generate multi-scale region prototypes for both low-level and high-level features. This process explicitly enhances the inter-class distinction while implicitly guiding the model to mine intra-class consistency across different granularities. Finally, observing that existing similarity comparison-based segmentation methods lead to segmentation ambiguity when foreground and background are excessively similar, we introduce a Prototype Confidence-modulated Thresholding (PCMT) module, which utilizes prototype confidence-weighted adaptive thresholds to segment foreground con-

fidence maps, resulting in more accurate query masks. In summary, our contributions are as follows:

- We are the first to focus on segmentation granularity gaps between source and target domains in CD-FSS, and propose a DSR module and an HSM module to learn hierarchical semantic features, enabling the model to generalize better to target domains with varying segmentation granularities.
- We propose a PCMT module to mitigate segmentation ambiguity when foreground and background are excessively similar, thereby improving segmentation accuracy.
- Extensive experiments on four standard target domain datasets demonstrate that our method significantly outperforms current state-of-the-art CD-FSS methods.

Related Works

Few-shot Segmentation

FSS aims to segment novel classes in query images using only a few annotated images. Existing FSS methods can be broadly categorized into two types: prototype-based methods and matching-based methods. Prototype-based methods typically extract single (Zhang et al. 2019; Fan et al. 2022; Sun et al. 2024) or multiple prototypes (Yang et al. 2020a; Liu et al. 2020; Li et al. 2021) from support images as category guidance and use feature concatenation or distance calculation to segment query images. Due to spatial structure information loss during prototype computation, numerous matching-based methods (Peng et al. 2023; Xu et al. 2023, 2024a) have been proposed to fully exploit information in support images. These methods focus on calculating pixel-to-pixel dense correlations between support and query features, decoding them to obtain query masks. However, when base classes and novel classes originate from source and target domains with significant domain gaps, these FSS methods often fail to generalize effectively to novel classes, resulting in a sharp decline in segmentation performance.

Cross-domain Few-shot Segmentation

To overcome the limitations of FSS methods, many CD-FSS methods have been proposed, aiming to ensure that the model trained solely on the source domain can generalize to target domains with large domain gaps. Some methods (Su et al. 2024; Liu et al. 2025) introduce style perturbations during training to simulate style variations across domains, extracting features with better generalization. Other methods design complex support-query matching mechanisms to better exploit the correlations between support and query images, though they often require additional strategies to mitigate overfitting, such as transforming features extracted by the backbone into domain-agnostic spaces (Lei et al. 2022; He et al. 2024) or fine-tuning with a few annotated samples from target domains (Nie et al. 2024; Tong et al. 2024; Peng et al. 2025; Tong et al. 2025). Some works (Herzog 2024) completely abandon the training stage and directly fine-tune with limited annotated data. In contrast, we focus on extracting hierarchical semantic features that exhibit intra-class consistency and inter-class distinction at different

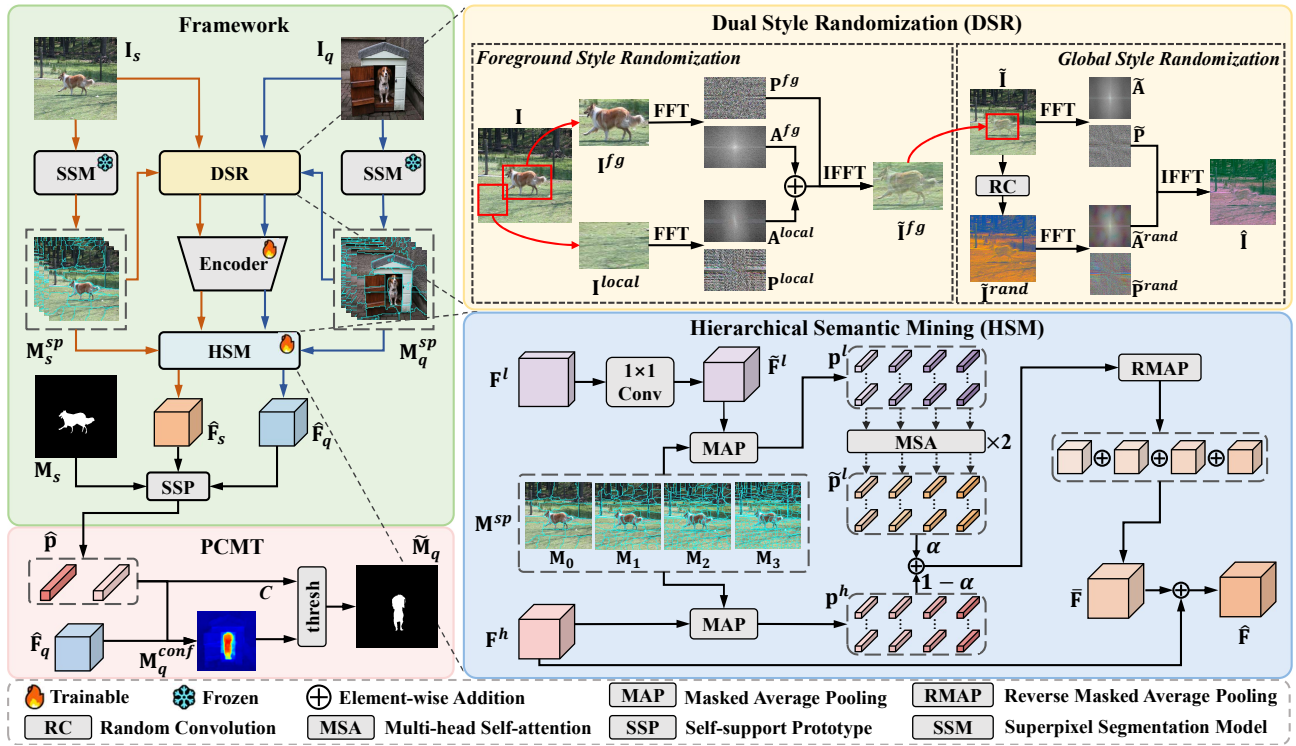


Figure 2: Overview of our method. We first use the SSM to extract multi-scale superpixel masks for both support and query images. These masks are applied to the DSR and HSM to assist in extracting hierarchical semantic features. Subsequently, support and query images are sequentially fed into the DSR, image encoder, and HSM for data augmentation, feature extraction, and feature enhancement, resulting in hierarchical semantic features. Finally, we compute foreground and background prototypes through SSP. These prototypes and query features are fed into the PCMT to perform query image segmentation.

granularities, without the need for complicated fine-tuning processes, while simultaneously adapting to inter-domain style variations and segmentation granularity differences.

Methodology

Problem Definition

In the CD-FSS task, we have two distinct data domains: a source domain $\mathcal{D}_s = \{X_s, Y_s\}$ and a target domain $\mathcal{D}_t = \{X_t, Y_t\}$, where X_s and X_t represent data distributions, and Y_s and Y_t denote corresponding label spaces. The task satisfies $X_s \neq X_t$ and $Y_s \cap Y_t = \emptyset$, indicating no overlap in both data distributions and class labels.

The CD-FSS framework follows a meta-learning paradigm, where only source domain data $\mathcal{D}_{train} \subseteq \mathcal{D}_s$ is used during training, while the model’s generalization ability to novel classes is evaluated on $\mathcal{D}_{test} \subseteq \mathcal{D}_t$ during testing. Each training or testing task is organized as an episode, containing a support set $\mathcal{S} = \{(\mathbf{I}_s^i, \mathbf{M}_s^i)\}_{i=1}^K$ and a query set $\mathcal{Q} = \{(\mathbf{I}_q, \mathbf{M}_q)\}$, where K denotes the number of support samples, \mathbf{I}_s and \mathbf{I}_q represent support and query images, and \mathbf{M}_s and \mathbf{M}_q are their corresponding binary masks. In each episode, the model predicts the query mask based on the support set \mathcal{S} and the query image \mathbf{I}_q .

Method Overview

The architecture of our method is illustrated in Fig. 2, which consists of three key modules to address style gaps and granularity gaps: a Dual Style Randomization (DSR) module, a Hierarchical Semantic Mining (HSM) module, and a Prototype Confidence-modulated Thresholding (PCMT) module.

First, multi-scale superpixel masks for support and query images are obtained through a superpixel segmentation model (Xu et al. 2024b). These masks are applied to the DSR and HSM modules to assist in extracting hierarchical semantic features. Next, DSR sequentially performs foreground and global style randomization on the input images. The randomized images are then fed into an image encoder to extract their corresponding features, which are further enhanced in their hierarchical discrimination ability through HSM. Subsequently, the enhanced features and support mask are input into the SSP module (Fan et al. 2022) to extract foreground and background prototypes. Finally, PCMT utilizes these prototypes and the query feature to compute the foreground confidence map, which is segmented using a prototype confidence-weighted adaptive threshold to generate the final query prediction. We use binary cross entropy (BCE) loss to train our model, as detailed in the appendix. Note that DSR is used only during training, while PCMT is employed exclusively during testing.

Dual Style Randomization

To mitigate performance degradation caused by domain gaps, prior works (Su et al. 2024; Liu et al. 2025) randomize the global style of input images or features to simulate target domain data with diverse styles. Building upon this, we propose a DSR module, which introduces an additional foreground style randomization process to enhance the model’s ability to capture varying degrees of inter-class differences.

Foreground Style Randomization. We aim to simulate target domain data with varying inter-class differences, allowing the model to adapt to segmentation granularity gaps. To achieve this, we use styles from randomly selected local regions to modulate the foreground style without destroying the semantic content of images.

First, we extract the foreground image and the local region image. Given an original image $\mathbf{I} \in \mathbb{R}^{3 \times H \times W}$ and its corresponding foreground mask $\mathbf{M} \in \mathbb{R}^{H \times W}$, where H and W represent the image height and width, we preprocess \mathbf{M} via max pooling and average pooling with kernel size $K \times K$ to obtain the mask \mathbf{M}' . Subsequently, we compute the bounding box of the foreground from \mathbf{M}' and crop the corresponding region from \mathbf{I} , thus obtaining the foreground region image \mathbf{I}^{fg} . Simultaneously, given multi-scale superpixel masks $\mathcal{M}^{sp} = \{\mathbf{M}_i\}_{i=1}^L$ for \mathbf{I} , where L represents the number of superpixel scales, $\mathbf{M}_i \in \mathbb{R}^{H \times W}$ divides the image into n_i superpixel regions $\{R_{i1}, R_{i2}, \dots, R_{in_i}\}$, and $n_i < n_{i+1}$. We randomly select the j -th superpixel region R_{0j} from the coarsest superpixel mask \mathbf{M}_0 . By computing the bounding box corresponding to R_{0j} and cropping the corresponding region from \mathbf{I} , we obtain the randomly selected local region image \mathbf{I}^{local} , which is then resized to the same size as \mathbf{I}^{fg} .

The Fast Fourier Transformation (FFT) can convert images from the spatial domain to the frequency domain and further decompose them into style-representing amplitude spectrum and content-representing phase spectrum. Some works (Chen et al. 2021; Liu et al. 2025) achieve style transformation of images or features by retaining the phase and replacing the amplitude. Here, we compute the random weighted sum of the amplitude of \mathbf{I}^{fg} and \mathbf{I}^{local} , allowing the foreground region’s style to approach or diverge from the local region’s style to varying degrees. Specifically, we transform \mathbf{I}^{fg} and \mathbf{I}^{local} to the frequency domain via FFT and decompose them into phase spectrum \mathbf{P}^{fg} , \mathbf{P}^{local} and amplitude spectrum \mathbf{A}^{fg} , \mathbf{A}^{local} :

$$[\mathbf{A}^{fg}, \mathbf{P}^{fg}] = \text{FFT}(\mathbf{I}^{fg}), \quad (1)$$

$$[\mathbf{A}^{local}, \mathbf{P}^{local}] = \text{FFT}(\mathbf{I}^{local}). \quad (2)$$

Next, we obtain the amplitude spectrum \mathbf{A}^{fusion} representing the randomized foreground style by randomly weighting \mathbf{A}^{fg} and \mathbf{A}^{local} :

$$\mathbf{A}^{fusion} = \omega \mathbf{A}^{local} + (1 - \omega) \mathbf{A}^{fg}, \quad (3)$$

where the weight parameter ω is sampled from a normal distribution $N(0, \sigma_\omega^2)$. We recombine \mathbf{P}^{fg} and \mathbf{A}^{fusion} through the Inverse Fast Fourier Transform (IFFT), generating the style-randomized foreground region image $\tilde{\mathbf{I}}^{fg}$:

$$\tilde{\mathbf{I}}^{fg} = \text{IFFT}(\mathbf{A}^{fusion} e^{i\mathbf{P}^{fg}}), \quad (4)$$

and further pad $\tilde{\mathbf{I}}^{fg}$ to the same size as the original image \mathbf{I} .

Finally, we use $\tilde{\mathbf{I}}^{fg}$ to replace the foreground region of \mathbf{I} :

$$\tilde{\mathbf{I}} = \mathbf{M}'\mathbf{I} + (1 - \mathbf{M}')\tilde{\mathbf{I}}^{fg}, \quad (5)$$

where $\tilde{\mathbf{I}}$ is the foreground style-randomized image. Since the mask \mathbf{M}' has been preprocessed with max pooling and average pooling, it ensures a smooth transition between the foreground and background regions, preventing obvious foreground-background edges in $\tilde{\mathbf{I}}$.

Global Style Randomization. Global style randomization aims to perturb the overall style of each training image, enabling the model to adapt to style variations across domains. Our global style randomization is similar to (Liu et al. 2025), but we apply it to image $\tilde{\mathbf{I}}$ instead of the middle-layer features. We first input $\tilde{\mathbf{I}}$ into a random convolution (RC) layer, generating an image $\tilde{\mathbf{I}}^{rand}$ with random style by perturbing the local textures of $\tilde{\mathbf{I}}$:

$$\tilde{\mathbf{I}}^{rand} = \text{RC}(\tilde{\mathbf{I}}), \quad (6)$$

where the parameters $\Theta \in \mathbb{R}^{3 \times 3 \times 3 \times 3}$ of the convolution kernel are sampled from a normal distribution $N(0, \sigma_\Theta^2)$.

Although random convolution can randomize the overall style of an image, it inevitably destroys some content details. Therefore, instead of directly using $\tilde{\mathbf{I}}^{rand}$ for training, we use the style of $\tilde{\mathbf{I}}^{rand}$ to modulate the style of $\tilde{\mathbf{I}}$. Specifically, following Eq. 1, we convert $\tilde{\mathbf{I}}$ and $\tilde{\mathbf{I}}^{rand}$ to the frequency domain via FFT and decompose them into phase spectrum $\tilde{\mathbf{P}}$, $\tilde{\mathbf{P}}^{rand}$ and amplitude spectrum $\tilde{\mathbf{A}}$, $\tilde{\mathbf{A}}^{rand}$. Then, following Eq. 4, we directly recombine $\tilde{\mathbf{P}}$ and $\tilde{\mathbf{A}}^{rand}$ via IFFT to generate the style-randomized image $\hat{\mathbf{I}}$.

Hierarchical Semantic Mining

In addition to performing data augmentation in DSR, we also process the features to further enhance their hierarchical discrimination ability. We observe that superpixel segmentation (Yang et al. 2020b; Wang et al. 2021; Xu et al. 2024b) can divide an image into multiple local regions, each consisting of a set of adjacent pixels with similar features or semantics. By introducing superpixel masks at multiple scales, we can approximate semantic regions at different granularities, thereby guiding the model to mine intra-class consistency and inter-class distinction at different granularities.

Specifically, given multi-scale superpixel masks $\mathcal{M}^{sp} = \{\mathbf{M}_i\}_{i=1}^L$ corresponding to the original image \mathbf{I} , we first generate binary masks $\{\tilde{\mathbf{M}}_{i1}, \tilde{\mathbf{M}}_{i2}, \dots, \tilde{\mathbf{M}}_{in_i}\}$ for the n_i superpixel regions corresponding to \mathbf{M}_i , where the calculation process for the mask $\tilde{\mathbf{M}}_{ij}$ of the j -th region is formulated as:

$$\tilde{\mathbf{M}}_{ij}(x, y) = \begin{cases} 1, & \mathbf{M}_i(x, y) = j \\ 0, & \text{otherwise} \end{cases}, \quad (7)$$

where (x, y) is the spatial coordinate of the mask. Concurrently, we use an image encoder to extract the shallow low-level feature $\mathbf{F}^l \in \mathbb{R}^{c_l \times h_l \times w_l}$ and deep high-level feature $\mathbf{F}^h \in \mathbb{R}^{c_h \times h_h \times w_h}$, where c_* , h_* and w_* denote the channel depth, height and width of the feature, respectively.

Next, we compute low-level and high-level prototypes corresponding to multi-scale regions and leverage these prototypes to enhance features. High-level region prototypes encode semantics, while low-level ones contain details such as color and texture to enhance the distinction of features. Specifically, we first input \mathbf{F}^l into a 1×1 convolution layer and downsample it to match the dimensions of \mathbf{F}^h :

$$\tilde{\mathbf{F}}^l = \text{Down}(\text{Conv}(\mathbf{F}^l)) \in \mathbb{R}^{c_h \times h_h \times w_h}. \quad (8)$$

Then, we use masked average pooling (MAP) to generate prototypes for each region:

$$\mathbf{p}_{ij}^l = \text{MAP}(\tilde{\mathbf{F}}^l, \tilde{\mathbf{M}}_{ij}), \quad (9)$$

$$\mathbf{p}_{ij}^h = \text{MAP}(\mathbf{F}^h, \tilde{\mathbf{M}}_{ij}), \quad (10)$$

where $\mathbf{p}_{ij}^l, \mathbf{p}_{ij}^h \in \mathbb{R}^{c_h}$ are the low-level and high-level region prototypes corresponding to the j -th region of the i -th scale superpixel mask, respectively.

Since low-level features typically contain noise and have a significant semantic gap with high-level features, we sequentially enhance low-level prototypes at each scale through two multi-head self-attention (MSA) layers:

$$\tilde{\mathcal{P}}_i^l = \text{MSA}_2(\text{MSA}_1(\mathcal{P}_i^l)), \quad (11)$$

where $\mathcal{P}_i^l = \{\mathbf{p}_{ij}^l\}_{j=1}^{n_i}$ and $\tilde{\mathcal{P}}_i^l = \{\tilde{\mathbf{p}}_{ij}^l\}_{j=1}^{n_i}$ are the sets of all low-level prototypes at the i -th scale. Subsequently, we perform weighted fusion of $\tilde{\mathbf{p}}_{ij}^l$ and \mathbf{p}_{ij}^h :

$$\mathbf{p}_{ij} = \alpha \tilde{\mathbf{p}}_{ij}^l + (1 - \alpha) \mathbf{p}_{ij}^h, \quad (12)$$

where α is a weight parameter.

To apply each region prototype to its corresponding pixels, we use the inverse process of MAP, namely RMAP (Peng et al. 2025), to restore region prototypes for each scale into feature maps. This process is formulated as:

$$\bar{\mathbf{F}}_i(x, y) = \sum_j^{n_i} \mathbf{p}_{ij} \tilde{\mathbf{M}}_{ij}(x, y), \quad (13)$$

where $\bar{\mathbf{F}}_i \in \mathbb{R}^{c_h \times h_h \times w_h}$ is the region feature map at scale i . Finally, we use region feature maps at all scales to enhance \mathbf{F}^h , thereby obtaining hierarchical semantic features $\hat{\mathbf{F}}$:

$$\hat{\mathbf{F}} = \mathbf{F}^h + \sum_i^L \bar{\mathbf{F}}_i. \quad (14)$$

Since each pixel’s segmentation is influenced by multi-scale region prototypes, this guides the model to focus on relationships between multi-scale regions across views during training, thereby enhancing the intra-class consistency and inter-class distinction of $\hat{\mathbf{F}}$ at different granularities.

Prototype Confidence-modulated Thresholding

After obtaining the hierarchical semantic features $\hat{\mathbf{F}}_s$ and $\hat{\mathbf{F}}_q$, we compute the support foreground prototype $\hat{\mathbf{p}}_s^{fg}$ and background prototype $\hat{\mathbf{p}}_s^{bg}$ through MAP, further using the SSP module (Fan et al. 2022) to compute the query prototypes $\hat{\mathbf{p}}_q^{fg}$ and $\hat{\mathbf{p}}_q^{bg}$, as well as the fused prototypes $\hat{\mathbf{p}}_q^{fg}$

and $\hat{\mathbf{p}}_q^{bg}$. Subsequently, we upsample the query feature $\hat{\mathbf{F}}_q$ to the size of \mathbf{I}_q , computing the cosine similarity between the fused prototypes and $\hat{\mathbf{F}}_q$ to obtain a foreground similarity map $\mathbf{M}_q^{fg} = \cos(\hat{\mathbf{p}}_q^{fg}, \hat{\mathbf{F}}_q) \in \mathbb{R}^{H \times W}$ and a background similarity map $\mathbf{M}_q^{bg} = \cos(\hat{\mathbf{p}}_q^{bg}, \hat{\mathbf{F}}_q) \in \mathbb{R}^{H \times W}$.

Previous works (Fan et al. 2022; Su et al. 2024; Wang et al. 2019) obtain final query predictions based on similarity comparison, directly classifying pixels as foreground where foreground similarity exceeds background similarity. We found that this approach easily leads to segmentation ambiguity in cross-domain scenarios. Specifically, for images with excessively similar foreground and background, the distinction of their foreground and background features is limited. However, due to cross-view intra-class variations, it is inevitable that both query foreground and background features are more similar to the foreground prototype (or background prototype). In such cases, the similarity comparison-based method results in incorrect segmentation.

Inspired by (Herzog 2024), we introduce the idea of threshold-based segmentation, using prototype confidence-weighted adaptive thresholds to segment the foreground confidence maps. Specifically, we first compute the query foreground confidence map \mathbf{M}_q^{conf} :

$$\mathbf{M}_q^{conf} = \mathbf{M}_q^{fg} - \mathbf{M}_q^{bg}. \quad (15)$$

Next, we compute the adaptive threshold t via OTSU (Otsu et al. 1975) to segment \mathbf{M}_q^{conf} into a binary mask.

Experiments reveal that the adaptive threshold t effectively mitigates segmentation ambiguity, but it introduces additional errors on data where no ambiguity exists. To address this, we introduce prototype confidence to modulate t . The prototype confidence C is calculated as:

$$C = \cos(\hat{\mathbf{p}}_q^{fg}, \hat{\mathbf{p}}_s^{fg}) - \frac{\cos(\hat{\mathbf{p}}_q^{fg}, \hat{\mathbf{p}}_s^{bg}) + \cos(\hat{\mathbf{p}}_s^{fg}, \hat{\mathbf{p}}_q^{bg})}{2}, \quad (16)$$

which approximates the probability of segmentation ambiguity based on cross-view prototype similarity. We then segment \mathbf{M}_q^{conf} using the adaptive threshold weighted by C :

$$\tilde{\mathbf{M}}_q = \mathbf{M}_q^{conf} > \frac{1}{1 + e^{\beta(C+\gamma)}} t, \quad (17)$$

where β and γ are hyperparameters that adjust the weight.

In fact, when the threshold is fixed to 0, Eq. 17 is equivalent to the similarity comparison-based method. For images with higher prototype confidence, we tend to use the threshold 0, whereas for images with lower prototype confidence, we prefer the adaptive threshold t .

Experiments

Datasets and Metric

Following the settings in PATNet (Lei et al. 2022), we train our model on the PASCAL VOC 2012 (Everingham et al. 2010) dataset augmented with SBD (Hariharan et al. 2011), then evaluate the trained model on four target domain datasets: Deepglobe (Demir et al. 2018), ISIC (Codella et al. 2019; Tschandl, Rosendahl, and Kittler 2018), Chest X-ray

Methods	Backbone	Mark	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
Few-shot Semantic Segmentation Methods												
SSP (Fan et al. 2022)	Res-50	ECCV-22	40.00	48.68	35.49	45.86	74.44	74.26	78.91	80.59	57.21	62.35
FPTrans (Zhang et al. 2022)	ViT-base	NIPS-22	38.36	49.30	48.65	60.37	80.92	82.91	80.74	83.65	62.17	69.06
HDMNet (Peng et al. 2023)	Res-50	CVPR-23	25.40	39.10	33.00	35.00	30.60	31.30	75.10	78.60	41.00	46.00
PerSAM (Zhang et al. 2023)	ViT-base	ICLR-24	36.08	40.65	23.27	25.33	29.95	30.05	60.92	66.53	37.56	40.64
VRP-SAM (Sun et al. 2024)	ViT-base	CVPR-24	40.43	44.75	28.21	31.96	30.54	29.99	80.78	83.18	44.99	47.47
Cross-domain Few-shot Semantic Segmentation Methods												
PATNet* (Lei et al. 2022)	Res-50	ECCV-22	37.89	42.97	41.16	53.58	66.61	70.20	78.59	81.23	56.06	61.99
ABCDFSS* (Herzog 2024)	Res-50	CVPR-24	42.60	49.00	45.70	53.30	79.80	81.40	74.60	76.20	60.67	64.97
DRA (Su et al. 2024)	Res-50	CVPR-24	41.29	50.12	40.77	48.87	82.35	82.31	79.05	80.40	60.86	65.42
APSeg (He et al. 2024)	ViT-base	CVPR-24	35.94	39.98	45.43	53.98	84.10	84.50	79.71	81.90	61.30	65.09
APM* (Tong et al. 2024)	Res-50	NIPS-24	40.86	44.92	41.71	51.16	78.25	82.81	79.29	81.83	60.03	65.18
LoEC (Liu et al. 2025)	Res-50	CVPR-25	44.10	49.67	38.21	47.04	81.02	82.73	78.51	80.60	60.46	65.01
LoEC (Liu et al. 2025)	ViT-base	CVPR-25	42.12	51.48	<u>52.91</u>	<u>62.43</u>	83.94	84.12	<u>81.05</u>	<u>83.69</u>	<u>65.01</u>	<u>70.43</u>
SDRC* (Tong et al. 2025)	ViT-base	ICML-25	43.15	46.83	46.57	55.02	82.86	84.79	80.31	82.55	63.22	67.30
HSL(Ours)	Res-50	Ours	46.13	53.80	48.01	55.56	84.57	85.34	78.22	80.36	64.23	68.77
HSL(Ours)	ViT-base	Ours	<u>45.77</u>	54.56	59.36	64.62	85.95	86.25	81.89	83.84	68.24	72.32

Table 1: Mean-IoU of 1-shot and 5-shot results compared with previous FSS and CD-FSS methods. The best and second-best methods are highlighted in **bold** and underlined, respectively. * indicates fine-tuning with target domain data.

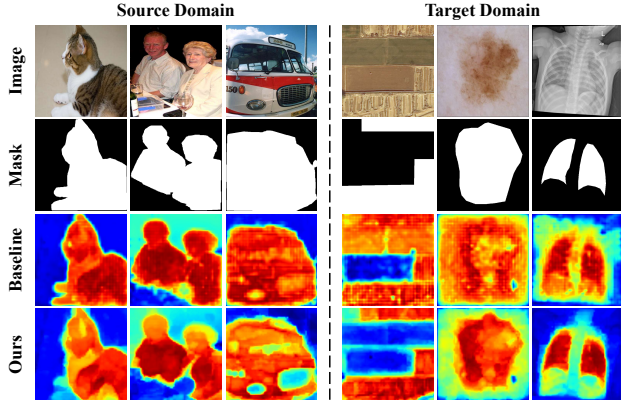


Figure 3: The heatmaps of foreground similarity maps for the source domain and target domains demonstrate that our method can extract hierarchical semantic features.

(Candemir et al. 2013; Jaeger et al. 2013), and FSS-1000 (Li et al. 2020). Please refer to the appendix for more details.

We use mean Intersection over Union (mIoU) as the evaluation metric, and report average results over 5 random seeds. Each run consists of 1200 episodes sampled from Deepglobe, ISIC, and Chest X-ray respectively, and 2400 episodes from FSS-1000.

Implementation Details

Following (Liu et al. 2025), we use ResNet-50 (He et al. 2016) and ViT-B/16 (Dosovitskiy et al. 2020) as our backbones. For the ResNet-50 backbone, we follow SSP (Fan et al. 2022) by discarding the last backbone stage for better generalization, while resizing all images to 400×400 as per the default practice of previous CD-FSS methods and extracting low-level features at the end of stage-0. For the ViT-

DSR	HSM	PCMT	Res-50	ViT-base
×	×	×	57.82	62.24
✓	×	×	60.44	64.55
×	✓	×	60.92	65.29
✓	✓	×	62.97	67.05
✓	✓	✓	64.23	68.24

Table 2: Effects of each proposed component.

B/16 backbone, we follow FPTrans (Zhang et al. 2022) by resizing all images to 480×480 , while extracting low-level features after patch embedding. All images are augmented following FPTrans during training. We use SGD to optimize our model for 5 epochs, with a momentum of 0.9, a weight decay of $5e-4$, and a constant learning rate of $1e-3$. All experiments are implemented on two NVIDIA GeForce RTX 4090 GPUs, with a batch size of 8 for 1-shot and 6 for 5-shot settings. We set the number of scales L for superpixel masks to 4, with the corresponding number of superpixels (local regions) at each scale being $\{5^2, 10^2, 15^2, 20^2\}$. The standard deviation σ_f in the DSR is set to 0.25, and the standard deviations σ_g are set to 0.1 for the ResNet-50 backbone and 0.6 for the ViT-B/16 backbone. Hyperparameters K , α , β , and γ are set to 9, 0.2, 40.0, and 0.1, respectively.

Comparison with State-of-the-art Methods

In Tab. 1, we compare our method with existing FSS and CD-FSS approaches, including CNN-based and ViT-based methods. The results demonstrate significant improvements achieved by our method across different backbones under both 1-shot and 5-shot settings. Specifically, with the ResNet-50 backbone, our method surpasses the state-of-the-art DRA (Su et al. 2024) by 3.37% and 3.35% under 1-shot and 5-shot settings, respectively. With the ViT-B/16

5×5	10×10	15×15	20×20	mean-IoU
✓	✓	✓	✓	67.05
×	✓	✓	✓	66.34
×	×	✓	✓	65.84
×	×	×	✓	65.43

Table 3: Effects of multi-scale superpixel masks.

Strategy	Deepglobe	ISIC	Chest	FSS	Average
Fixed at 0	44.54	55.79	85.93	81.93	67.05
OTSU	45.57	59.98	85.81	80.43	67.95
PCMT (Ours)	45.77	59.36	85.95	81.89	68.24

Table 4: Effects of threshold computation strategies.

backbone, our method surpasses the state-of-the-art LoEC (Liu et al. 2025) by 3.23% and 1.89% under 1-shot and 5-shot settings. Furthermore, our method outperforms existing methods across all four target domain datasets. These results clearly demonstrate the effectiveness of our method.

Ablation Studies and Visualizations

Effects of each component. As shown in Tab. 2, we conduct ablation studies on each proposed component. The first row represents our CNN-based and ViT-based baselines. DSR enhances the model’s generalization ability by simulating data with diverse foreground-background style differences and overall style variations, improving the average mIoU of the CNN-based baseline by 2.62% and the ViT-based baseline by 2.31% under 1-shot setting. HSM guides the model to mine intra-class consistency and inter-class distinction at different granularities, respectively improving performance by 3.10% and 3.05% for the CNN-based and ViT-based baselines. PCMT introduces prototype confidence-weighted adaptive thresholds to mitigate segmentation ambiguity, further improving the performance.

Effects of multi-scale superpixel masks. We use ViT-B/16 as the backbone and conduct ablation studies under 1-shot setting to validate the necessity of multi-scale superpixel masks. Specifically, we use four scales of superpixel masks in HSM with the number of regions being $\{5^2, 10^2, 15^2, 20^2\}$ and sequentially remove them from coarse to fine. As shown in Tab. 3, the best performance is achieved when using all four scales, and each removal results in performance degradation. These results demonstrate that using multi-scale superpixel masks facilitates the model to capture category relationships at different granularities, thereby better adapting to target domains.

Effects of different threshold computation strategies. To validate the effects of our PCMT, we use ViT-B/16 as the backbone and compare the performance of different thresholding strategies under the 1-shot setting, with the results shown in Tab. 4. The adaptive threshold calculated by OTSU (Otsu et al. 1975) demonstrates significant performance improvement on the ISIC dataset, which is prone to segmen-

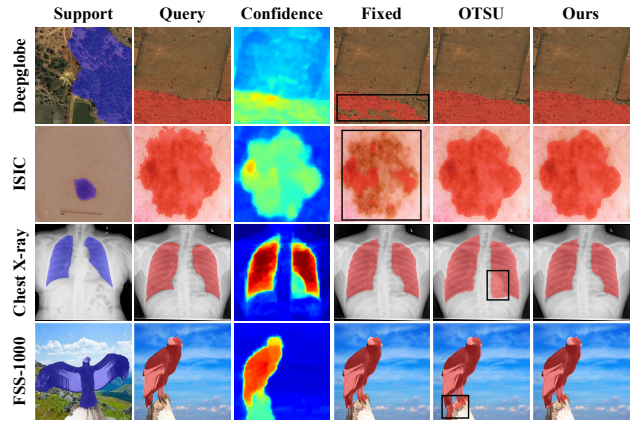


Figure 4: The heatmaps of foreground confidence maps and segmentation results using different thresholding methods.

tation ambiguity. However, it shows a notable performance drop on the FSS-1000 dataset, which is close to the source domain. When the threshold is fixed at 0, the opposite result is observed. Our PCMT dynamically adjusts thresholds based on the prototype confidence of each sample, achieving a better balance across different target domain data.

Visualizations. We use ground truth masks to extract foreground prototypes and calculate similarity maps, with the corresponding heatmaps visualized in Fig. 3. Our method extracts more hierarchical features from the source domain, enabling discrimination of semantics at varying granularities in target domains. Additionally, we visualize the heatmaps of the foreground confidence maps and the segmentation results using different thresholding methods, as shown in Fig. 4. For cases with segmentation ambiguity (top two rows), our method effectively alleviates these ambiguities. For cases without segmentation ambiguity (bottom two rows), our method avoids introducing additional errors at the boundaries. This demonstrates the flexibility and effectiveness of our method. We also present extensive ablation studies and more visualizations in the appendix.

Conclusion

In this paper, we propose an HSL framework to address both inter-domain style gaps and segmentation granularity gaps in CD-FSS. Specifically, we propose a DSR module that sequentially performs foreground style randomization and global style randomization on training data, simulating target domain data with diverse foreground-background style differences and overall style variations. Additionally, we propose an HSM module that leverages multi-scale superpixel masks to extract multi-scale region prototypes to enhance features. DSR and HSM facilitate the model in extracting hierarchical semantic features, enabling better generalization to target domains. Furthermore, we propose a PCMT module to mitigate segmentation ambiguity, further improving segmentation performance. Extensive experiments show that our method achieves state-of-the-art performance on four datasets with different domain gaps.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under the Grant No. 62371235, in part by the Key Research and Development Plan of Jiangsu Province under Grant No. BE2023008-2.

References

- Bai, Y.; Zhao, Z.; Gong, Y.; and Wei, X. 2024. Artrackv2: Prompting autoregressive tracker where to look and how to describe. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 19048–19057.
- Candemir, S.; Jaeger, S.; Palaniappan, K.; Musco, J. P.; Singh, R. K.; Xue, Z.; Karagyris, A.; Antani, S.; Thoma, G.; and McDonald, C. J. 2013. Lung segmentation in chest radiographs using anatomical atlases with nonrigid registration. *IEEE transactions on medical imaging*, 33(2): 577–590.
- Chen, G.; Peng, P.; Ma, L.; Li, J.; Du, L.; and Tian, Y. 2021. Amplitude-phase recombination: Rethinking robustness of convolutional neural networks in frequency domain. In *Proceedings of the IEEE/CVF international conference on computer vision*, 458–467.
- Codella, N.; Rotemberg, V.; Tschandl, P.; Celebi, M. E.; Dusza, S.; Gutman, D.; Helba, B.; Kalloo, A.; Liopyris, K.; Marchetti, M.; et al. 2019. Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic). *arXiv preprint arXiv:1902.03368*.
- Demir, I.; Koperski, K.; Lindenbaum, D.; Pang, G.; Huang, J.; Basu, S.; Hughes, F.; Tuia, D.; and Raskar, R. 2018. Deepglobe 2018: A challenge to parse the earth through satellite images. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 172–181.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88: 303–338.
- Fan, Q.; Pei, W.; Tai, Y.-W.; and Tang, C.-K. 2022. Self-support few-shot semantic segmentation. In *European Conference on Computer Vision*, 701–719. Springer.
- Hariharan, B.; Arbeláez, P.; Bourdev, L.; Maji, S.; and Malik, J. 2011. Semantic contours from inverse detectors. In *2011 international conference on computer vision*, 991–998. IEEE.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- He, W.; Zhang, Y.; Zhuo, W.; Shen, L.; Yang, J.; Deng, S.; and Sun, L. 2024. Apspeg: auto-prompt network for cross-domain few-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23762–23772.
- Herzog, J. 2024. Adapt before comparison: A new perspective on cross-domain few-shot segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 23605–23615.
- Hu, J.; Huang, L.; Ren, T.; Zhang, S.; Ji, R.; and Cao, L. 2023. You only segment once: Towards real-time panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17819–17829.
- Huang, L.; Zhao, X.; and Huang, K. 2019. Got-10k: A large high-diversity benchmark for generic object tracking in the wild. *IEEE transactions on pattern analysis and machine intelligence*, 43(5): 1562–1577.
- Jaeger, S.; Karagyris, A.; Candemir, S.; Folio, L.; Siegelman, J.; Callaghan, F.; Xue, Z.; Palaniappan, K.; Singh, R. K.; Antani, S.; et al. 2013. Automatic tuberculosis screening using chest radiographs. *IEEE transactions on medical imaging*, 33(2): 233–245.
- Lei, S.; Zhang, X.; He, J.; Chen, F.; Du, B.; and Lu, C.-T. 2022. Cross-domain few-shot semantic segmentation. In *European conference on computer vision*, 73–90. Springer.
- Li, G.; Jampani, V.; Sevilla-Lara, L.; Sun, D.; Kim, J.; and Kim, J. 2021. Adaptive prototype learning and allocation for few-shot segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8334–8343.
- Li, X.; Wei, T.; Chen, Y. P.; Tai, Y.-W.; and Tang, C.-K. 2020. Fss-1000: A 1000-class dataset for few-shot segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2869–2878.
- Liao, Y.; Xie, J.; and Geiger, A. 2022. Kitti-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3): 3292–3310.
- Liu, Y.; Zhang, X.; Zhang, S.; and He, X. 2020. Part-aware prototype network for few-shot semantic segmentation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX 16*, 142–158. Springer.
- Liu, Y.; Zou, Y.; Li, Y.; and Li, R. 2025. The devil is in low-level features for cross-domain few-shot segmentation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 4618–4627.
- Nie, J.; Xing, Y.; Zhang, G.; Yan, P.; Xiao, A.; Tan, Y.-P.; Kot, A. C.; and Lu, S. 2024. Cross-domain few-shot segmentation via iterative support-query correspondence mining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3380–3390.
- Otsu, N.; et al. 1975. A threshold selection method from gray-level histograms. *Automatica*, 11(285-296): 23–27.
- Peng, B.; Tian, Z.; Wu, X.; Wang, C.; Liu, S.; Su, J.; and Jia, J. 2023. Hierarchical dense correlation distillation for few-shot segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23641–23651.

- Peng, S.-F.; Sun, G.; Li, Y.; Wang, H.; and Xie, G.-S. 2025. SAM-Aware Graph Prompt Reasoning Network for Cross-Domain Few-Shot Segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 6488–6496.
- Su, J.; Fan, Q.; Pei, W.; Lu, G.; and Chen, F. 2024. Domain-rectifying adapter for cross-domain few-shot segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24036–24045.
- Sun, Y.; Chen, J.; Zhang, S.; Zhang, X.; Chen, Q.; Zhang, G.; Ding, E.; Wang, J.; and Li, Z. 2024. VRP-SAM: SAM with visual reference prompt. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23565–23574.
- Tong, J.; Zou, Y.; Chen, G.; Li, Y.; and Li, R. 2025. Self-Disentanglement and Re-Composition for Cross-Domain Few-Shot Segmentation. *arXiv preprint arXiv:2506.02677*.
- Tong, J.; Zou, Y.; Li, Y.; and Li, R. 2024. Lightweight frequency masker for cross-domain few-shot semantic segmentation. *Advances in Neural Information Processing Systems*, 37: 96728–96749.
- Tschandl, P.; Rosendahl, C.; and Kittler, H. 2018. The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data*, 5(1): 1–9.
- Wang, A.; Chen, H.; Liu, L.; Chen, K.; Lin, Z.; Han, J.; et al. 2024. Yolov10: Real-time end-to-end object detection. *Advances in Neural Information Processing Systems*, 37: 107984–108011.
- Wang, K.; Liew, J. H.; Zou, Y.; Zhou, D.; and Feng, J. 2019. Panet: Few-shot image semantic segmentation with prototype alignment. In *proceedings of the IEEE/CVF international conference on computer vision*, 9197–9206.
- Wang, Y.; Wei, Y.; Qian, X.; Zhu, L.; and Yang, Y. 2021. AINet: Association implantation for superpixel segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7078–7087.
- Xu, Q.; Lin, G.; Loy, C. C.; Long, C.; Li, Z.; and Zhao, R. 2024a. Eliminating feature ambiguity for few-shot segmentation. In *European Conference on Computer Vision*, 416–433. Springer.
- Xu, Q.; Zhao, W.; Lin, G.; and Long, C. 2023. Self-calibrated cross attention network for few-shot segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 655–665.
- Xu, S.; Wei, S.; Ruan, T.; and Liao, L. 2024b. Learning invariant inter-pixel correlations for superpixel generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 6351–6359.
- Yang, B.; Liu, C.; Li, B.; Jiao, J.; and Ye, Q. 2020a. Prototype mixture models for few-shot semantic segmentation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VIII 16*, 763–778. Springer.
- Yang, F.; Sun, Q.; Jin, H.; and Zhou, Z. 2020b. Superpixel segmentation with fully convolutional networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13964–13973.
- Yu, F.; Chen, H.; Wang, X.; Xian, W.; Chen, Y.; Liu, F.; Madhavan, V.; and Darrell, T. 2020. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2636–2645.
- Zhang, C.; Lin, G.; Liu, F.; Yao, R.; and Shen, C. 2019. 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*, 5217–5226.
- Zhang, J.-W.; Sun, Y.; Yang, Y.; and Chen, W. 2022. Feature-proxy transformer for few-shot segmentation. *Advances in neural information processing systems*, 35: 6575–6588.
- Zhang, R.; Jiang, Z.; Guo, Z.; Yan, S.; Pan, J.; Ma, X.; Dong, H.; Gao, P.; and Li, H. 2023. Personalize segment anything model with one shot. *arXiv preprint arXiv:2305.03048*.