

Foreground-Covering Prototype Generation and Matching for SAM-Aided Few-Shot Segmentation

Suho Park*, SuBeen Lee*, Hyun Seok Seong, Jaejoon Yoo, Jae-Pil Heo[†]

Sungkyunkwan University
{shpa06972, leesb7426, gustjrd195, ujaejoon, jaepilheo}@gmail.com

Abstract

We propose Foreground-Covering Prototype Generation and Matching to resolve Few-Shot Segmentation (FSS), which aims to segment target regions in unlabeled query images based on labeled support images. Unlike previous research, which typically estimates target regions in the query using support prototypes and query pixels, we utilize the relationship between support and query prototypes. To achieve this, we utilize two complementary features: SAM Image Encoder features for pixel aggregation and ResNet features for class consistency. Specifically, we construct support and query prototypes with SAM features and distinguish query prototypes of target regions based on ResNet features. For the query prototype construction, we begin by roughly guiding foreground regions within SAM features using the conventional pseudo-mask, then employ iterative cross-attention to aggregate foreground features into learnable tokens. Here, we discover that the cross-attention weights can effectively alternate the conventional pseudo-mask. Therefore, we use the attention-based pseudo-mask to guide ResNet features to focus on the foreground, then infuse the guided ResNet feature into the learnable tokens to generate class-consistent query prototypes. The generation of the support prototype is conducted symmetrically to that of the query one, with the pseudo-mask replaced by the ground-truth mask. Finally, we compare these query prototypes with support ones to generate prompts, which subsequently produce object masks through the SAM Mask Decoder. Our state-of-the-art performances on various datasets validate the effectiveness of the proposed method for FSS.

Introduction

Semantic segmentation has achieved remarkable advancements in recent years (Chen et al. 2015; Badrinarayanan, Kendall, and Cipolla 2017; Strudel et al. 2021; Cheng et al. 2022; Xie et al. 2021). However, these successes heavily rely on large amounts of labeled data, which require significant human effort. To reduce this burden, Few-Shot Segmentation (FSS) has been introduced (Shaban et al. 2017a; Shi et al. 2022; Zhang et al. 2021; Peng et al. 2023; Wang, Sun, and Zhang 2023). FSS aims to segment an unlabeled image,

*Equal contribution

[†]Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

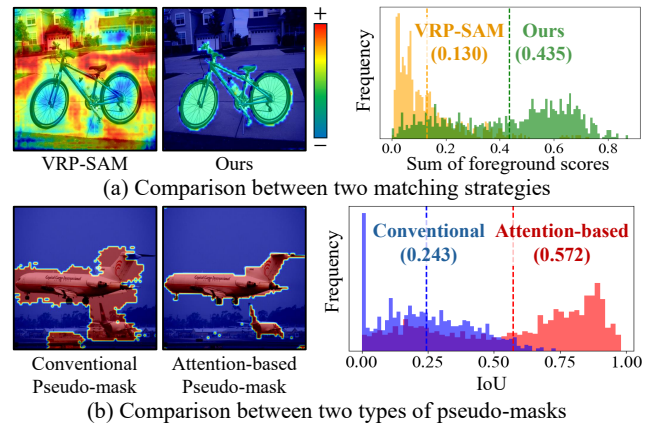


Figure 1: Comparison between VRP-SAM and Ours. (a) (left) We visualize pixel-wise attention maps of query image compared to support prototype. (right) Summing the scores corresponding to the foreground, our prototype-to-prototype matching achieves a higher average score than prototype-to-pixel matching (VRP-SAM). (b) We compare the conventional and attention-based pseudo masks generated by VRP-SAM and our method. The visualizations and IoU distribution with query foreground validate the effectiveness of our attention-based pseudo-mask.

known as the query image, using a small number of labeled images, referred to as the support image.

The Segment Anything Model (SAM) (Kirillov et al. 2023) has demonstrated remarkable versatility across various segmentation tasks. Although SAM excels at generating object masks with appropriate prompts within an image, effectively utilizing SAM for FSS remains challenging. This limitation arises from the absence of class consistency across images in SAM, which means it cannot classify the foreground regions of the query image based on the support images.

Recently, VRP-SAM (Sun et al. 2024) introduced an alternative approach to leverage SAM for FSS by generating visual reference prompts that encode query foreground pixels corresponding to support foreground pixels, and producing an object mask via the SAM Mask Decoder. Specifically, VRP-SAM encapsulates support foreground pixels into learnable tokens, which serve as support prototypes. It then generates

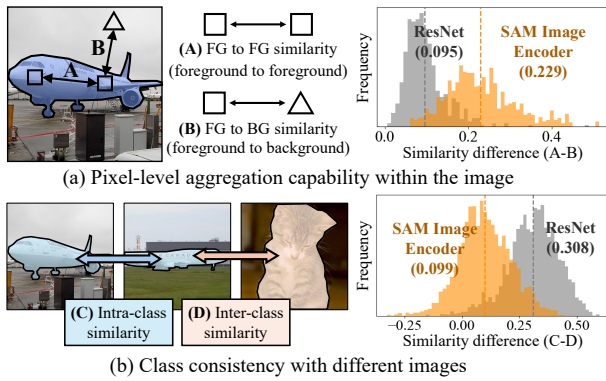


Figure 2: Comparison between ResNet and SAM Image Encoder features on 1000 PASCAL VOC images in test dataset, demonstrating their complementary strengths. (a) Similarity difference between FG-to-FG and FG-to-BG within the image. The higher similarity difference shown in SAM Image Encoder features reflects its superior pixel-level aggregation for prototype construction. (b) Difference between the FG-to-FG similarity from the same class (intra-class) and different classes (inter-class). Compared to SAM features, ResNet features convey better class consistency across different images.

the visual reference prompts by identifying query pixels that match these support prototypes. Note that, this process is built upon ResNet feature, ensuring class consistency in distinguishing foreground query pixels based on the support prototypes.

Although the pipeline by VRP-SAM is practical, it still has limitations. Since the relationship between prototypes and pixels is sub-optimal as depicted in Fig. 1 (a), this could result in visual reference prompts that either lack sufficient foreground information or contain background elements. Furthermore, as VRP-SAM uses conventional pseudo-masks based on a simple pixel-pixel similarity, the quality of a conventional pseudo-mask to enhance the query foreground-specific information is low, as described in Fig. 1 (b). Therefore, VRP-SAM struggles to selectively enhance only the query foreground pixels, potentially blurring the distinction between foreground and background pixels.

To address these issues, we propose Foreground-Covering Prototype Generation and Matching, a method that constructs prototypes from both query and support images that are representative features of each and compares between prototypes to generate reliable visual reference prompts. Since our method can produce foreground-covering and class-consistent prototypes even within the query, the introduced prototype-to-prototype matching produces more reliable results than prototype-pixel matching as depicted in Fig. 1 (a). This leads to an improvement in the quality of visual reference prompts. To make prototypes, our approach leverages two types of complementary features: SAM Image Encoder features to group foreground features with its superior aggregation capability (Fig. 2 (a)), and ResNet features, which ensures class consistency (Fig. 2 (b)). Specifically, we employ cross-attention with learnable tokens to sequentially

encode aggregated foreground information from the SAM features and class-consistent properties from ResNet features. To effectively aggregate foreground information in the query image, we first roughly guide the foreground features from the SAM features using conventional pseudo-masks. Then, our iterative cross-attention between the SAM features and the learnable tokens ensures a gradual concentration of query foreground features. For instance, even if the initial pseudo-mask captures only parts of the foreground, the robust aggregation of the SAM features can progressively uncover more of the foreground. Conversely, if the pseudo-mask over-predicts and includes background elements, the SAM features can gradually refine it to better focus more on the foreground. In the meantime, we discover that the cross-attention weights between learnable tokens and SAM features effectively alternate to the conventional pseudo-mask, as visualized in Fig. 1 (b). Therefore, we utilize this attention-based pseudo-mask to highlight the foreground information in the ResNet features. Specifically, we infuse foreground-centric class consistency into learnable tokens guided by the attention-based pseudo-mask and ResNet features, to create query prototypes. This approach allows the prototypes to maintain the class consistency of ResNet features while leveraging the aggregation power of the SAM features.

The process of generating support prototypes is symmetrical to the query prototype generation process, except that the pseudo-masks are replaced by a ground-truth mask. Consequently, we construct visual reference prompts by matching the query prototypes with the support prototypes.

Overall, our contributions are summarized as follows:

- We propose a Foreground-Covering Prototype Generation and Matching for few-shot segmentation, which constructs prototypes in both query and support images and compares them to generate visual reference prompts and produce an object mask of query image via the SAM Mask Decoder.
- To effectively generate foreground-centric prototype, we leverage two types of complementary features: SAM Image Encoder features for its superior aggregation capability and ResNet features having class consistency.
- We propose the attention-based pseudo-mask that can effectively replace the conventional pseudo-mask by leveraging SAM Image Encoder features.
- The effectiveness of our Prototype2Prototype Matching is validated by achieving new state-of-the-art performances across diverse datasets for Few-Shot Segmentation.

Related Work

Vision Foundation Model

Vision foundation models have demonstrated remarkable adaptability across a wide range of tasks (Radford et al. 2021; Zhou, Loy, and Dai 2022; Caron et al. 2021; Rao et al. 2022). Among them, Segment Anything Model (SAM) (Kirillov et al. 2023) stands out as a foundation model for segmentation. Due to its promptable design and powerful zero-shot capability, SAM has gained significant attention for its versatility across diverse applications (Li et al. 2023; Chen et al.

2023; Yu et al. 2023; Ma et al. 2024) and domains (Roy et al. 2023; Tang, Xiao, and Li 2023; Huang et al. 2023).

Few-Shot Segmentation

Few-Shot Segmentation (FSS) aims to segment the foreground mask for novel classes within a query image with only a few labeled support images. Specifically, given support images and masks for a specific object, FSS aims to find the mask for the same semantic object within a query image. FSS has been studied through two mainstreams: prototype-based methods and affinity learning methods. The prototype-based methods (Li et al. 2021; Zhang et al. 2020; Wang et al. 2019; Liu et al. 2020) represent the support foreground as prototypes and utilize them for prediction. On the other hand, the affinity learning methods (Park et al. 2024; Xu et al. 2023; Hong et al. 2022; Xu et al. 2025; Min, Kang, and Cho 2021; Xu et al. 2024) predict query mask by utilizing pixel-level dense correlation between the support and query images.

Recently, leveraging the zero-shot segmentation capability of the Segment Anything Model (SAM), VRP-SAM (Sun et al. 2024) introduced a method to generate appropriate prompts as input for the SAM’s mask decoder in the context of FSS. Nevertheless, the pixel-wise comparison in VRP-SAM has limitations in effectively distinguishing foreground from background. In this regard, we show that prototype-wise comparison is more effective in generating foreground-specific prompts, given SAM’s strong grouping capabilities within a single image.

Method

Problem Definition

Few-Shot Segmentation (FSS) aims to segment target regions in an unlabeled image by referencing labeled areas in example images. To accomplish this, FSS typically employs a meta-learning approach. In this framework, two separate datasets are used: $\mathcal{D}_{\text{train}}$ for training and $\mathcal{D}_{\text{test}}$ for evaluation. These datasets are composed of distinct classes, $\mathcal{C}_{\text{base}}$ and $\mathcal{C}_{\text{novel}}$, which do not overlap ($\mathcal{C}_{\text{base}} \cap \mathcal{C}_{\text{novel}} = \emptyset$). Under the above configuration, FSS carries out the training and evaluation through multiple episodes. In each episode, there are a support set with K labeled images $S = \{(I_i^S, M_i^S)\}_{i=1}^K$ and a query set with an unlabeled image $Q = (I^Q, M^Q)$. Images in both sets contain a randomly sampled category corresponding to the current phase. I represents the image, and M denotes the ground-truth mask, where 1 indicates the foreground region (the sampled category) and 0 indicates the background region. As a result, the FSS model is learned to predict the mask of the query set M^{pred} based on the support set and query image. For simplicity, we assume that the support set contains only one image in the below sections.

Overview

The overall pipeline of our method is described in Fig. 3. We aim to construct both support and query prototypes and match those prototypes to generate reliable Visual Reference Prompts (VRPs) for SAM Mask Decoder. To achieve this, we utilize two types of complementary features: SAM Image

Encoder features to gather foreground features with its superior aggregation capability (Fig. 2 (a)) and ResNet features to ensure class consistency (Fig. 2 (b)). In the following, we will describe each component of our method: Support Prototypes, Query Prototypes, and Prototype-Prototype Matching.

Support Prototypes

To produce reliable VRPs by matching support and query prototypes, we first generate Support Prototypes (SPs) that encode foreground information from the support features. Since the SPs will be utilized to distinguish foreground and background information of the query features, they should focus on object-specific details and maintain class consistency with other images. To achieve this, we aggregate foreground information with the SAM features, and then fuse the class consistency of the ResNet features into those aggregated ones.

Given support SAM features $G^S \in \mathbb{R}^{C \times H \times W}$, we first guide foreground features using the ground-truth mask for support $M^S \in \mathbb{R}^{H \times W}$ to facilitate foreground aggregation, as follows:

$$\bar{G}^S = \text{Conv}^G(\text{Concat}(G^S, M^S, \text{MP}(G^S, M^S))), \quad (1)$$

where $\bar{G}^S \in \mathbb{R}^{C \times H \times W}$ is guided SAM features while $\text{Conv}(\cdot)$, $\text{Concat}(\cdot)$, and $\text{MP}(\cdot)$ represent a 1×1 convolution layer, concatenate function, and mask-average pooling with expansion, respectively. Then, we gather foreground information of those guided features into N learnable tokens $P^S \in \mathbb{R}^{N \times C}$ by employing a cross-attention methodology. This aggregation process leverages the high pixel-level class consistency within a single image of SAM features and is implemented in a progressive refinement strategy over $T - 1$ steps. Step by step, we encode foreground-specific information $T - 1$ times by repeatedly applying cross-attention between the learnable tokens and the SAM features. Formally, the t -th aggregated support SAM features with the learnable tokens $P_t^S \in \mathbb{R}^{N \times C}$ is defined as follows:

$$P_t^S = \text{MaskedCrossAttn}(P_{t-1}^S, \bar{G}^S, \bar{G}^S; M^S), \quad (2)$$

where $\text{MaskedCrossAttn}(\cdot)$ denotes cross-attention with a mask condition while P_0^S is the same with the initial learnable tokens P^S .

However, these aggregated features struggle to distinguish foreground and background information of the query features since SAM features lack class consistency with different images compared to ResNet features (Fig. 2). Therefore, we adopt ResNet features for support $F^S \in \mathbb{R}^{C \times H \times W}$ to grant class-consistent properties to aggregated support features. To facilitate this, we guide the support ResNet features with the support ground-truth mask, as follows:

$$\bar{F}^S = \text{Conv}^G(\text{Concat}(F^S, M^S, \text{MP}(F^S, M^S))), \quad (3)$$

where $\bar{F}^S \in \mathbb{R}^{C \times H \times W}$ denotes the guided support ResNet features. Subsequently, we generate SPs $P_T^S \in \mathbb{R}^{N \times C}$ by applying the cross-attention with the aggregated support SAM features P_{T-1}^S , the guided support SAM features \bar{G}^S , and the guided support ResNet features \bar{F}^S as *query*, *key*, and *value*, respectively. This can be formulated as follows:

$$P_T^S = \text{MaskedCrossAttn}(P_{T-1}^S, \bar{G}^S, \bar{F}^S; M^S). \quad (4)$$

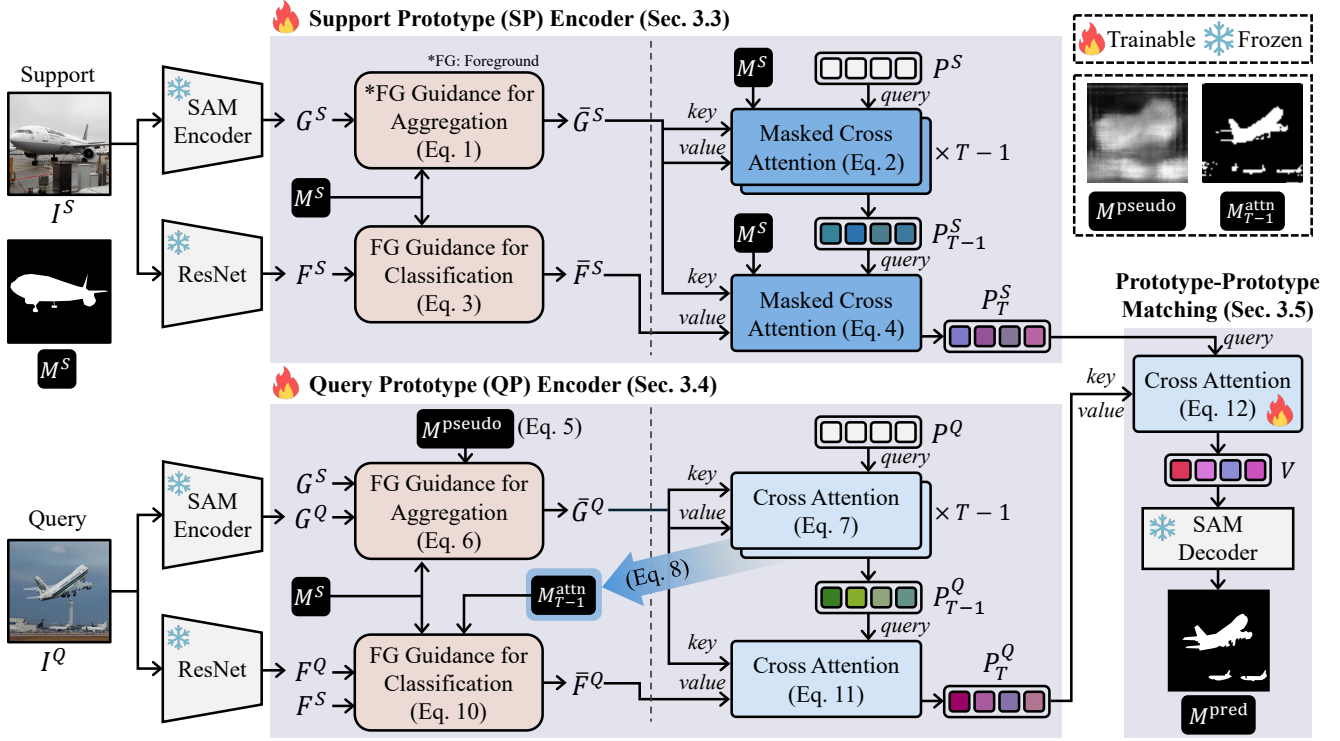


Figure 3: Overall procedure of Foreground-Covering Prototype Generation and Matching. Given the SAM Image Encoder features G , we start by guiding the foreground features using the ground-truth mask for the support M^S and a conventional pseudo-mask for the query M^{pseudo} , then gather these guided features \bar{G} into learnable tokens P through iterative cross-attention. However, SAM features lack class consistency across different images, making it challenging to directly construct prototypes. To address this, we utilize ResNet features F to infuse class-consistent properties into the tokens. We first guide the ResNet features to enhance the foreground-specific information with the ground-truth mask of the support and an attention-based pseudo-mask for the query M_{T-1}^{attn} . The attention-based pseudo-mask, benefiting from SAM’s high aggregation capability, provides better precision compared to the conventional pseudo-mask, as shown in the upper right. By infusing the class consistency of the guided ResNet features \bar{F} into the learnable prompts, we obtain both support and query prototypes (P_T^S and P_T^Q). As a result, the visual reference prompts are generated by matching the query prototypes with the support ones and these prompts are passed to the SAM Decoder to predict a query mask M^{pred} .

Through the above process, we produce SPs that consist of foreground features and possess class consistency with other images.

Query Prototypes

The relationship between the SPs and query ResNet features is still insufficient for a reliable comparison, as described in Fig 1 (a). Therefore, we also construct Query Prototypes (QPs) that each encode query features in the same category to improve trustworthiness and generate reliable VRPs through prototype-prototype comparison. Note that, the largest difference between SPs and QPs is the absence of the ground truth mask.

Given query ResNet features $F^Q \in \mathbb{R}^{C \times H \times W}$ and query SAM features $G^Q \in \mathbb{R}^{C \times H \times W}$, the QPs generation process consists of foreground features aggregation with SAM features and infusing class consistency using ResNet features, similar to SPs. However, unlike the support ground-truth mask, we cannot access the query ground-truth mask. There-

fore, we first compute a conventional pseudo-mask for the query $M^{\text{pseudo}} \in \mathbb{R}^{H \times W}$ by utilizing ResNet features, which have class-consistency between different images, following the VRP-SAM (Sun et al. 2024), as follows:

$$M_{h,w}^{\text{pseudo}} = \max_{1 \leq h' \leq H, 1 \leq w' \leq W} M_{h',w'}^S (F_{h,w}^Q \cdot F_{h',w'}^S). \quad (5)$$

where (\cdot) means the cosine similarity. Using this pseudo-mask, we guide the query SAM features to assist the foreground aggregation, as follows:

$$\bar{G}^Q = \text{Conv}^G(\text{Concat}(G^Q, M^{\text{pseudo}}, \text{MP}(G^S, M^S))), \quad (6)$$

where $\bar{G}^Q \in \mathbb{R}^{C \times H \times W}$ is the guided query SAM features. Here, we use $\text{MP}(G^S, M^S)$ instead of $\text{MP}(G^Q, M^{\text{pseudo}})$ since the quality of the pseudo-mask is low. Subsequently, we employ cross-attention to encode foreground-specific information of guided query SAM features into N learnable tokens $P^Q \in \mathbb{R}^{N \times C}$. Formally, the t -th aggregated query SAM features $P_t^Q \in \mathbb{R}^{N \times C}$ is defined as follows:

$$P_t^Q = \text{CrossAttn}(P_{t-1}^Q, \bar{G}^Q, \bar{G}^Q), \quad (7)$$

where $\text{CrossAttn}(\cdot)$ refers to the cross-attention while P_0^Q is the same with the initial learnable tokens P^Q .

However, unlike SPs that are guided to focus exclusively on foreground regions using the ground truth mask as shown in Eq. 2 and 4, QPs are likely to encode both background and foreground information due to the absence of the ground truth mask. To minimize background information, we explicitly encourage QPs to focus on the foreground by introducing a guide loss.

First, we define an attention-based mask for query $M_t^{\text{attn}} \in \mathbb{R}^{H \times W}$, which is used for defining the guide loss, as follows:

$$M_{t,h,w}^{\text{attn}} = \max_{1 \leq n \leq N} A_{t,n,h,w}^Q \quad (8)$$

where $A_t^Q \in \mathbb{R}^{N \times H \times W}$ represents attention weights of Eq. 7. Then, the guide loss $\mathcal{L}_{\text{guide}}$ can be formulated with the attention-based mask, as follows:

$$\mathcal{L}_{\text{guide}} = \frac{1}{T-1} \sum_{t=1}^{T-1} \mathcal{L}_{\text{BCE}}(M_t^{\text{attn}}, M^Q) + \mathcal{L}_{\text{DL}}(M_t^{\text{attn}}, M^Q), \quad (9)$$

where $M^Q \in \mathbb{R}^{H \times W}$ is the ground truth mask of the query, while \mathcal{L}_{BCE} and \mathcal{L}_{DL} denote the binary cross-entropy loss and the dice loss, respectively.

Furthermore, we observe that the attention-based mask during aggregation stages can effectively replace the conventional pseudo-mask. Thanks to the strong aggregation capability of SAM features, even if the pseudo-mask only captures part of the foreground object, the attention weights progressively reveal more details about the foreground regions. Therefore, when guiding the query ResNet features F^Q to enhance the foreground-specific information, $(T-1)$ -th attention-based mask M_{T-1}^{attn} was used instead of the conventional pseudo-mask, as follows:

$$\bar{F}^Q = \text{Conv}^F(\text{Concat}(F^Q, M_{T-1}^{\text{attn}}, \text{MP}(F^S, M^S))), \quad (10)$$

where $\bar{F}^Q \in \mathbb{R}^{C \times H \times W}$ is the guided query ResNet features. Consequently, we generate QPs $Q_T^P \in \mathbb{R}^{N \times C}$ by applying the cross-attention, as follows:

$$P_T^Q = \text{CrossAttn}(P_{T-1}^Q, \bar{G}^Q, \bar{F}^Q). \quad (11)$$

Although some QPs may encode background information due to the absence of the ground-truth mask in the cross-attention process (Eq. 2 and 4), they still inherit class consistency from ResNet features and benefit from the strong aggregation capability of the SAM features. This facilitates the classification of QPs representing foreground, based on SPs.

Prototype-Prototype Matching

The SAM Mask Decoder requires a reference in the query image to indicate positive or negative areas for segmenting the target region. To provide this reference, we explore QPs corresponding to SPs containing foreground information and produce VRPs $V \in \mathbb{R}^{N \times C}$ based on those results, as follows:

$$V = \text{CrossAttn}(P_T^S, P_T^Q, P_T^Q). \quad (12)$$

Then, we convey VRPs and query SAM features to SAM Mask Decoder to predict the query mask $M^{\text{pred}} \in \mathbb{R}^{H \times W}$.

To ensure that the predicted query mask closely matches the ground truth mask, we define a prompt loss $\mathcal{L}_{\text{prompt}}$ between the predicted and ground-truth masks, as follows:

$$\mathcal{L}_{\text{prompt}} = \mathcal{L}_{\text{BCE}}(M^{\text{pred}}, M^Q) + \mathcal{L}_{\text{DL}}(M^{\text{pred}}, M^Q). \quad (13)$$

where \mathcal{L}_{BCE} and \mathcal{L}_{DL} are the binary cross-entropy and dice loss, respectively.

Additionally, we introduce an orthogonal loss to ensure that the object prompts represent diverse regions by encouraging each SP and QP to encode distinct information. The orthogonal loss $\mathcal{L}_{\text{ortho}}$ is computed as follows:

$$\mathcal{L}_{\text{ortho}} = \frac{1}{(T-1)} \sum_{t=1}^{T-1} \sum_{i \neq j} (A_{t,i}^S \cdot A_{t,j}^S) + (A_{t,i}^Q \cdot A_{t,j}^Q), \quad (14)$$

where $A_{t,n}^S \in \mathbb{R}^{H \times W}$ and $A_{t,n}^Q \in \mathbb{R}^{H \times W}$ represent the cross-attention weights for N -th object prompts from Eq. 2, Eq. 4, Eq. 7 and Eq. 11, respectively, while (\cdot) is the cosine-similarity. To sum up, the total loss for training is defined as follows:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{prompt}} + \lambda_{\text{ortho}} \mathcal{L}_{\text{ortho}} + \lambda_{\text{guide}} \mathcal{L}_{\text{guide}}, \quad (15)$$

where λ_{ortho} and λ_{guide} are coefficients.

Experiments

Datasets To validate our work, we utilize PASCAL-5ⁱ (Shaban et al. 2017b) and COCO-20ⁱ (Nguyen and Todorovic 2019) following the prior works (Sun et al. 2024; Min, Kang, and Cho 2021; Park et al. 2024). Specifically, PASCAL-5ⁱ consists of 20 classes from PASCAL VOC 2012 (Williams 2010) and SDS (Hariharan et al. 2014), while COCO-20ⁱ comprises 80 categories from COCO (Lin et al. 2014) dataset. Each dataset is divided into 4 folds and each fold does not share categories with other folds. To evaluate generalization capability, we use one of the folds as novel classes for testing and others as base categories for training. As a result, we use 15 base and 5 novel classes in PASCAL-5ⁱ while 60 base and 20 novel categories in COCO-20ⁱ. To measure the performance of the model, we randomly sample 1000 support-query pairs from novel classes and compute mean Intersection over Union (mIoU).

Implementation Details Following VRP-SAM, we utilize SAM Image Encoder features and other backbone features with class-consistent properties (VGG and ResNet), both pretrained on ImageNet. We employ the AdamW optimizer, adjusted by a cosine annealing schedule. For PASCAL-5ⁱ, we use 100 epochs and an initial learning rate of 2e-4 for training while 50 epochs and an initial learning rate of 1e-4 in COCO-20ⁱ, and the batch size of both datasets is 8. The number of learnable tokens is 50 for both support and query. Additionally, the number of layers for constructing prototypes is 3, *i.e.*, $T = 3$. The loss coefficients are 0.05 and 0.5 for λ_{ortho} and λ_{guide} , respectively.

Comparison with State-of-the-Art Methods We validate our method using standard benchmarks by comparing it against various FSS models. These include HDM-Net (Peng et al. 2023) and VRP-SAM (Sun et al. 2024)

Method	IE	PASCAL-5 ⁱ										COCO-20 ⁱ									
		1-shot					5-shot					1-shot					5-shot				
		F-0	F-1	F-2	F-3	Avg	F-0	F-1	F-2	F-3	Avg	F-0	F-1	F-2	F-3	Avg	F-0	F-1	F-2	F-3	Avg
HDMNet	VGG	64.8	71.4	67.7	56.4	65.1	68.1	73.1	71.8	64.0	69.3	40.7	50.6	48.2	44.0	45.9	47.0	56.5	54.1	51.9	52.4
VRP-SAM [†]		69.9	73.9	67.6	62.3	68.4	72.8	75.0	67.5	63.2	69.6	39.4	52.0	50.6	47.5	47.4	43.6	57.7	54.7	51.8	51.9
FCP (Ours)		71.9	75.0	69.8	65.5	70.5	73.5	76.4	71.6	65.8	71.8	44.5	54.3	53.5	47.7	50.0	49.0	58.3	55.2	52.3	53.7
PFENet	ResNet	61.7	69.5	55.4	56.3	60.8	63.1	70.7	55.8	57.9	61.9	36.5	38.6	34.5	33.8	35.8	36.5	43.3	37.8	38.4	39.0
CyCTR		65.7	71.0	59.5	59.7	64.0	69.3	73.5	63.8	63.5	67.5	38.9	43.0	39.6	39.8	40.3	41.1	48.9	45.2	47.0	45.6
SSP		60.5	67.8	66.4	51.0	61.4	67.5	72.3	75.2	62.1	69.3	35.5	39.6	37.9	36.7	37.4	40.6	47.0	45.1	43.9	44.1
BAM		69.0	73.6	67.6	61.1	67.8	70.6	75.1	70.8	67.2	70.9	39.4	49.9	46.2	45.2	45.2	43.2	53.4	49.4	48.1	48.5
HDMNet		71.0	75.4	68.9	62.1	69.4	71.3	76.2	71.3	68.5	71.8	43.8	55.3	51.6	49.4	50.0	50.6	61.6	55.7	56.0	56.0
VRP-SAM [†]		74.5	77.3	69.5	65.8	71.8	76.3	76.8	69.5	63.1	71.4	44.3	54.3	52.3	50.0	50.2	50.5	59.5	56.9	54.9	55.5
FCP (Ours)		74.9	77.4	71.8	68.8	73.2	77.2	78.8	72.2	67.7	74.0	46.4	56.4	55.3	51.8	52.5	52.6	63.3	59.8	56.1	58.0

Table 1: Experimental results on the PASCAL-5ⁱ and COCO-20ⁱ. VGG and ResNet denote VGG-16 and ResNet-50, respectively. IE means the type of image encoder while † indicates the reproduced version for the fair comparison.

with a VGG Image Encoder, as well as PFENet (Tian et al. 2020), CyCTR (Zhang et al. 2021), SSP (Fan et al. 2022), BAM (Lang et al. 2022), and VRP-SAM (Sun et al. 2024) with a ResNet Image Encoder. As shown in Table. 1, our method achieves new state-of-the-art performance, proving its effectiveness. This success is consistent across different datasets, shot numbers, and backbone networks, indicating the robustness of our approach. Moreover, the superior performance over VRP-SAM indicates that our generated VRP is more advantageous in producing a higher-quality query mask.

Qualitative Results. We compare our qualitative results with VRP-SAM in Figure 4 and also display both the conventional and attention-based pseudo-masks in VRP-SAM and our method, respectively. Our approach consistently outperforms the baseline across different scene types, whether simple (1st and 2nd columns), multi-object (3rd column), or complex (4th column). This improvement is largely due to the precision of the attention-based pseudo-mask in our method.

Further Analysis

Ablation Study

We perform ablation studies to evaluate the contributions of the main components and key hyperparameters of our method. Most experiments are progressed with ResNet-50 on PASCAL-5ⁱ with all folds except for T ablation study which conducted with only fold-3 of PASCAL-5ⁱ.

Main Components Tab. 2 presents the results of the ablation study for the main components. (a) shows the result of prototype-pixel matching with ResNet-50, which constructs only support prototypes and matches them with the query, which is our baseline model VRP-SAM. (b) shows the result when using SAM features instead of ResNet features under the same condition. There is a significant performance drop because SAM features cannot ensure class-consistent properties. (c) and (d) represent the results when applying our prototype-prototype matching to (a) and (b), respectively. For instance, for (c), the inferior pixel-level aggregation of

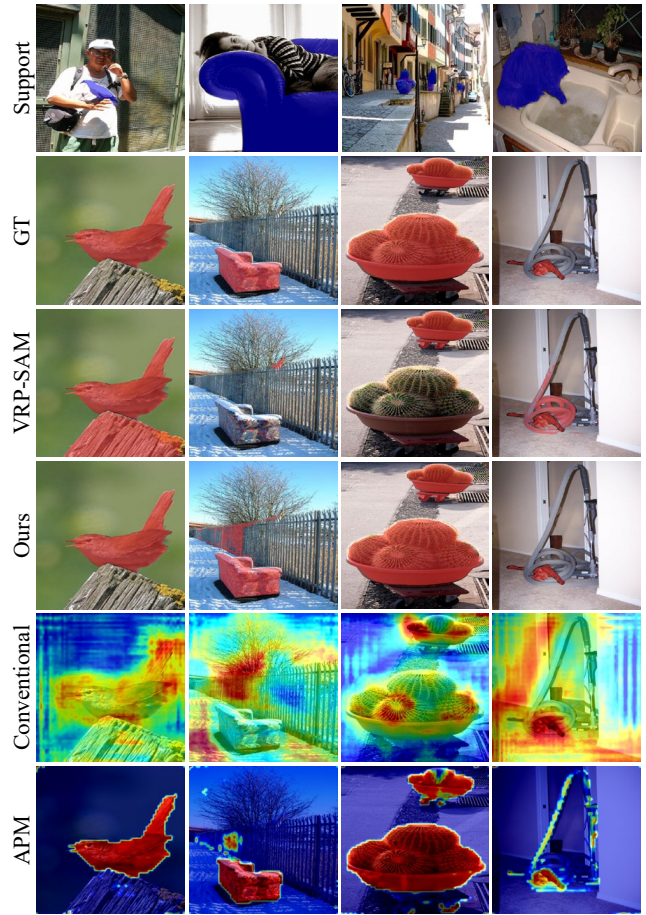


Figure 4: Qualitative comparison results of Ours and VRP-SAM on the PASCAL-5ⁱ dataset. Conventional and APM refer to the conventional pseudo-mask of VRP-SAM and our Attention-based Pseudo-Mask, respectively.

ResNet features causes a performance degradation compared to the baseline. And for (d), the performance still remains poor. Furthermore, by incorporating prototype-prototype match-

	Res	SAM	PPM	APM	mIoU	Δ
(a)	✓				71.8	0.0
(b)		✓			66.2	-5.6
(c)	✓		✓		66.5	-5.3
(d)		✓	✓		66.1	-5.7
(e)	✓	✓	✓		72.6	+0.8
(f)	✓	✓	✓	✓	73.2	+1.4

Table 2: Ablation studies for different components and architectures. Res, SAM, PPM and APM denote ResNet-50, SAM Image Encoder, Prototype-Prototype Matching, and Attention-based Pseudo Mask, respectively.




Figure 5: Ablation study for varying the number of aggregation steps for prototype construction. Prediction and APM denote our model’s query mask prediction and Attention-based Pseudo Mask, respectively.

ing by combining SAM’s strong aggregation capabilities with ResNet’s class-consistent properties in (e), the performance significantly surpasses the baseline. This emphasizes the importance of using complementary features to construct prototypes. However, as the conventional pseudo-mask still limits the potential gains, we employ our attention-based pseudo-mask, which leverages SAM’s superior aggregation capability. In specific, (f) further boosts performance, demonstrating that the attention-based pseudo-mask offers precise guidance for the model in locating the target object.

Number of Step T We present the impact of the number of aggregation step for prototype construction in Fig. 5. For instance, 3 steps of prototype construction achieves optimal performance while using more steps leads to a performance drop. This is because the learnable tokens increasingly focus on smaller areas as they repeatedly apply the softmax function in the cross-attention. This is further supported by a reduction in the recall of the attention-based pseudo-mask. A proper T can prevent excessive shrinkage of the pseudo-mask and generate appropriate prototypes.

Loss Functions In addition to the prompt loss L_{prompt} , which is essential for model training, we employ two supplementary losses: the guide loss L_{guide} and the orthogonal loss L_{ortho} . We validate the effectiveness of each loss function, as presented in Tab 3. (a) represents the performance when we use only the prompt loss. (b) and (c) show the results when we adopt the guide loss and the orthogonal loss with the

	$\mathcal{L}_{\text{prompt}}$	$\mathcal{L}_{\text{guide}}$	$\mathcal{L}_{\text{ortho}}$	mIoU	Δ
(a)	✓			72.3	0.0
(b)	✓	✓		72.7	+0.4
(c)	✓		✓	72.4	+0.1
(d)	✓	✓	✓	73.2	+0.9

Table 3: Ablation studies about each loss function.

Pseudo-Mask	mIoU	Prec	Rec
Conventional	32.4	46.5	53.6
Attention-based	60.9	69.1	79.4

Table 4: Comparison between conventional and attention-based pseudo-masks. Prec and Rec denote precision and recall with respect to the ground-truth mask, respectively.

prompt loss, respectively. In the case of (b), we can improve the performance by guiding the attention-based pseudo-mask to represent the object regions. For (c), each learnable token can be induced to capture various parts, which brings a slight performance gain. When we utilize all losses for training as shown in (d), we can get a large performance enhancement. This verifies the appropriateness of the proposed losses.

Improvement of Pseudo Mask

The attention-based pseudo-mask is a key component of our method, used to enhance foreground-specific information. As such, its quality directly affects the overall performance of the model. Table 4 compares the quality of pseudo-masks and highlights the performance improvements achieved by our attention-based pseudo mask. First, the conventional pseudo-mask, computed by a pixel-level cosine similarity, produces a low mIoU with respect to the ground-truth mask. On the other hand, we introduced our attention-based pseudo-mask by leveraging the strong aggregation capability of the SAM features. As a result, our attention-based pseudo-mask substantially outperforms the conventional ones, demonstrating the superior foreground-covering capability of our method.

Conclusion

In this paper, we proposed Foreground-Covering Prototype Generation and Matching, which constructs support and query prototypes and matches them to generate reliable prompts. To build prototypes, we utilized two complementary features: SAM Image Encoder for pixel aggregation and ResNet for class consistency. We guide SAM features using a pseudo-mask for query prototypes, then employ iterative cross-attention to aggregate foreground features into learnable tokens. Here, we discovered that the cross-attention weights effectively replace the conventional pseudo mask, allowing attention-based pseudo-masks to guide ResNet features to focus on the foreground. Then, we encode the guided ones into the learnable tokens to generate class-consistency query prototypes. Finally, we generated reliable prompts by comparing query and support prototypes, achieving state-of-the-art performance on the few-shot segmentation task.

Acknowledgments

This work was supported in part by MSIT/IITP (No. 2022-0-00680, 2020-0-01821, 2019-0-00421, RS-2024-00459618, RS-2024-00360227, RS-2024-00437102, RS-2024-00437633), and MSIT/NRF (No. RS-2024-00357729).

References

- Badrinarayanan, V.; Kendall, A.; and Cipolla, R. 2017. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(12): 2481–2495.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9650–9660.
- Chen, L.-C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; and Yuille, A. L. 2015. Semantic image segmentation with deep convolutional nets and fully connected crfs. *ICLR*.
- Chen, T.; Zhu, L.; Deng, C.; Cao, R.; Wang, Y.; Zhang, S.; Li, Z.; Sun, L.; Zang, Y.; and Mao, P. 2023. Sam-adapter: Adapting segment anything in underperformed scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3367–3375.
- Cheng, B.; Misra, I.; Schwing, A. G.; Kirillov, A.; and Girdhar, R. 2022. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1290–1299.
- 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.; Girshick, R.; and Malik, J. 2014. Simultaneous detection and segmentation. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part VII 13*, 297–312. Springer.
- Hong, S.; Cho, S.; Nam, J.; Lin, S.; and Kim, S. 2022. Cost aggregation with 4d convolutional swin transformer for few-shot segmentation. In *European Conference on Computer Vision*, 108–126. Springer.
- Huang, Y.; Cao, Y.; Li, T.; Juefei-Xu, F.; Lin, D.; Tsang, I. W.; Liu, Y.; and Guo, Q. 2023. On the robustness of segment anything. *arXiv preprint arXiv:2305.16220*.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.
- Lang, C.; Cheng, G.; Tu, B.; and Han, J. 2022. 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*, 8057–8067.
- Li, F.; Zhang, H.; Sun, P.; Zou, X.; Liu, S.; Yang, J.; Li, C.; Zhang, L.; and Gao, J. 2023. Semantic-sam: Segment and recognize anything at any granularity. *arXiv preprint arXiv:2307.04767*.
- 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.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, 740–755. Springer.
- Liu, Y.; Zhang, X.; Zhang, S.; and He, X. 2020. Part-Aware Prototype Network for Few-Shot Semantic Segmentation. In *European Conference on Computer Vision*, 142–158.
- Ma, J.; He, Y.; Li, F.; Han, L.; You, C.; and Wang, B. 2024. Segment anything in medical images. *Nature Communications*, 15(1): 654.
- Min, J.; Kang, D.; and Cho, M. 2021. Hypercorrelation squeeze for few-shot segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6941–6952.
- Nguyen, K.; and Todorovic, S. 2019. Feature weighting and boosting for few-shot segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 622–631.
- Park, S.; Lee, S.; Hyun, S.; Seong, H. S.; and Heo, J.-P. 2024. Task-Disruptive Background Suppression for Few-Shot Segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 4442–4449.
- 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.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Rao, Y.; Zhao, W.; Chen, G.; Tang, Y.; Zhu, Z.; Huang, G.; Zhou, J.; and Lu, J. 2022. Densclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 18082–18091.
- Roy, S.; Wald, T.; Koehler, G.; Rokuss, M. R.; Disch, N.; Holzschuh, J.; Zimmerer, D.; and Maier-Hein, K. H. 2023. Sam. md: Zero-shot medical image segmentation capabilities of the segment anything model. *arXiv preprint arXiv:2304.05396*.
- Shaban, A.; Bansal, S.; Liu, Z.; Essa, I.; and Boots, B. 2017a. One-shot learning for semantic segmentation. *arXiv preprint arXiv:1709.03410*.
- Shaban, A.; Bansal, S.; Liu, Z.; Essa, I.; and Boots, B. 2017b. One-shot learning for semantic segmentation. *arXiv preprint arXiv:1709.03410*.
- Shi, X.; Wei, D.; Zhang, Y.; Lu, D.; Ning, M.; Chen, J.; Ma, K.; and Zheng, Y. 2022. Dense cross-query-and-support

attention weighted mask aggregation for few-shot segmentation. In *European Conference on Computer Vision*, 151–168. Springer.

Strudel, R.; Garcia, R.; Laptev, I.; and Schmid, C. 2021. Segmenter: Transformer for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7262–7272.

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.

Tang, L.; Xiao, H.; and Li, B. 2023. Can sam segment anything? when sam meets camouflaged object detection. *arXiv preprint arXiv:2304.04709*.

Tian, Z.; Zhao, H.; Shu, M.; Yang, Z.; Li, R.; and Jia, J. 2020. Prior guided feature enrichment network for few-shot segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 44(2): 1050–1065.

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.; Sun, R.; and Zhang, T. 2023. Rethinking the Correlation in Few-Shot Segmentation: A Buoy's View. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7183–7192.

Williams, E. M. V. G. L. 2010. CK Winn J Zisserman A The pascal visual object classes (VOC) challenge. *Int. J. Comput. Vis.*, 88(2): 303.

Xie, E.; Wang, W.; Yu, Z.; Anandkumar, A.; Alvarez, J. M.; and Luo, P. 2021. SegFormer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34: 12077–12090.

Xu, Q.; Lin, G.; Loy, C. C.; Long, C.; Li, Z.; and Zhao, R. 2025. Eliminating feature ambiguity for few-shot segmentation. In *European Conference on Computer Vision*, 416–433. Springer.

Xu, Q.; Liu, X.; Zhu, L.; Lin, G.; Long, C.; Li, Z.; and Zhao, R. 2024. Hybrid mamba for few-shot segmentation. *arXiv preprint arXiv:2409.19613*.

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.

Yu, T.; Feng, R.; Feng, R.; Liu, J.; Jin, X.; Zeng, W.; and Chen, Z. 2023. Inpaint anything: Segment anything meets image inpainting. *arXiv preprint arXiv:2304.06790*.

Zhang, G.; Kang, G.; Yang, Y.; and Wei, Y. 2021. Few-shot segmentation via cycle-consistent transformer. *Advances in Neural Information Processing Systems*, 34: 21984–21996.

Zhang, X.; Wei, Y.; Yang, Y.; and Huang, T. S. 2020. Sg-one: Similarity guidance network for one-shot semantic segmentation. *IEEE transactions on cybernetics*, 50(9): 3855–3865.

Zhou, C.; Loy, C. C.; and Dai, B. 2022. Extract free dense labels from clip. In *European Conference on Computer Vision*, 696–712. Springer.