

# Rethinking Prior Information Generation with CLIP for Few-Shot Segmentation

Jin Wang<sup>1</sup> Bingfeng Zhang<sup>1\*</sup> Jian Pang<sup>1</sup> Honglong Chen<sup>1</sup> Weifeng Liu<sup>1\*</sup>  
<sup>1</sup>China University of Petroleum (East China)

{wangjin, jianpang}@s.upc.edu.cn, {bingfeng.zhang, wfliu, chenhl}@upc.edu.cn

## Abstract

Few-shot segmentation remains challenging due to the limitations of its labeling information for unseen classes. Most previous approaches rely on extracting high-level feature maps from the frozen visual encoder to compute the pixel-wise similarity as a key prior guidance for the decoder. However, such a prior representation suffers from coarse granularity and poor generalization to new classes since these high-level feature maps have obvious category bias. In this work, we propose to replace the visual prior representation with the visual-text alignment capacity to capture more reliable guidance and enhance the model generalization. Specifically, we design two kinds of training-free prior information generation strategy that attempts to utilize the semantic alignment capability of the Contrastive Language-Image Pre-training model (CLIP) to locate the target class. Besides, to acquire more accurate prior guidance, we build a high-order relationship of attention maps and utilize it to refine the initial prior information. Experiments on both the PASCAL-5i and COCO-20i datasets show that our method obtains a clearly substantial improvement and reaches the new state-of-the-art performance. The code is available on the project website<sup>1</sup>.

## 1. Introduction

With the development of deep learning [3, 4], semantic segmentation [16, 29, 34, 51, 63] has made a great progress. Traditional semantic segmentation relies on intensive annotations which is time-consuming and labour-intensive, once the segmentation model encounters samples with limited labeled data, it cannot output accurate prediction, making it difficult to apply in practice. Few-shot segmentation [19, 50, 57] is proposed to address the above problem, which aims to segment novel classes with a few annotated samples during inference. To achieve this, it divides data into a support set and a query set, the images in the query set are segmented using the provided information in

\*Corresponding author.

<sup>1</sup><https://github.com/vangjin/PI-CLIP>

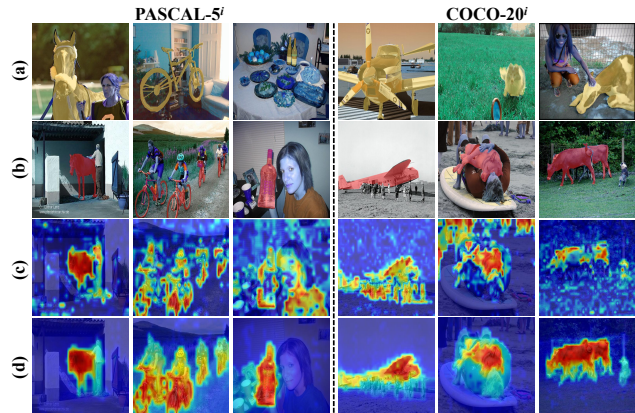


Figure 1. Comparison of prior information. (a) Support images with ground-truth masks; (b) Query images with ground-truth masks; (c) Prior information from previous approaches generated based on the frozen ImageNet [5] weights, which are biased towards some classes, such as the ‘Person’ class; (d) Our prior information, which is generated utilizing the text and visual alignment ability of the frozen CLIP model. Our prior information is finer-grained and mitigates the bias of the class.

the support set as a reference. Existing few-shot segmentation methods can be roughly categorized into two types: pixel-level matching [12, 43, 54, 54, 58] and prototype-level matching [8, 19, 39, 50, 53, 57]. Pixel-level matching uses the pixel-to-pixel matching mechanism [41, 43, 60] to enforce the few-shot model mine pixel-wise relationship [1, 49, 60]. Prototype-level matching methods extract prototypes [14, 26, 37, 52] from the support set to perform similarity [11, 21, 38, 39] or dense comparisons [2, 6, 20, 22, 47] with the query image features to make predictions. No matter pixel-level matching or prototype-level matching, most recent approaches [8, 19, 22, 39, 45] introduce the prior masks [19, 25, 62] as a coarse localization map to guide the matching or segmentation process to concentrate on the located regions. However, such prior masks are mainly generated through interacting fixed high-dimensional features from the visual pre-trained models, *i.e.*, CNN with ImageNet [5] pre-train initialization, causing several insolvable problems as shown in Fig. 1: 1) incorrect target location

response due to original ImageNet [5] pre-training weights being insensitive to category information, which misleads the segmentation process and thus restricting generalization of the model. 2) coarse prior mask shapes, caused by undistinguished vision features between the target and non-target pixels, make the prior information locate many non-target regions, which further confuses the segmentation process.

To address the aforementioned drawbacks, we rethink the prior mask generation strategy and attempt to use Contrastive Language-Image Pre-training (CLIP) [40] to generate more reliable prior information for few-shot segmentation. A large amount of text-image training data pairs make the CLIP model sensitive to category due to the forced text-image alignment, which enables better localization of the target class [32, 61, 64]. Besides, the success in the zero-shot task [9, 24, 40] also demonstrates the powerful generalization ability of the CLIP model. Based on this, we attempt to utilize the CLIP model to generate better prior guidance.

Finally, in this paper, we propose Prior Information Generation with CLIP (PI-CLIP), a training-free CLIP-based approach, to extract prior information to guide the few-shot segmentation. Specifically, we propose two kinds of prior information generation, the first one is called visual-text prior information (VTP) which aims to provide accurate prior location based on the strong visual-text alignment ability of the CLIP model, we re-design target and non-target prompts and force the model to perform category selection for each pixel, thus locating more accurate target regions. The other one is called visual-visual prior information (VVP) which focuses on providing more general prior guidance using the matching map extracted from the CLIP model between the support set and the query image.

However, as a training-free approach, the forced alignment of visual information and text information makes VTP excessively focus on local target regions instead of the expected whole target regions, the incomplete original global structure information only highlights local target regions which reduces the quality of guidance. Based on this, we build a high-order attention matrix based on the attention maps of the CLIP model, called Prior Information Refinement (PIR), to refine the initial VTP, which makes full use of the original pixel-pair structure relationship to highlight the whole target area and reduce the response to the non-target area, thus clearly improving the quality of the prior mask. Note that VVP is not refined to keep its generalization ability. Without any training, the generated prior masks overcome the drawback caused by inaccurate prior information in existing methods, significantly improving the performance of different few-shot approaches.

Our contributions are as follows:

- We rethink the prior information generation for few-shot segmentation, proposing a training-free strategy based on the CLIP model to provide more accurate prior guidance

by mining visual-text alignment information and visual-visual matching information.

- To generate finer-grained prior information, we build a high-order attention matrix to refine the initial prior information based on the frozen CLIP attention maps to extract the relationship of different pixels, clearly improving the quality of the prior information.
- Our method has a significant improvement over existing methods on both PASCAL-5<sup>i</sup> [42] and COCO-20<sup>i</sup> [36] datasets and achieves state-of-the-art performance.

## 2. Related Work

### 2.1. Few-Shot Segmentation

Few-shot segmentation aims to generate dense predictions for new classes using a small number of labeled samples. Most existing few-shot segmentation methods followed the idea of metric-based meta-learning [13, 48]. Depending on the object of the metric, current approaches can be divided into pixel-level matching mechanism [12, 43, 54] and prototype-level matching mechanism [8, 22, 25, 45, 62]. No matter pixel-level matching or prototype-level matching mechanism, most recent approaches [19, 33, 50, 55, 57] utilized prior information to guide the segmentation process.

PCN [30] fused the scores from base and novel classifier to prevent base class bias. CWT [31] adapted the classifier’s weights to each query image in an inductive way. PFENet [50] first proposed to utilize prior information extracted from pixel relationship between support set and query image to guide the decoder and designed a module to aggregate contextual information at different scales. PFENet++ [33] rethought the prior information and proposed to utilize the additional nearby semantic cues for a better location ability of the prior information. BAM [19] further optimized the prior information and proposed to leverage the segmentation of new classes by suppressing the base classes learned by the model. SCL [57] proposed a self-guided learning approach to mine the lost critical information on the prototype and utilize the prior information as guidance for the decoder. IPMT [28] mined useful information by interacting prototype and mask to mitigate the category bias and design an intermediate prototype to mine more accurate prior guidance by an iterative approach. MM-Former [59] utilized a class-specific segmenter to decompose the query image into a single possible prediction and extracted support information as prior to matching the single prediction which can improve the flexibility of the segmentation network. MIANet [55] proposed to use general prior information from semantic word embedding and instance information to perform an accurate segmentation. HDMNet [39] mined pixel-level correlation with transformer based on two kinds of prior information between support set and query image to avoid overfitting.

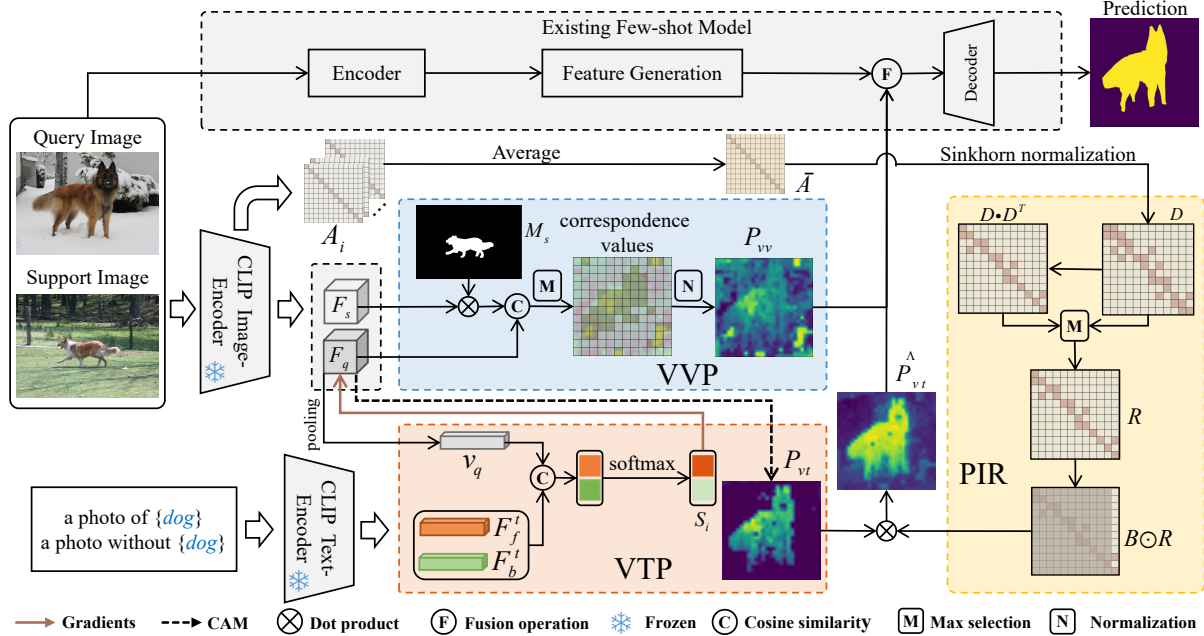


Figure 2. Overview of our proposed PI-CLIP for few-shot segmentation. We design a group of text prompts for a certain class to attract more attention to target regions. The VTP module generates the visual-text prior information by aligning the visual information and text information with the help of softmax-GradCAM. The VVP module generates the visual-visual prior information by a pixel-level similarity calculation. The PIR module is proposed to refine the coarse initial prior information. Finally, the original prior information in the existing few-shot model is directly replaced by VVP and refined VTP, after passing the decoder, the final prediction is generated.

Most recent existing methods utilized coarse masks to guide segmentation, our approach attempts to generate finer-grained masks with the help of CLIP models.

## 2.2. Contrastive Language-Image Pretraining

Contrastive Language-Image Pretraining (CLIP) [40] is able to map text and image into high-dimensional space by text-encoder and image-encoder respectively. Trained on a large amount of text-image data makes the CLIP [15, 40] model has a strong feature extraction capability, which is used in many downstream applications such as detection [17], segmentation [24, 44, 56], and so on. CLIPSeg [32] first attempted to introduce the CLIP model into few-shot segmentation. However, CLIPSeg is more like to use the CLIP model as a validation method to show the powerful capability of the CLIP model in few-shot tasks. In this paper, we design a new prior information generation strategy using the CLIP model for few-shot segmentation through the visual-text relationship and the visual-visual relationship to perform a more efficient guidance.

## 3. Method

### 3.1. Task Description

Few-shot segmentation aims to segment novel classes by using the model trained on base classes. Most existing

few-shot segmentation approaches follow the meta-learning paradigm. The model is optimized with multiple meta-learning tasks in the training phase and evaluates the performance of the model in the testing phase. Given a dataset  $D$ , dividing it into a training set  $D_{train}$  and a test set  $D_{test}$ , there has no crossover between the class set  $C_{train}$  in the training set and class set  $C_{test}$  in the test set ( $C_{train} \cap C_{test} = \emptyset$ ). The model is expected to transfer the knowledge in  $D_{train}$  with restricted labeled data to the  $D_{test}$ . Both training set  $D_{train}$  and test set  $D_{test}$  are composed of support set  $S$  and query set  $Q$ , support set  $S$  contains  $K$  samples  $S = \{S_1, S_2, \dots, S_K\}$ , each  $S_i$  contains an image-mask pair  $\{I_s, M_s\}$  and query set  $Q$  contains  $N$  samples  $Q = \{Q_1, Q_2, \dots, Q_N\}$ , each  $Q_i$  contains an image-mask pair  $\{I_q, M_q\}$ . During training, the few-shot model is optimized with training set  $D_{train}$  by epochs where the model performs prediction for query image  $I_q$  with the guidance of the support set  $S$ . During inference, the performance will be acquired with the test set  $D_{test}$ , and the model is no longer optimized.

### 3.2. Method Overview

In order to enhance the ability of prior information to localize target categories as well as to produce more generalized prior information, we propose to mine visual-text and visual-visual information instead of purely visual feature

similarity to guide the segmentation process. Besides, to further improve the quality of the prior information to get finer-grained guidance, we design an attention map-based high-order matrix to refine the initial prior information by pixel-pairs relationships, Fig. 2 shows our framework of the one-shot case with the following steps:

1. Given a support image and a query image with the target class name, we first input the query image and support image to the CLIP image encoder to generate corresponding visual support and query features. Meanwhile, the target class name is used to build two text prompts, *i.e.*, target prompt and non-target prompt, which are then input to the CLIP text encoder to generate two text embeddings.
2. Then, two text embeddings and the query visual features are input to the visual-text prior (VTP) module to generate the initial VTP information by enforcing a classification process for each pixel.
3. Meanwhile, the support visual features and query visual features are input to the visual-visual prior (VVP) module where the VVP information is generated through the pixel-level relationship.
4. After that, we extract attention maps from the clip model, which are input to our prior information refinement (PIR) module to build a high-order attention matrix for refining the above initial VTP information.
5. Finally, the original prior information in the existing method is directly replaced by our VVP and refined VTP to generate the final prediction for the query image.

### 3.3. Visual-Text Prior Information Generation

Few-shot Segmentation (FSS) remains one major challenge that an image might have more than one class, but the model is required to segment only one class at each episode. This challenge means that once the prior information is unable to provide the correct target region, *e.g.*, a true target region is “dog” but the prior information provides a “cat” region, it will confuse the FSS model to segment the true target pixels, especially for the untrained novel class. To correctly locate target regions, we utilize the visual-text alignment information from the CLIP model to produce a new prior information called VTP. We innovatively define a group of text prompts of the target class as a guidance to the model, in which the target (foreground) text prompts  $t_f$  is defined as “a photo of  $\{target\ class\}$ ” and the non-target (background) text prompts  $t_b$  is “a photo without  $\{target\ class\}$ ”.

Based on the designed text prompts, a pixel-level classification is performed for the query image so as to locate the true target foreground regions. To force the model to decide whether one pixel is the target or not, we use softmax-GradCAM [24] to generate the prior information using the relationship between the visual and text features. Specifically, the designed target and non-target prompts,

*i.e.*, “a photo of  $\{target\ class\}$ ” and “a photo without  $\{target\ class\}$ ”, are sent to the CLIP text encoder to get the high dimensional text features, represented as  $F_f^t$  and  $F_b^t$ . Suppose the query image is  $I_q$ , after passing the CLIP visual encoder, the query features  $F_q^v \in \mathbb{R}^{d \times (hw+1)}$ , after removing the class token in  $F_q^v$ , visual query feature  $F_q \in \mathbb{R}^{d \times hw}$  is generated, then the query token  $v_q$  is obtained through global average pooling:

$$v_q = \frac{1}{hw} \sum_{i=1}^{hw} F_q(i), v_q \in \mathbb{R}^{d \times 1}. \quad (1)$$

Then classification scores are obtained by performing a distance calculation between the text features and the query token after the softmax operation:

$$S_i = softmax\left(\frac{v_q^T F_i^T}{\|v_q\| \|F_i^T\|} / \tau\right), i \in \{f, b\}, \quad (2)$$

where  $T$  represents the matrix transposition and  $\tau$  is a temperature parameter. Then the gradient is calculated based on the final classification score:

$$w_m = \frac{1}{hw} \sum_i \sum_j \frac{\partial S_f}{\partial F_q^m(i, j)}, \quad (3)$$

where  $w_m$  is the weight for  $m$ -th feature map of the foreground regions,  $F_q^m$  means the activation value for  $m$ -th feature map and  $(i, j)$  means the pixel position.

Finally, the visual-text prior information  $P_{vt} \in \mathbb{R}^{1 \times h \times w}$  is obtained:

$$P_{vt} = ReLU\left(\sum_m w_m F_q^m\right), \quad (4)$$

$ReLU$  means the ReLU activation function to filter the negative response. Due to the forced alignment of the semantic information from the visual modal and text modal with softmax-GradCAM, the generated visual-text prior information clearly locates accurate target regions, which avoids the confusion of the segmentation process.

### 3.4. Visual-Visual Prior Information Generation

We enforce VTP to make a classification for each pixel so that it can locate the correct region. However, we observe that VTP tends to locate a discriminative local region, *e.g.*, the “head” region of a “dog” rather than the whole region. To overcome this drawback, we attempt to take advantage of the support information that is naturally present in few-shot segmentation and get region-larger and location-rougher prior information to give more generalized guidance to the model.

We design VVP to mine more general target information by performing matching on the visual-visual relationship



between the support image feature and the query image feature. Suppose the support image is  $I_s$ , after passing through the CLIP image encoder, its high dimensional image feature is generated and the visual support feature is  $F_s^v \in \mathbb{R}^{d \times hw}$  (class token is removed). To get more target-focused support information, we extract the target information from the support image:

$$F_s = F_s^v \odot M_s, \quad (5)$$

where  $M_s$  represents the support mask, which is required to downsample to the same height and width as the feature map. Then we perform a cosine similarity calculation between all pixel pairs for  $f_s^i \in F_s$  and  $f_q^j \in F_q$  as:

$$\cos(f_s^i, f_q^j) = \frac{(f_s^i)^T f_q^j}{\|f_s^i\| \|f_q^j\|} \quad i, j \in \{0, 1, 2, \dots, hw\}. \quad (6)$$

For each pixel in  $F_q$ , the maximum similarity is selected from all pixels in the support feature as the correspondence value:

$$P_{vv}(j) = \max_{i, j \in \{1, 2, \dots, hw\}} \cos(f_s^i, f_q^j). \quad (7)$$

After computing all correspondence value by the above equation, prior information is generated, the values in  $P_{vv}$  are normalized by a min-max normalization to generate the initial visual-visual prior information,  $P_{vv} \in \mathbb{R}^{1 \times h \times w}$ :

$$P_{vv} = \frac{P_{vv} - \min(P_{vv})}{\max(P_{vv}) - \min(P_{vv}) + \varepsilon}, \quad (8)$$

where  $\varepsilon$  is set to  $10^{-7}$ . We utilize the feature from the CLIP model which contains more reliable semantic information to acquire the visual-visual prior information, thus matching support information with the query image, the model can provide more general location information as the prior guidance.

### 3.5. Prior Information Refinement

The above prior information is generated by the visual and textual features extracted from the frozen CLIP weights. As a training-free method, the representation of the prior information can not adaptively guide the model to perform an efficient segmentation. To generate finer-grained prior information that focuses more target regions, we propose a Prior Information Refinement (PIR) module to refine the initial prior information. PIR builds a high-order matrix based on the attention map from the query image, which can accurately build the pixel-wise relationship and retain the original global structure information, thus efficiently capturing spatial information and details of semantics to refine the prior information. In this way, the refined prior information pays more attention to the whole target regions and focuses less on non-target regions.

Specifically, suppose  $A_i \in \mathbb{R}^{hw \times hw}$  is the multi-head self-attention map generated from CLIP with the  $i$ -th block,

to acquire more accurate attention maps for each image, we first compute the average attention map by:

$$\bar{A} = \frac{1}{l} \sum_{i=n-l}^n A_i, \quad (9)$$

where  $l$  and  $n$  are the block number of the vision transformer in CLIP and  $l < n$ . Based on the average attention map, in order to eliminate as much as possible the influence of the background region while preserving the intrinsic structural information, we design a high-order refinement matrix  $R \in \mathbb{R}^{1 \times h \times w}$  follows:

$$R = \max(D, (D \cdot D^T)), D = \text{Sinkhorn}(\bar{A}), \quad (10)$$

where *Sinkhorn* means Sinkhorn normalization [46] to aligning data from rows and columns. We then utilize the refinement matrix  $R$  to refine the initial coarse prior information from VTP and VVP by:

$$\hat{P}_i = B \odot R \cdot P_i, \{i \in vt, vv\}, \quad (11)$$

where  $B$  is a box mask generated from the prior mask following [24] and  $\odot$  represents the Hadamard product. We experimentally found that only refining the visual-text prior  $P_{vt}$  is enough since the refinement matrix will make  $P_{vt}$  and  $P_{vv}$  produce similar responses, which will damage the generalization of the model. Therefore, we select the refined text-visual prior and initial visual-visual prior, *i.e.*,  $\hat{P}_{vt}$  and  $P_{vv}$ , as the final prior information.

Finally, we directly replace the prior information in existing methods with the concatenation of our visual-visual prior information  $P_{vv}$  and refined visual-text prior information  $\hat{P}_{vt}$ , to generate the final prediction.

## 4. Experiments

**Datasets and Evaluation Metrics.** We utilize the PASCAL-5<sup>i</sup> [42] and COCO-20<sup>i</sup> [36] to evaluate the performance of our proposed method. PASCAL-5<sup>i</sup> is built on PASCAL VOC 2012 [7] with the complement of SDS [10] which is a classical computer vision dataset for segmentation tasks including 20 different object classes such as people, cars, cats, dogs, chairs, aeroplanes, etc. COCO-20<sup>i</sup> is built on MSCOCO [23] consists of more than 120,000 images from 80 categories and is a more challenging dataset. To evaluate the performance of our proposed method, we adopt mean intersection-over-union (mIoU) and foreground-background IoU (FB-IoU) as the evaluation metrics following previous works [19, 39, 50].

### 4.1. Implementation details.

We utilize HDMNet [39], BAM [19] and PFENet [50] as the baseline to test our performance. In all experiments on PASCAL-5<sup>i</sup> and COCO-20<sup>i</sup>, the images are set to  $473 \times 473$

Table 1. Performance comparisons with mIoU (%) as a metric on PASCAL-5<sup>i</sup>, “ours-PI-CLIP (PFENet)”, “ours-PI-CLIP (BAM)” and “ours-PI-CLIP (HDMNet)” represent the baseline is PFENet [50], BAM [19] and HDMNet [39] respectively.

Method	Backbone	1-shot					5-shot				
		Fold0	Fold1	Fold2	Fold3	Mean	Fold0	Fold1	Fold2	Fold3	Mean
SCL (CVPR’21) [57]	resnet50	63.0	70.0	56.5	57.7	61.8	64.5	70.9	57.3	58.7	62.9
SSP (ECCV’22) [8]	resnet50	60.5	67.8	66.4	51.0	61.4	67.5	72.3	75.2	62.1	69.3
DCAMA (ECCV’22) [43]	resnet50	67.5	72.3	59.6	59.0	64.6	70.5	73.9	63.7	65.8	68.5
NERTNet (CVPR’22) [27]	resnet50	65.4	72.3	59.4	59.8	64.2	66.2	72.8	61.7	62.2	65.7
IPMT (NeurIPS’22) [28]	resnet50	72.8	73.7	59.2	61.6	66.8	73.1	74.7	61.6	63.4	68.2
ABCNet (CVPR’23) [53]	resnet50	68.8	73.4	62.3	59.5	66.0	71.7	74.2	65.4	67.0	69.6
MIANet (CVPR’23) [55]	resnet50	68.5	75.8	67.5	63.2	68.8	70.2	77.4	70.0	68.8	71.6
MSI (ICCV’23) [35]	resnet50	71.0	72.5	63.8	65.9	68.3	73.0	74.2	66.6	70.5	71.1
PFENet (TPAMI’20) [50]	resnet50	61.7	69.5	55.4	56.3	60.8	63.1	70.7	55.8	57.9	61.9
BAM (CVPR’22) [19]	resnet50	68.9	73.6	67.6	61.1	67.8	70.6	75.1	70.8	67.2	70.9
HDMNet (CVPR’23) [39]	resnet50	71.0	75.4	68.9	62.1	69.4	71.3	76.2	71.3	68.5	71.8
ours-PI-CLIP (PFENet)	resnet50	67.4	76.5	71.3	69.4	71.2	70.4	78.2	72.4	70.2	72.8
ours-PI-CLIP (BAM)	resnet50	72.4	80.2	71.6	70.5	73.7	72.6	80.6	73.5	72.0	74.7
ours-PI-CLIP (HDMNet)	resnet50	<b>76.4</b>	<b>83.5</b>	<b>74.7</b>	<b>72.8</b>	<b>76.8</b>	<b>76.7</b>	<b>83.8</b>	<b>75.2</b>	<b>73.2</b>	<b>77.2</b>

Table 2. Performance comparisons on COCO-20<sup>i</sup>, “ours-PI-CLIP (HDMNet)” represent the baseline is HDMNet [39].

Method	Backbone	1-shot					5-shot				
		Fold0	Fold1	Fold2	Fold3	Mean	Fold0	Fold1	Fold2	Fold3	Mean
SCL (CVPR’21) [57]	resnet50	36.4	38.6	37.5	35.4	37.0	38.9	40.5	41.5	38.7	39.9
SSP (ECCV’22) [8]	resnet101	39.1	45.1	42.7	41.2	42.0	47.4	54.5	50.4	49.6	50.2
DCAMA (ECCV’22) [43]	resnet50	41.9	45.1	44.4	41.7	43.3	45.9	50.5	50.7	46.0	48.3
BAM (CVPR’22) [18]	resnet50	43.4	50.6	47.5	43.4	46.2	49.3	54.2	51.6	49.5	51.2
NERTNet (CVPR’22) [27]	resnet101	38.3	40.4	39.5	38.1	39.1	42.3	44.4	44.2	41.7	43.2
IPMT (NeurIPS’22) [28]	resnet50	41.4	45.1	45.6	40.0	43.0	43.5	49.7	48.7	47.9	47.5
ABCNet (CVPR’23) [53]	resnet50	42.3	46.2	46.0	42.0	44.1	45.5	51.7	52.6	46.4	49.1
MIANet (CVPR’23) [55]	resnet50	42.5	53.0	47.8	47.4	47.7	45.9	58.2	51.3	52.0	51.7
MSI (ICCV’23) [35]	resnet50	42.4	49.2	49.4	46.1	46.8	47.1	54.9	54.1	51.9	52.0
PFENet (TPAMI’20) [50]	resnet101	34.3	33.0	32.3	30.1	32.4	38.5	38.6	38.2	34.3	37.4
HDMNet (CVPR’23) [39]	resnet50	43.8	55.3	51.6	49.4	50.0	50.6	61.6	55.7	56.0	56.0
ours-PI-CLIP (PFENet)	resnet50	36.1	42.3	37.3	37.7	38.4	40.4	45.6	39.9	38.6	41.1
ours-PI-CLIP (HDMNet)	resnet50	<b>49.3</b>	<b>65.7</b>	<b>55.8</b>	<b>56.3</b>	<b>56.8</b>	<b>56.4</b>	<b>66.2</b>	<b>55.9</b>	<b>58.0</b>	<b>59.1</b>

pixels and the CLIP pre-trained model is ViT-B-16 [40]. For COCO-20<sup>i</sup>, setting higher resolution can get higher performance but with more computing cost, the temperature parameter  $\tau$  in VTP is set to 0.01 and the selected layer  $l$  in PIR is set to 8. For the 5-shot case, we directly concatenate 5 VVP rather than using the average of them as the prior information. For fair comparisons, other settings like data augmentation technique, learning rate and optimizer, *e.g.*, all follow the corresponding baselines. All experiments are run on NVIDIA V100 GPUs.

With the help of the accurate visual-text prior information and the generalized visual-visual prior information, our proposed PI-CLIP method can able to reach better performance quickly, so PI-CLIP is only trained for 30 epochs on both PASCAL-5<sup>i</sup> and COCO-20<sup>i</sup> which needs less time than any previous methods and the batch sizes are set to 4 on

1-shot and 2 on 5-shot respectively, the model can perform better if can be trained for more epochs.

## 4.2. Comparison with state-of-the-art

**Quantitative results.** Table 1 shows the performance of our method and existing state-of-the-art methods for few-shot segmentation on PASCAL-5<sup>i</sup>, our approach greatly improves the performance of the model over the 1-shot task compared to different baselines and achieves new state-of-the-art performance, with mIoU increases of 5.9% for BAM [19] and 7.4% for HDMNet [39]. For the 5-shot segmentation task, our approach outperforms other approaches by a clear margin, with mIoU gain of 3.8% for BAM [19] and 5.4% for HDMNet [39], respectively. Besides, we also experimented by plugging our method into PFENet[50], a different baseline from BAM [19] and HDMNet [39] that



Figure 3. Qualitative results of the proposed PI-CLIP and baseline (HDMNet [39]) approach under 1-shot setting. Each row from top to bottom represents the support images with ground-truth (GT) masks (green), query images with GT masks (blue), baseline results (red), and our results (yellow), respectively.

does not use a base learner, it can be seen even without the inhibition of base classes by the base learner, our approach also improves the mIoU of 10.4% and 10.9% for 1-shot and 5-shot tasks respectively. The performance improvement of the different baseline methods shows that our method is a plug-and-play module with high flexibility. The main reasons for our success with different approaches are the accurate localization of VTP and the strong generalization of VVP.

In Table 2, we compare the performance of our approach and others on COCO-20<sup>i</sup> dataset. Our approach also exhibits strong performance and achieves new state-of-the-art performance. Specifically, our approach improves the baseline by 6.8% and 3.1% mIoU for 1-shot and 5-shot tasks.

**Qualitative results.** In order to better show the effect of our proposed model on the existing methods, we visualize the results of the baseline and our proposed method in Fig. 3, it can be found that our method (yellow part) has a much stronger target localization ability than the baseline (red part), and the bias on the base class is greatly reduced.

Fig. 4 shows the visualization of our proposed VTP and VVP to help understand the localization capabilities of VTP and the generalization capabilities of VVP. VTP focuses more on the accurate target regions, which are localized in a local region compared to the whole object. VVP, on the other hand, focuses on larger regions of the target class than VTP, but the details provided by VVP are tougher than VTP. Fig. 4 also shows that synchronous refining VVP and VTP information makes them similar which is harmful to the generalization of the few-shot segmentation model.

Table 3. Ablation study about our proposed VTP and VVP on the PASCAL-5<sup>i</sup>, “baseline” represents the HDMNet[39], VTP and VVP represent the proposed VTP module and VVP module.

baseline	VTP	VVP	mIoU (%)	FB-IoU (%)
✓			71.00	85.86
✓	✓		75.30	86.77
✓	✓	✓	<b>76.40</b>	<b>87.57</b>

Table 4. Ablation study about our proposed PIR on PASCAL-5<sup>i</sup>,  $P_{vv}$  and  $P_{vt}$  represent the initial information generated by VTP and VVP,  $PIR_{vv}$  and  $PIR_{vt}$  represent the refinement on VVP and VTP.

$P_{vv}$	$P_{vt}$	$PIR_{vv}$	$PIR_{vt}$	mIoU (%)	FB-IoU (%)
✓	✓			75.40	86.61
✓	✓	✓		74.82	86.05
✓	✓		✓	<b>76.40</b>	<b>87.57</b>
✓	✓	✓	✓	75.70	86.93

### 4.3. Ablation Study

We conduct a series of ablation studies to investigate the impact of each module on the PASCAL-5<sup>i</sup> dataset using HDMNet [39] as the baseline.

**Ablation Study on VVP and VTP.** The prior information has a large impact on the performance of the model, so we conduct relevant ablation studies to separately verify the validity of the prior information for the two modules we



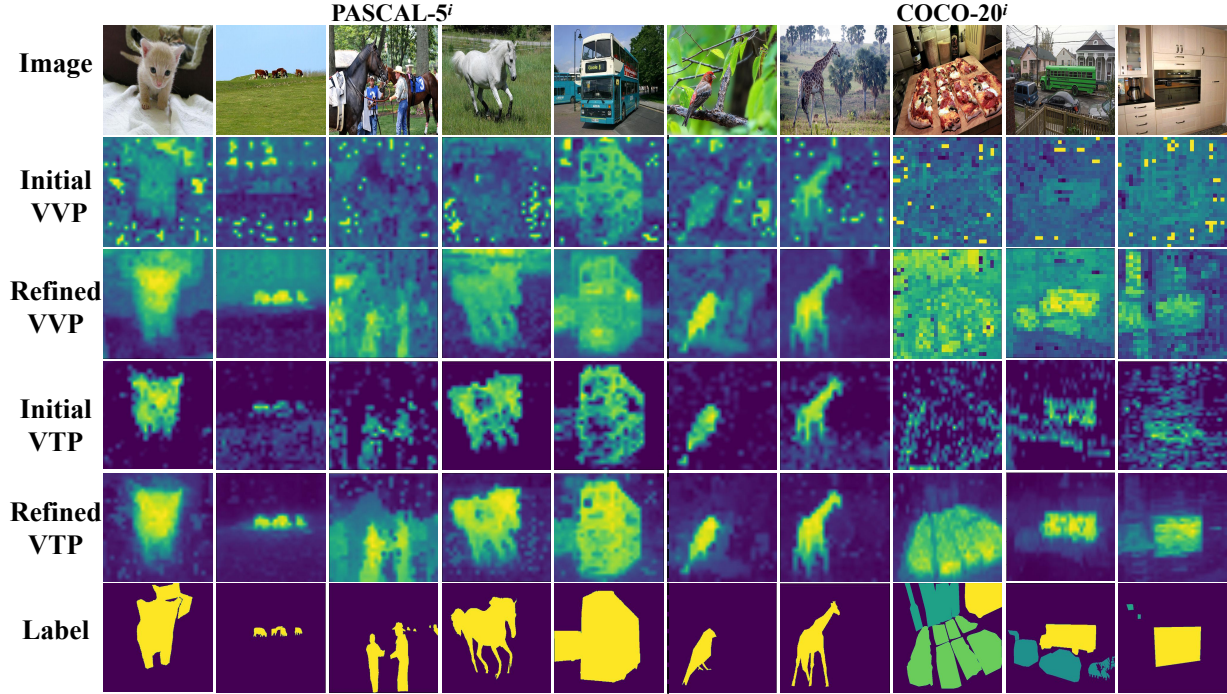


Figure 4. Visualization of the different prior information generated by our proposed method. The left is sampled from PASCAL-5<sup>i</sup> [42] and the right is selected from COCO-20<sup>i</sup> [36]. Each row from top to bottom represents the query image, initial visual-visual prior information, refined visual-visual prior information, initial visual-text prior information and refined visual-text prior information. The  $P_{vv}$  has more general localization regions and the  $P_{vt}$  has more local target regions. With the refinement of the designed high-order matrix, more accurate prior information can be extracted.

designed. As can be seen in Table 3, VTP yields a performance improvement of 4.3% and VVP yields a performance improvement of 1.1%.

**Ablation Study on PIR.** In the PIR module we designed a high order matrix to maintain the structural information of the original features and used it to refine the initial prior information, we conduct ablation experiments on the refinement ability of PIR as shown in Table 4. It can be found that the refinement of PIR using only the VTP information is able to get an enhancement of 1.0%, but the refinement of PIR using only the VVP information as prior information reduces model performance by 0.58%, this is due to the fact that refining VVP and VTP based on the same matrix will make them produce a similar response, which will reduce the generalization of the prior guidance. When both the initial VVP and refined VTP are used, the model is able to achieve the highest performance of 76.40%.

## 5. Conclusion

In this paper, we rethink the prior information for few-shot segmentation and realize that CLIP is able to achieve more accurate localization of the target class without further training. The proposed prior information generation with CLIP (PI-CLIP) can give more accurate and generalized

prior information which facilitates the segmentation performance. Furthermore, we design two prior information generation modules, one is VTP which aligns the semantic information from the visual modal and text modal to generate accurate prior information, and the other is VVP which performs a matching on visual feature between support image and query image to mine more useful target information and give a regionally larger prior information. To extract more useful information, the PIR module is designed to refine the initial prior information. Extensive experiments demonstrate the effectiveness of our proposed module. In the future, we will explore how to better extract the useful information from the CLIP model.

**Acknowledge:** This work was supported by National Natural Science Foundation of China (No. 62301613, 62372468), the Taishan Scholar Program of Shandong (No. tsqn202306130), the Shandong Natural Science Foundation (No. ZR2023QF046, ZR2023MF008), the Major Basic Research Projects in Shandong Province (Grant No.ZR2023ZD32), the Qingdao Natural Science Foundation (Grant No. 23-2-1-161-zyyd-jch), Qingdao Postdoctoral Applied Research Project (No. QDBSH20230102091) and Independent Innovation Research Project of China University of Petroleum (East China) (No. 22CX06060A).



## References

- [1] Hanbo Bi, Yingchao Feng, Zhiyuan Yan, Yongqiang Mao, Wenhui Diao, Hongqi Wang, and Xian Sun. Not just learning from others but relying on yourself: A new perspective on few-shot segmentation in remote sensing. *IEEE Transactions on Geoscience and Remote Sensing*, 2023. 1
- [2] Jiacheng Chen, Bin-Bin Gao, Zongqing Lu, Jing-Hao Xue, Chengjie Wang, and Qingmin Liao. Apanet: adaptive prototypes alignment network for few-shot semantic segmentation. *arXiv preprint arXiv:2111.12263*, 2021. 1
- [3] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34:17864–17875, 2021. 1
- [4] Jiequan Cui, Zhisheng Zhong, Zhuotao Tian, Shu Liu, Bei Yu, and Jiaya Jia. Generalized parametric contrastive learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 1
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 1, 2
- [6] Henghui Ding, Hui Zhang, and Xudong Jiang. Self-regularized prototypical network for few-shot semantic segmentation. *Pattern Recognition*, 133:109018, 2023. 1
- [7] 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. 5
- [8] Qi Fan, Wenjie Pei, Yu-Wing Tai, and Chi-Keung Tang. Self-support few-shot semantic segmentation. In *European Conference on Computer Vision*, pages 701–719. Springer, 2022. 1, 2, 6
- [9] Ziyu Guo, Renrui Zhang, Longtian Qiu, Xianzheng Ma, Xupeng Miao, Xuming He, and Bin Cui. Calip: Zero-shot enhancement of clip with parameter-free attention. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 746–754, 2023. 2
- [10] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *2011 international conference on computer vision*, pages 991–998. IEEE, 2011. 5
- [11] Shuting He, Xudong Jiang, Wei Jiang, and Henghui Ding. Prototype adaption and projection for few-and zero-shot 3d point cloud semantic segmentation. *IEEE Transactions on Image Processing*, 2023. 1
- [12] Sunghwan Hong, Seokju Cho, Jisu Nam, Stephen Lin, and Seungryong Kim. Cost aggregation with 4d convolutional swin transformer for few-shot segmentation. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIX*, pages 108–126. Springer, 2022. 1, 2
- [13] Tao Hu, Pengwan Yang, Chiliang Zhang, Gang Yu, Yadong Mu, and Cees GM Snoek. Attention-based multi-context guiding for few-shot semantic segmentation. In *Proceedings of the AAAI conference on artificial intelligence*, pages 8441–8448, 2019. 2
- [14] Kai Huang, Feigege Wang, Ye Xi, and Yutao Gao. Prototypical kernel learning and open-set foreground perception for generalized few-shot semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19256–19265, 2023. 1
- [15] Tianyu Huang, Bowen Dong, Yunhan Yang, Xiaoshui Huang, Rynson WH Lau, Wanli Ouyang, and Wangmeng Zuo. Clip2point: Transfer clip to point cloud classification with image-depth pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22157–22167, 2023. 3
- [16] Siyu Jiao, Yunchao Wei, Yaowei Wang, Yao Zhao, and Humphrey Shi. Learning mask-aware clip representations for zero-shot segmentation. *Advances in Neural Information Processing Systems*, 36:35631–35653, 2023. 1
- [17] Chen Ju, Peisen Zhao, Siheng Chen, Ya Zhang, Xiaoyun Zhang, Yanfeng Wang, and Qi Tian. Adaptive mutual supervision for weakly-supervised temporal action localization. *IEEE Transactions on Multimedia*, 2022. 3
- [18] Dahyun Kang and Minsu Cho. Integrative few-shot learning for classification and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9979–9990, 2022. 6
- [19] 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. 1, 2, 5, 6
- [20] Chunbo Lang, Binfei Tu, Gong Cheng, and Junwei Han. Beyond the prototype: Divide-and-conquer proxies for few-shot segmentation. *arXiv preprint arXiv:2204.09903*, 2022. 1
- [21] Yuan-Hao Lee, Fu-En Yang, and Yu-Chiang Frank Wang. A pixel-level meta-learner for weakly supervised few-shot semantic segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2170–2180, 2022. 1
- [22] 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. 1, 2
- [23] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 5
- [24] Yuqi Lin, Minghao Chen, Wenxiao Wang, Boxi Wu, Ke Li, Binbin Lin, Haifeng Liu, and Xiaofei He. Clip is also an efficient segmenter: A text-driven approach for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15305–15314, 2023. 2, 3, 4, 5
- [25] Jie Liu, Yanqi Bao, Guo-Sen Xie, Huan Xiong, Jan-Jakob Sonke, and Efstratios Gavves. Dynamic prototype convolution network for few-shot semantic segmentation. In *Pro-*

- ceedings of the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11553–11562, 2022. [1](#), [2](#)
- [26] Nian Liu, Kepan Nan, Wangbo Zhao, Yuanwei Liu, Xiwen Yao, Salman Khan, Hisham Cholakkal, Rao Muhammad Anwer, Junwei Han, and Fahad Shahbaz Khan. Multi-grained temporal prototype learning for few-shot video object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18862–18871, 2023. [1](#)
- [27] Yuanwei Liu, Nian Liu, Qinglong Cao, Xiwen Yao, Junwei Han, and Ling Shao. Learning non-target knowledge for few-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11573–11582, 2022. [6](#)
- [28] 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](#), [6](#)
- [29] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. [1](#)
- [30] 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. [2](#)
- [31] Zhihe Lu, Sen He, Da Li, Yi-Zhe Song, and Tao Xiang. Prediction calibration for generalized few-shot semantic segmentation. *IEEE Transactions on Image Processing*, 2023. [2](#)
- [32] Timo Lüddecke and Alexander Ecker. Image segmentation using text and image prompts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7086–7096, 2022. [2](#), [3](#)
- [33] Xiaoli Luo, Zhuotao Tian, Taiping Zhang, Bei Yu, Yuan Yan Tang, and Jiaya Jia. Pfenet++: Boosting few-shot semantic segmentation with the noise-filtered context-aware prior mask. *arXiv preprint arXiv:2109.13788*, 2021. [2](#)
- [34] Shervin Minaee, Yuri Boykov, Fatih Porikli, Antonio Plaza, Nasser Kehtarnavaz, and Demetri Terzopoulos. Image segmentation using deep learning: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3523–3542, 2021. [1](#)
- [35] Seonghyeon Moon, Samuel S Sohn, Honglu Zhou, Sejong Yoon, Vladimir Pavlovic, Muhammad Haris Khan, and Mubbasir Kapadia. Msi: Maximize support-set information for few-shot segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19266–19276, 2023. [6](#)
- [36] 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. [2](#), [5](#), [8](#)
- [37] Atsuro Okazawa. Interclass prototype relation for few-shot segmentation. In *European Conference on Computer Vision*, pages 362–378. Springer, 2022. [1](#)
- [38] Prashant Pandey, Aleti Vardhan, Mustafa Chasmai, Tanuj Sur, and Brejesh Lall. Adversarially robust prototypical few-shot segmentation with neural-odes. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 77–87. Springer, 2022. [1](#)
- [39] 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](#), [5](#), [6](#), [7](#)
- [40] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [2](#), [3](#), [6](#)
- [41] Kate Rakelly, Evan Shelhamer, Trevor Darrell, Alexei A Efros, and Sergey Levine. Few-shot segmentation propagation with guided networks. *arXiv preprint arXiv:1806.07373*, 2018. [1](#)
- [42] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots. One-shot learning for semantic segmentation. *arXiv preprint arXiv:1709.03410*, 2017. [2](#), [5](#), [8](#)
- [43] Xinyu Shi, Dong Wei, Yu Zhang, Donghuan Lu, Munan Ning, Jiashun Chen, Kai Ma, and Yefeng Zheng. Dense cross-query-and-support attention weighted mask aggregation for few-shot segmentation. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XX*, pages 151–168. Springer, 2022. [1](#), [2](#), [6](#)
- [44] Chen Shuai, Meng Fanman, Zhang Runtong, Qiu Heqian, Li Hongliang, Wu Qingbo, and Xu Linfeng. Visual and textual prior guided mask assemble for few-shot segmentation and beyond. *arXiv preprint arXiv:2308.07539*, 2023. [3](#)
- [45] Mennatullah Siam, Boris Oreshkin, and Martin Jagersand. Adaptive masked proxies for few-shot segmentation. *arXiv preprint arXiv:1902.11123*, 2019. [1](#), [2](#)
- [46] Richard Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *The annals of mathematical statistics*, 35(2):876–879, 1964. [5](#)
- [47] Haoliang Sun, Xiankai Lu, Haochen Wang, Yilong Yin, Xiantong Zhen, Cees GM Snoek, and Ling Shao. Attentional prototype inference for few-shot segmentation. *Pattern Recognition*, page 109726, 2023. [1](#)
- [48] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1199–1208, 2018. [2](#)
- [49] Antonio Tavera, Fabio Cermelli, Carlo Masone, and Barbara Caputo. Pixel-by-pixel cross-domain alignment for few-shot semantic segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1626–1635, 2022. [1](#)
- [50] 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, 5, 6
- [51] Zhuotao Tian, Jiequan Cui, Li Jiang, Xiaojuan Qi, Xin Lai, Yixin Chen, Shu Liu, and Jiaya Jia. Learning context-aware classifier for semantic segmentation. *arXiv preprint arXiv:2303.11633*, 2023. 1
- [52] 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. 1
- [53] Yuan Wang, Rui Sun, and Tianzhu Zhang. Rethinking the correlation in few-shot segmentation: A buoys view. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7183–7192, 2023. 1, 6
- [54] Xianghui Yang, Bairun Wang, Kaige Chen, Xinchu Zhou, Shuai Yi, Wanli Ouyang, and Luping Zhou. Brinet: Towards bridging the intra-class and inter-class gaps in one-shot segmentation. *arXiv preprint arXiv:2008.06226*, 2020. 1, 2
- [55] Yong Yang, Qiong Chen, Yuan Feng, and Tianlin Huang. Mianet: Aggregating unbiased instance and general information for few-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7131–7140, 2023. 2, 6
- [56] Yuhuan Yang, Chaofan Ma, Chen Ju, Ya Zhang, and Yanfeng Wang. Multi-modal prototypes for open-set semantic segmentation. *arXiv preprint arXiv:2307.02003*, 2023. 3
- [57] 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. 1, 2, 6
- [58] 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. 1
- [59] Gengwei Zhang, Shant Navasardyan, Ling Chen, Yao Zhao, Yunchao Wei, Honghui Shi, et al. Mask matching transformer for few-shot segmentation. *Advances in Neural Information Processing Systems*, 35:823–836, 2022. 2
- [60] Miao Zhang, Miaojing Shi, and Li Li. Mfnet: Multiclass few-shot segmentation network with pixel-wise metric learning. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(12):8586–8598, 2022. 1
- [61] Renrui Zhang, Wei Zhang, Rongyao Fang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free adaption of clip for few-shot classification. In *European Conference on Computer Vision*, pages 493–510. Springer, 2022. 2
- [62] Xiaolin Zhang, Yunchao Wei, Zhao Li, Chenggang Yan, and Yi Yang. Rich embedding features for one-shot semantic segmentation. *IEEE Transactions on Neural Networks and Learning Systems*, 33(11):6484–6493, 2021. 1, 2
- [63] Zekang Zhang, Guangyu Gao, Jianbo Jiao, Chi Harold Liu, and Yunchao Wei. Coinseg: Contrast inter-and intra-class representations for incremental segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 843–853, 2023. 1
- [64] Xiangyang Zhu, Renrui Zhang, Bowei He, Aojun Zhou, Dong Wang, Bin Zhao, and Peng Gao. Not all features matter: Enhancing few-shot clip with adaptive prior refinement. *arXiv preprint arXiv:2304.01195*, 2023. 2