

Iterative Few-shot Semantic Segmentation from Image Label Text

Haohan Wang^{1*}, Liang Liu^{2*}, Wuhao Zhang², Jiangning Zhang², Zhenye Gan²,
Yabiao Wang^{2†}, Chengjie Wang^{2†} and Haoqian Wang^{1†}

¹Shenzhen International Graduate School, Tsinghua University

²Tencent Youtu Lab

wang-hh20@mails.tsinghua.edu.cn, {leoneliu, wuhaozhang, vtzhang, wingzygan, caseywang, jasoncjwang}@tencent.com, wanghaoqian@tsinghua.edu.cn

Abstract

Few-shot semantic segmentation aims to learn to segment unseen class objects with the guidance of only a few support images. Most previous methods rely on the pixel-level label of support images. In this paper, we focus on a more challenging setting, in which only the image-level labels are available. We propose a general framework to firstly generate coarse masks from image label text with the help of the powerful vision-language model CLIP, and then refine the mask predictions of support and query images iteratively and mutually. During the refinement, we design an Iterative Mutual Refinement (IMR) module to adapt to the varying quality of the coarse support mask. Extensive experiments on PASCAL-5ⁱ and COCO-20ⁱ datasets demonstrate that our method not only outperforms the state-of-the-art weakly supervised approaches by a significant margin, but also achieves comparable or better results to recent supervised methods. Moreover, our method owns an excellent generalization ability for the images in the wild and uncommon classes. Code will be available at <https://github.com/Whileherham/IMR-HSNet>.

1 Introduction

In recent years, semantic segmentation has reached remarkable progress. However, a large amount of data is required for training such a task to segment unseen class objects. It differs from the human perception that a human is able to handle novel objects after learning from only a few samples.

Few-shot semantic segmentation is proposed to mimic the property of humans. The segmentation model is expected to segment objects with an unseen class in the query image with a few support images and pixel-wise annotations, as shown in Figure 1 (a). Although it is more like humans and helps reduce the work of data annotation, pixel-wise mask annotation is still unavoidable. In general, compared with image-level or box-level annotation, generating the precise pixel-level label is more time-consuming. Therefore, once faced with unseen

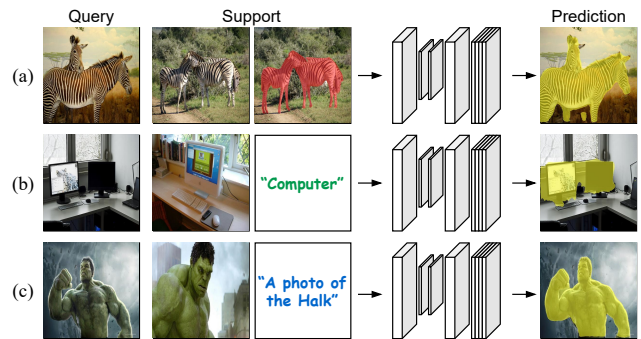


Figure 1: Illustration of (a) regular few-shot segmentation, (b) previous weakly supervised method, and (c) our method. We replace the simple class word used in previous work with a sentence of label text. Joint modeling vision and language contents and specially designed refining module enable our method to handle fine-grained novel classes with higher precision.

categories, it takes much time to prepare the dense annotations, which restricts the practical application in the real-time open world, such as robotic engineering or consumer electronics products.

In this work, we concentrate on the task of few-shot segmentation with image-level label support, namely weakly supervised few-shot semantic segmentation, as depicted in Figure 1 (b). We argue that pixel-wise examples are not necessary for pixel-wise recognition, as humans are able to imagine the shape and structure of objects with the help of prior knowledge such as attribute and low-level feature. Therefore, it is meaningful to study the weakly supervised setting.

However, due to the severer data scarcity, *i.e.*, the lack of precise support mask, weakly supervised few-shot semantic segmentation is rarely explored. [Raza *et al.*, 2019] simplify the problem into the condition where only one novel class exists in the support image, but the model would collapse in the more complex scene. Recent works employ the Word2Vec embedding to associate the image feature with the class label explicitly [Lee *et al.*, 2022] or implicitly [Siam *et al.*, 2020]. Nevertheless, the performance still falls far behind that in the standard few-shot setting.

We summarize that the performance gap is mainly because a model designed for weakly supervised few-shot segmentation is faced with two daunting challenges. **The first chal-**

*Equal contribution.

†Corresponding authors.

lence is to locate specific objects in the support image. Given a novel class label, the model has to filter out the corresponding objects of unseen categories from the background and other foreground instances in the support images. It always results in an imprecise or even wrong support mask, rather than the ground-truth mask. Therefore, *another challenge is to learn from imprecise labels.* In other words, the model has to adapt to the imprecise input support masks.

In this work, we propose a two-stage framework to address the challenges. In the first stage, we tackle the first challenge, *i.e.*, locating specific objects guided by the image label text. We fall back on the recent progress of vision-language model CLIP [Radford *et al.*, 2021], which is trained with massive image-text pairs from the web in a self-supervised way. CLIP encodes images and texts into the joint embedding space, and thus it builds a strong correlation between them. To leverage its potential, we inject the class name into a simple prompt template, as shown in Figure 1 (c), and compute the cosine similarity of the output from vision and text encoder in CLIP. This process could be regarded as the classification task, and hence we can generate the support mask according to the high-response regions of its Grad CAM [Selvaraju *et al.*, 2017].

Then in the second stage, we aim at handling the second challenge, *i.e.*, the imprecise input support mask. We experimentally find that due to the robustness of the well-trained few-shot segmentation model, when segmenting the query mask with an imprecise support mask, the quality of the predicted query mask is much higher than that of the input support mask. Based on this observation, we propose to refine the support and query mask iteratively and mutually. Specifically, we continue to refine the support mask with the help of a higher-quality predicted query mask, and then update the query mask with the better support mask in turn. However, due to the varying quality of the input object masks, the regular few-shot segmentation model would struggle in judging the confidence of the input mask, which leads to a suboptimal result. Hence we propose a progressive mask refinement module Iterative Mutual Refinement (IMR), to capture and adapt to the varying input distribution, *i.e.*, how well the input mask has been estimated. During the refinement, IMR provides an effective strategy to leverage the information of all past steps with memory, so that the support and query mask could be optimized adaptively. We incorporate the IMR into the pipeline of few-shot segmentation, so as to better explore the potential of the iterative refinement. Our contributions are summarized as follows:

- We propose a two-stage framework for weakly supervised few-shot segmentation, consisting of initial mask generation from the image label text and iterative support and query mask refinement.
- We present an Iterative Mutual Refinement module to progressively refine the coarse masks, which is better for the case where the input mask quality varies. It is designed to be model-agnostic, and could serve to extend common few-shot segmentation models to the weakly supervised setting with minor modifications.
- Extensive experiments show that our method could not

only outperform the state-of-the-art weakly supervised counterparts by a significant margin, but also close to or beyond recent regular few-shot segmentation methods.

2 Related Work

Few-shot Semantic Segmentation. The dominant approaches of few-shot segmentation are devised based on meta-learning. The information of support images is aggregated into the query feature for further prediction in different ways. Most methods achieve this by first extracting single [Tian *et al.*, 2020; Xie *et al.*, 2021] or multiple [Zhang *et al.*, 2021; Li *et al.*, 2021a] class prototypes from support feature by masked average pooling, then concatenating the prototypes and query feature as the input of a well-designed decoder. However, this simple strategy is faced with severe support information loss unavoidably. Some recent works alleviate this by exploring the potential of the pair-wise feature correlations between support and query feature, and the more powerful transformer architecture [Sun *et al.*, 2021] or 4D convolutions [Min *et al.*, 2021] is thus utilized to boost the performance further. Despite these signs of progress, few-shot segmentation requires the precise support mask during testing, which is superfluous and hardly available in customer-oriented real-time applications. In this work, we tackle this problem by replacing the ground-truth mask with the image label text.

Weakly Supervised Few-shot Segmentation. Weakly supervised few-shot segmentation is rarely explored due to the severer challenge of data scarcity. [Wang *et al.*, 2019; Zhang *et al.*, 2019] utilize bounding box as the annotation of support images to segment the query images during testing. However, their performance is suboptimal as these models are not specifically designed for the weakly supervised condition. The first few-shot segmentation method with image-level support annotation is proposed by [Raza *et al.*, 2019], where they generate the pseudo masks of support images by retaining the pixels not classified as background. Nevertheless, it cannot deal with the support images containing more than one novel class. Recently, [Siam *et al.*, 2020] incorporate the semantic word embedding of the novel category into the visual feature, and exploited a multi-stage attention mechanism to predict the query mask. [Lee *et al.*, 2022] regard the pseudo Classification Activation Map (CAM) of the unseen classes as the support mask and ground truth query mask during training, which is estimated by the CAMs of semantically similar categories. However, the performance of these approaches is far away from that using a precise support mask, which restricts its practical application. On the contrary, our method could perform close to or even beyond the state-of-the-art few-shot segmentation models, and show strong generalization ability for images in the wild.

CLIP in Classification and Segmentation. CLIP [Radford *et al.*, 2021] is a milestone of vision-language pretraining model, which learns a multi-modal embedding space by jointly encoding images and texts from the web in contrastive learning. CLIP shows strong capacity in zero-shot or few-shot recognition [Gao *et al.*, 2021; Radford *et al.*, 2021], and

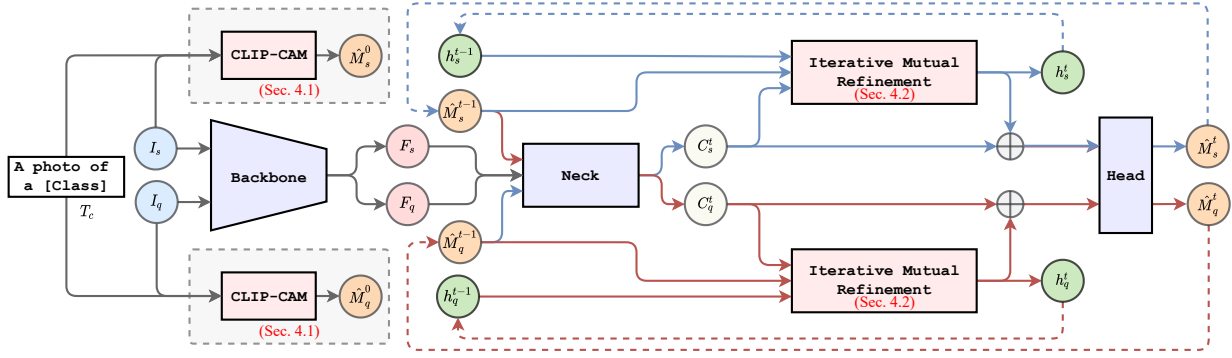


Figure 2: Pipeline of our weakly supervised few-shot segmentation method. Given a support image I_s and a query image I_q for a specific class c , we estimate coarse masks for support and query with the label prompt T_c as the initial masks firstly. Then, an iterative refinement with a symmetrical data flow is used to improve the segmentation quality gradually. Note that “Backbone”, “Neck”, and “Head” are abstract from the regular few-shot segmentation methods so that our method is able to integrate into most of the previous well-designed networks.

recent works have focused on the extension into dense prediction tasks like semantic segmentation. [Rao *et al.*, 2021] first employ CLIP in fully supervised semantic segmentation in a model-agnostic way, and [Zhou *et al.*, 2021] extend CLIP into zero-shot segmentation. [Li *et al.*, 2022] combine CLIP with a well-designed transformer to further improve the precision of segmentation. However, these methods rely heavily on complicated prompt engineering, and the upper bound of their segmentation performance is limited due to the lack of support images. Therefore, in this work, we explore the potential of CLIP in few-shot segmentation for the first time with a simple text prompting only, along with an Iterative Mutual Refinement module to obtain the precise prediction.

3 Problem Definition

Few-shot semantic segmentation aims at training a model which can segment unseen class objects given extremely limited annotated samples with the same class. Formally, the dominant episode based meta-learning paradigm employs a *base* dataset $\mathcal{D}_{\text{base}}$ with category set $\mathcal{C}_{\text{base}}$ for training and a *novel* dataset $\mathcal{D}_{\text{novel}}$ with unseen category set $\mathcal{C}_{\text{novel}}$ for testing. During training, a large number of episodes are sampled from $\mathcal{D}_{\text{base}}$. Each episode consists of a *support* set \mathcal{S} and a *query* set \mathcal{Q} for a specific class $c \in \mathcal{C}_{\text{base}}$ as follows:

$$\mathcal{S} = \{(I_s^k, M_s^k)\}_{k=1}^K, \quad \mathcal{Q} = \{(I_q, M_q)\}. \quad (1)$$

where I_* and M_* denote image and mask respectively. After training, the model $f(\cdot, \theta)$ learned to predict a binary mask \hat{M}_q for query image I_q with the guidance of K support images and masks for a specific novel class $c \in \mathcal{C}_{\text{novel}}$,

$$\hat{M}_q = f\left(\{(I_s^k, M_s^k)\}_{k=1}^K, I_q, \theta \mid c\right). \quad (2)$$

In the standard few-shot segmentation setting [Min *et al.*, 2021; Tian *et al.*, 2020], both support masks M_s and query masks M_q for class c are available for training, while query masks are inaccessible for testing. On the contrary, we focus on a more challenging setting, named weakly supervised

few-shot segmentation, in which even support masks are inaccessible in testing. Thus, the prediction becomes:

$$\hat{M}_q = f\left(\{I_s^k\}_{k=1}^K, I_q, \theta \mid c\right). \quad (3)$$

Although a weaker annotation replaces the pixel-wise annotated mask, we show that it can still reach a comparable or even better performance.

4 Method

We propose a framework for weakly supervised few-shot segmentation. The pipeline is shown in Figure 2, which contains a CLIP-CAM module to generate coarse masks for support and query images (Section 4.1), an Iterative Mutual Refinement module to refine the coarse masks progressively (Section 4.2), and some basic blocks from previous well-designed models. Notably, our framework is plug-and-play and model-agnostic, and we show how to extend regular few-shot segmentation models to the weakly supervised setting in Section 4.3.

4.1 Coarse Mask from Label Text

In the weakly supervised setting, dense annotations of supported images are not available, so we propose to generate segmentation maps as an alternative from images and class labels only. To some extent, this subtask degrades into zero-shot segmentation, for which the goal comes to segment novel objects in an image with the class label only. Nevertheless, most existing zero-shot approaches require complex architectures and training strategies, which might be inefficient to serve as only one stage of our whole pipeline.

Recent vision-language models, especially CLIP [Radford *et al.*, 2021], have shown surprisingly powerful performance in associating visual and text concepts. CLIP consists of a vision encoder $E_v(\cdot)$ and a text encoder $E_t(\cdot)$, where the image and text are encoded into $f_v, f_t \in \mathbb{R}^{1 \times 1 \times C}$. Since it learns a joint embedding space for images and texts, one can simply compute the cosine similarity of f_v and f_t for image-level zero-shot classification without extra training and modules.

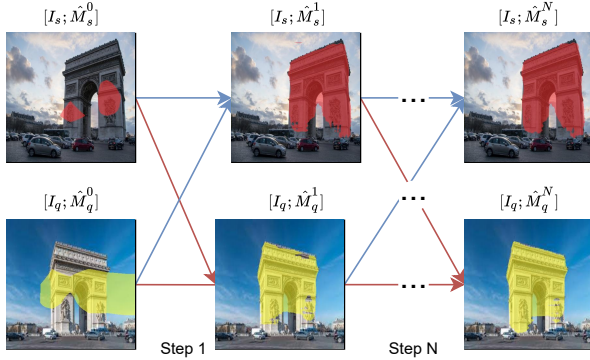


Figure 3: Illustration of the iterative mutual mask refinement. In the step t , the posterior query mask estimation \hat{M}_q^t is refined with the guidance of corresponding support mask \hat{M}_s^{t-1} and the prior estimation \hat{M}_q^{t-1} . The posterior support mask estimation \hat{M}_s^t is simultaneously updated in a symmetric way. During the refinement, the quality of the object masks is improving gradually.

In this work, we extend it to pixel-level zero-shot segmentation in a training-free way.

The original architecture of CLIP is not suitable for dense prediction tasks, so a naive idea is to adjust the architecture of the vision encoder $E_v(\cdot)$. Specifically, similar to [Zhou *et al.*, 2021], we drop the *query* and *key* embedding layers, and reformulate the *value* embedding layer and the last linear layer into two convolutional layers, so as to project the image feature of each position to the embedding space of the text feature. After that, we can conduct pixel-level prediction by computing the cosine similarity between text feature and dense visual feature.

An alternative is the heatmap of Grad CAM [Selvaraju *et al.*, 2017], which aims at looking for the region contributing most to the classification. Since CLIP computes the cosine distance between f_v and f_t , it could be regarded as the classification output of a certain category, where f_t works as the weight of the classifier. Thus we utilize the heatmap output by Grad CAM as the coarse object mask, as the region with a high response makes f_v resemble f_t .

With the guidance of the class label text, both strategies above could generate the coarse object masks from the images. We select the latter one due to its better performance. Please refer to the supplementary materials for more details of the implementation and comparison. It is worth mentioning that it is unnecessary to pursue a precise initial mask, since we adopt an iterative strategy to refine the estimation of the mask. Therefore, we do not depend on the complicated prompt engineering and other various zero-shot methods, and still achieve a satisfying performance.

4.2 Iterative Mutual Mask Refinement

With a coarse support mask \hat{M}_s^0 , an intuitive solution is to feed it into the few-shot segmentation model to obtain a more delicate query mask. Similarly, the finer query mask could then result in the finer support mask, which is more suitable as the indeed support input. Therefore, an iterative strategy could improve the final prediction gradually by support-query

mutual refinement, as illustrated in Figure 3.

Nevertheless, the distribution of input masks from different steps varies, which degrades the performance of the model. Therefore, we propose an Iterative Mutual Refinement (IMR) module to capture and adapt to various input distributions.

Specifically, we first introduce a hidden state h to reflect how well the input mask is estimated. Thus in the t step, we already have the image feature F_s, F_q , the prior estimation of support and query masks $\hat{M}_s^{t-1}, \hat{M}_q^{t-1}$, as well as the stored previous state h_s^{t-1}, h_q^{t-1} . In order to predict \hat{M}_q^t , we first encode the information of support images and masks into the query feature, and obtain the correlation volume C_q^t , which is consistent with the regular few-shot segmentation. Then $F_q, C_q^t, \hat{M}_q^{t-1}$ are concatenated as x_q^t , along with h_q^{t-1} as the input to update the h_q^t . The whole process of state updating derived from the gated recurrent unit is formulated as:

$$\begin{aligned} z_q^t &= \sigma(W_{\text{conv}1,1} \cdot [x_q^t, h_q^{t-1}]), \\ r_q^t &= \sigma(W_{\text{conv}1,2} \cdot [x_q^t, h_q^{t-1}]), \\ \tilde{h}_q^t &= \tanh(W_{\text{conv}2} \cdot [r_q^t \otimes h_q^{t-1}, x_q^t]), \\ h_q^t &= (1 - z_q^t) \otimes h_q^{t-1} + z_q^t \otimes \tilde{h}_q^t. \end{aligned} \quad (4)$$

where $W_{\text{conv}1,1}, W_{\text{conv}1,2}, W_{\text{conv}2}$ are implemented by 1×1 convolutions followed by 3×3 group convolutions with 16 groups, and $[\cdot], \otimes, \sigma$ denote concatenation, Hadamard product and sigmoid function, respectively. For brevity, we only list how h_q^t is updated, and h_s^t is updated in the symmetrical way. After that, the updated h_q^t , as well as the prior estimation \hat{M}_q^{t-1} , are utilized for further prediction. Detailed analysis is provided in the supplementary materials.

4.3 Model Architecture and Objective

Our main pipeline consists of the Iterative Mutual Refinement module, backbone, neck, and head. The last three modules are abstract from the regular few-shot segmentation methods according to their respective roles, and could be replaced with corresponding modules in different basic methods.

In this work we reuse the architecture of HSNet [Min *et al.*, 2021] for its excellent performance. Concretely, we adopt ResNet50 [He *et al.*, 2016] or VGG16 [Simonyan and Zisserman, 2015] as backbone to extract feature F_s, F_q . Then we adopt the well-designed 4D convolutions followed by multi-scale feature fusion as the neck to model the correlation of F_s and F_q , and obtain the intermediate output C_q . Please refer to [Min *et al.*, 2021] for more details.

For the mask prediction, we make a minor modification based on the original HSNet head, in order to leverage the prior mask estimation \hat{M}^{t-1} , which is formulated as follows:

$$\tilde{M}^t = \text{Head}_0(C^t + h^t), \quad (5)$$

$$\hat{M}^t = \text{Softmax} \left(W_{h2} \cdot \text{GeLU}(W_{h1} \cdot [\tilde{M}^t, \hat{M}^{t-1}]) \right). \quad (6)$$

where Head_0 is the original head in HSNet, and W_{h1}, W_{h2} is implemented by 1×1 convolutional layers.

In order to accelerate convergence, we add the intermediate supervision to the output mask of each step. As a result, the

Backbone	Method	A. Type	Pascal-5 ⁱ				
			5 ⁰	5 ¹	5 ²	5 ³	Mean
VGG16	RPMM [Yang <i>et al.</i> , 2020]	Pixel	47.1	65.8	50.6	48.5	53.0
	ASR [Liu <i>et al.</i> , 2021]	Pixel	49.2	65.4	52.6	51.3	54.6
	SS-PANet [Li <i>et al.</i> , 2021b]	Pixel	49.3	60.8	53.9	45.2	52.3
	SS-PFENet [Li <i>et al.</i> , 2021b]	Pixel	54.5	67.4	63.4	54.0	59.8
	HSNet [Min <i>et al.</i> , 2021]	Pixel	59.6	65.7	59.6	54.0	59.7
	PANet [Wang <i>et al.</i> , 2019]	Box	–	–	–	–	45.1
Ours (IMR-HSNet)	Image	36.5	51.7	45.9	35.6	42.4	
	Image	58.2	63.9	52.9	51.2	56.5	
	Image	58.2	63.9	52.9	51.2	56.5	
ResNet50	PFENet [Tian <i>et al.</i> , 2020]	Pixel	61.7	69.5	55.4	56.3	60.8
	ASR [Liu <i>et al.</i> , 2021]	Pixel	53.8	69.6	51.6	52.8	56.9
	ASGNet [Li <i>et al.</i> , 2021a]	Pixel	58.8	67.9	56.8	53.7	59.3
	SCL [Zhang <i>et al.</i> , 2021]	Pixel	63.0	70.0	56.5	57.7	61.8
	HSNet [Min <i>et al.</i> , 2021]	Pixel	64.3	70.7	60.3	60.5	64.0
	CANet [Zhang <i>et al.</i> , 2019]	Box	–	–	–	–	52.0
	[Siam <i>et al.</i> , 2020]+CANet	Image	49.5	65.5	50.0	49.2	53.5
	[Siam <i>et al.</i> , 2020]+PANet	Image	42.5	64.8	48.1	46.5	50.5
	Ours (IMR-HSNet)	Image	62.6	69.1	56.1	56.7	61.1
	Ours (IMR-HSNet)	Image	62.6	69.1	56.1	56.7	61.1

Table 1: Comparisons with regular and weakly supervised few-shot semantic segmentation methods on Pascal-5ⁱ.

overall loss function is defined as follows:

$$\mathcal{L}_{\text{all}} = \sum_{t=1}^N \lambda_t \cdot [\mathcal{L}_{\text{seg}}(\hat{M}_q^t, M_q) + w_s \cdot \mathcal{L}_{\text{seg}}(\hat{M}_s^t, M_s)]. \quad (7)$$

where N denotes the number of the total steps in the refinement, and M_s, M_q are the ground-truth support and query mask. We adopt cross entropy loss \mathcal{L}_{seg} during training and λ_t and w_s are the hyperparameters.

5 Experimental Results

5.1 Datasets and Evaluation Metrics

We follow almost the same settings with regular few-shot segmentation, and the only difference is that the ground-truth masks of support images are replaced by the label text. We evaluate our framework in two widely-used datasets, *i.e.*, Pascal-5ⁱ [Shaban *et al.*, 2017], COCO-20ⁱ [Lin *et al.*, 2014], and conduct the cross-validation over all the folds in each dataset. For each fold i , samples in the remaining folds serve as the training data, while the 1000 episodes (support-query pairs) are sampled randomly as the testing data. We adopt the mean intersection over union (mIoU) to measure the performance, which averages the IoU of all classes in the testing data to avoid the bias to the specific categories.

5.2 Implementation Details

We iteratively refine the estimation of support and query masks twice in Pascal-5ⁱ, and thrice in COCO-20ⁱ due to the more complex scene. It is actually a trade-off between effectiveness and efficiency, as we empirically find that the gain of performance gradually converges with the step number increasing. For a fair comparison, we utilize the ImageNet pre-trained VGG16 and ResNet50 as our backbone to extract feature, and the pre-trained CLIP model for initial mask generation. For the CLIP model, the vision and text encoder are a modified ResNet50 and transformer, respectively. The parameters of the backbone and CLIP are fixed during training to avoid the undesirable representation bias towards the base class.

Backbone	Method	A. Type	COCO-20 ⁱ				
			20 ⁰	20 ¹	20 ²	20 ³	Mean
VGG16	PFENet [Tian <i>et al.</i> , 2020]	Pixel	35.4	38.1	36.8	34.7	36.3
	SAGNN [Xie <i>et al.</i> , 2021]	Pixel	35.0	40.5	37.6	36.0	37.3
	SS-PANet [Li <i>et al.</i> , 2021b]	Pixel	29.8	21.2	26.5	28.5	26.2
	SS-PFENet [Li <i>et al.</i> , 2021b]	Pixel	35.6	39.2	37.6	37.3	37.5
	PANet [Wang <i>et al.</i> , 2019]	Box	12.7	8.7	5.9	4.8	8.0
Ours (IMR-HSNet)	Image	24.2	12.9	17.0	14.0	17.0	
	Image	34.9	38.8	37.0	40.1	37.7	
	Image	34.9	38.8	37.0	40.1	37.7	
ResNet50	ASR [Liu <i>et al.</i> , 2021]	Pixel	29.9	35.0	31.9	33.5	32.6
	RePRI [Boudiaf <i>et al.</i> , 2021]	Pixel	31.2	38.1	33.3	33.0	34.0
	ASGNet [Li <i>et al.</i> , 2021a]	Pixel	–	–	–	–	34.6
	[Yang <i>et al.</i> , 2021]	Pixel	46.8	35.3	26.2	27.1	33.9
	[Wu <i>et al.</i> , 2021]	Pixel	34.9	41.0	37.2	37.0	37.5
	CWT [Lu <i>et al.</i> , 2021]	Pixel	32.2	36.0	31.6	31.6	32.9
	HSNet [Min <i>et al.</i> , 2021]	Pixel	36.3	43.1	38.7	38.7	39.2
	[Siam <i>et al.</i> , 2020]	Image	–	–	–	–	15.0
	Ours (IMR-HSNet)	Image	39.5	43.8	42.4	44.1	42.4

Table 2: Comparisons with regular and weakly supervised few-shot semantic segmentation methods on COCO-20ⁱ.

Owing to our refinement in the second stage, we adopt the most superficial prompt text “a photo of a [class]” for the input to the text encoder, where [class] is replaced by the label of specific categories.

Consistent with HSNet, we resize the input into 400×400 in both the training and testing stage, and we do not adopt any data augmentation strategies. The network is trained with Adam optimizer with learning rate $2e^{-4}$, and all hyperparameters λ_t and w_s are set to one by default. However, our method is robust to the choice of these hyperparameters, as shown in the supplementary materials.

5.3 Comparison with State-of-the-art

In this subsection, we compare our method with the state-of-the-art weakly supervised and regular few-shot segmentation models in the 1-shot setting. Table 1 presents the results in the Pascal-5ⁱ. It could be seen that our method sets new state of the arts with both VGG16 and ResNet50 backbone in the weakly supervised setting, and outperforms its counterparts with either box-level or image-level annotation, by a large margin (11% with VGG16 and 7.6% with ResNet50). Furthermore, benefiting from our iterative strategy, our weakly supervised method surpasses most regular few-shot models.

Table 2 summarizes the results in the more challenge benchmark COCO-20ⁱ. When more objects with various scales appear in the query image, the performance of other weakly supervised models falls much more behind that of fully supervised models. However, our method turns out to be robust to this complex scene, and thus accordingly, the superiority of our method is more significant. Specifically, we achieve more than larger performance gains (20% and 27% for VGG16 and ResNet) compared to the state-of-the-art ones using weak annotation. Furthermore, adopting ResNet50 as the backbone, our method in the weakly supervised setting achieves 42.4 mIoU, which even outperforms the state-of-the-art model in the standard setting.

In order to show the comparison more intuitively, we visualize the results of some of the challenging query images, which are segmented by state-of-the-art HSNet (standard few-shot setting) and our method (weakly supervised setting). As shown in Figure 4, our method could perform

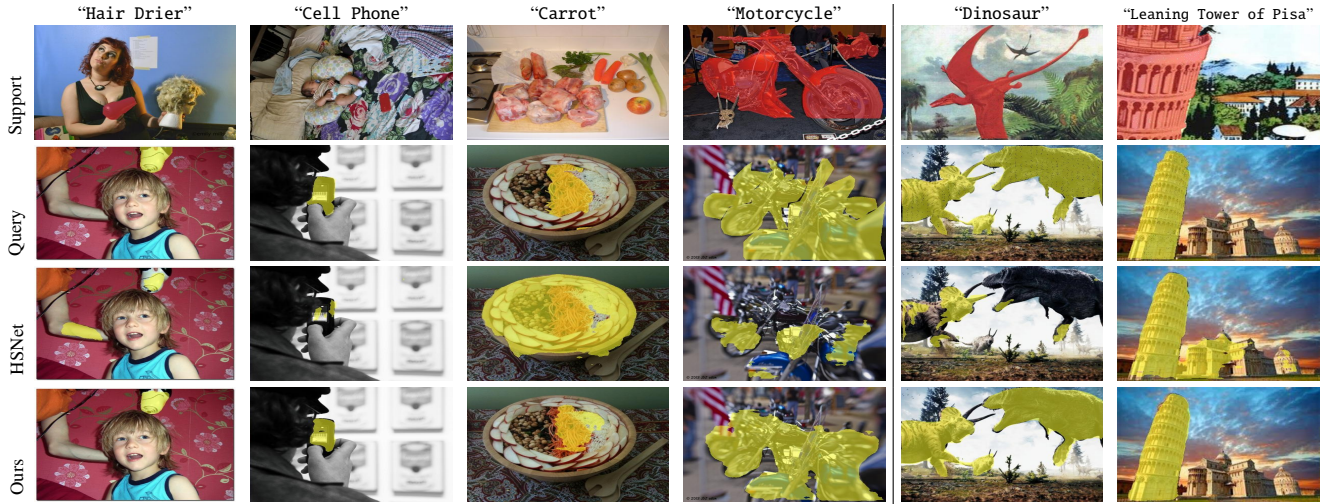


Figure 4: Qualitative visualization comparing with the regular few-shot segmentation method HSNNet [Min *et al.*, 2021]. The first four columns are from the COCO-20ⁱ, and the last two columns are from web to show the performance in-the-wild.

Method	Shot	Backbone	20 ⁰	20 ¹	20 ²	20 ³	Mean
ZS3Net [Bucher <i>et al.</i> , 2019]	Zero	Res101	18.8	20.1	24.8	20.5	21.1
L-Seg [Li <i>et al.</i> , 2022]	Zero	Res101	22.1	25.1	24.9	21.5	23.4
L-Seg [Li <i>et al.</i> , 2022]	Zero	ViT-L/16	28.1	27.5	30.0	23.2	27.2
Ours	Zero	Res50	37.9	41.8	41.3	43.3	41.1
Ours	One	Res50	39.5	43.8	42.4	44.1	42.4
Ours	Five	Res50	40.7	46.0	45.0	46.0	44.4

Table 3: Comparisons with zero-shot semantic segmentation methods on COCO-20ⁱ.

Methods	Shot	5 ⁰	5 ¹	5 ²	5 ³	Mean
Cosine Sim	Zero	21.4	39.9	27.9	32.1	30.3
CLIP-CAM	Zero	24.3	27.6	29.7	29.0	27.6
Cosine Sim	One	38.9	56.2	39.1	39.5	43.4
CLIP-CAM	One	50.2	58.6	49.0	47.5	51.3

Table 4: Comparisons of initial mask generation methods.

better in segmenting tiny objects (1st and 2nd columns), semantically confusing objects (3rd column), multiple objects (4th column), and objects in the wild (last two columns). This is probably because our method is able to associate support and query objects with a large difference in appearance with the help of label text, and more likely to predict the complete masks with complex shapes thanks to the prior mask estimation in each step of the refinement.

5.4 Comparison with Zero-shot Segmentation

Weakly supervised few-shot segmentation requires support label text rather than the support mask to segment the query image. From this perspective, it has a lot in common with zero-shot semantic segmentation (ZSS), where only the label text without the support image is available. Therefore, it is natural to compare our method with zero-shot counterparts. However, the comparison is non-trivial due to the inconsistent testing protocol between zero-shot and few-shot segmentation as well as the lack of zero-shot datasets and baselines.

Methods	5 ⁰	5 ¹	5 ²	5 ³	Mean
+HSNet (pretrained)	50.2	58.6	49.0	47.5	51.3
+HSNet (mutual)	59.5	67.5	54.1	53.8	58.7
+HSNet (retrained)	60.9	65.8	53.1	54.7	58.6
+HSNet (retrained+mutual)	56.2	63.8	47.0	49.6	54.2
Ours (IMR-HSNet)	62.6	69.1	56.1	56.7	61.1

Table 5: Comparisons of mask refinement methods.

Methods	5 ⁰	5 ¹	5 ²	5 ³	Mean
CLIP-CAM	24.3	27.6	29.7	29.0	27.6
CLIP-CAM / PFENet	44.2	51.1	49.0	47.9	48.0
CLIP-CAM / HSNNet	50.2	58.6	49.0	47.5	51.3
Ours (IMR-PFENet)	53.4	53.7	54.2	53.3	53.7
Ours (IMR-HSNet)	62.6	69.1	56.1	56.7	61.1

Table 6: Generalization ability with other few-shot segmentation architectures.

Thus, we still adopt the few-shot benchmark, *i.e.*, Pascal-5ⁱ and COCO-20ⁱ for evaluation.

In order to apply our framework to the zero-shot segmentation, we need to make some modifications. Therefore, we take the query image as the input of where the support image is required. Concretely, we regard $\{I_q, \hat{M}_q^{t-1}\}$ as the $\{I_s, \hat{M}_s^{t-1}\}$ to predict \hat{M}_q^t , and \hat{M}_s^t is equal to \hat{M}_q^t , then repeat the refinement to obtain final \hat{M}_q^N . In this way, the query image is segmented with the guidance by itself. We compare our method with competitive zero-shot methods, *e.g.* ZS3Net [Bucher *et al.*, 2019], and L-Seg [Li *et al.*, 2022].

As reported in Table 3, our method achieves superior performance in the zero-shot setting with a mean mIoU score of 41.1. Without being designed and trained specifically for zero-shot segmentation and strong backbone ResNet101, our method could outperform the recent zero-shot approaches by a significant margin. However, the zero-shot performance

still falls behind that of the few-shot version, which proves the importance of the support image, even without the precise object mask.

5.5 Ablation Study

We conduct the ablation study to investigate the effects of components in our framework. We adopt the ResNet50 as our backbone, and evaluate in the Pascal-5^t dataset.

Initial Mask Generation. In Section 4.1, we introduce two kinds of ways to generate the initial support and query mask \hat{M}_s^0, \hat{M}_q^0 , *i.e.*, the pixel-level cosine similarity between text and dense visual feature, and the Grad CAM of similarity between text and global visual feature. We further compare the two strategies in the zero-shot and few-shot settings. For zero-shot segmentation, we adopt the generated query mask \hat{M}_q^0 as the predicted query mask. And for few-shot segmentation, we feed the generated support mask \hat{M}_s^0 into the pre-trained HSNet as the ground-truth support mask to segment the query images.

Table 4 reports the results of the comparison. From the first two rows (zero-shot setting), it can be seen that the quality of \hat{M}^0 is similar. Nevertheless, when using the generated mask as the support mask, the results of Grad CAM show clear superiority with nearly 10%. The reason could be that the Grad CAM better leverages the architecture and weights of the pre-trained vision encoder, and thus reflects the strong generality. As a result, we utilize this strategy to generate the initial support and query mask estimation.

Iterative Mutual Refinement. We further compare different iterative strategies in Table 5 to evaluate the effectiveness of IMR. We set the method without Iterative Mutual Refinement as the baseline, which is implemented by feeding \hat{M}_s^0 to the pretrained HSNet, denoted as *+HSNet (pretrained)*. We start with evaluating a simple iterative strategy, where we first predict \hat{M}_q^1 with \hat{M}_s^0 , then update \hat{M}_s^2 with \hat{M}_q^1 , and get the final prediction \hat{M}_q^3 guided by \hat{M}_s^2 , denoted as *+HSNet (mutual)*. Simultaneously, in order to adapt to the condition where the clean ground-truth support mask is unavailable, we trained HSNet with coarse support masks generated by CLIP-CAM. Similarly, we feed \hat{M}_s^0 to the retrained HSNet (*+HSNet (retrained)*), and refine the estimation of support and query mask iteratively (*+HSNet (retrained + mutual)*).

As shown in Table 5, both iterative strategies outperform the simple baseline. Therefore, iterative refining mask estimation may be the key to weakly supervised few-shot segmentation. Moreover, our method with IMR still shows great superiority due to the architecture designed with hidden states to fit the varying input distribution during the refinement.

Generalization of Our Method. As mentioned above, our framework is plug-and-play and model-agnostic. In order to verify this, we evaluate our framework with another few-shot segmentation model PFENet [Tian *et al.*, 2020]. Concretely, we adopt the last convolutional layers of PFENet as the *head* in our method, and the other modules, including the prior generation and feature enrichment module (FEM), as the *neck*. We keep all the hyperparameters the same with the origin implementation of PFENet. We compare our method built on

PFENet with corresponding baselines, where we feed the \hat{M}_s^0 as the support mask to the pre-trained PFENet.

As shown in Table 6, with the well-designed iterative strategy, our method outperforms the baseline by around 6%. Hence, it could be concluded that our method is robust to the selection of the few-shot segmentation model.

6 Conclusion

In this work, we present a general framework to tackle the weakly supervised few-shot segmentation, which contains coarse masks generation from label text and iteratively and mutually refine support and query mask. Extensive experiments show that our method outperforms the state-of-the-art counterparts by a large margin. One limitation is that although our method is able to refine segmentation from coarse masks, it is unclear whether finer initial masks from prompt engineering or post-process such as dense CRF could help performance. We leave it for future work.

Acknowledgements

This work is partially supported by the NSFC fund (61831014), in part by the Shenzhen Science and Technology Project under Grant (ZDYBH201900000002, JCYJ20180508152042002, CJGJZD20200617102601004)

References

- [Boudiaf *et al.*, 2021] Malik Boudiaf, Hoel Kervadec, Ziko Imtiaz Masud, Pablo Piantanida, Ismail Ben Ayed, and Jose Dolz. Few-shot segmentation without meta-learning: A good transductive inference is all you need? In *CVPR*, 2021.
- [Bucher *et al.*, 2019] Maxime Bucher, Tuan-Hung Vu, Matthieu Cord, and Patrick Pérez. Zero-shot semantic segmentation. *NeurIPS*, 2019.
- [Gao *et al.*, 2021] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *arXiv:2110.04544*, 2021.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [Lee *et al.*, 2022] Yuan-Hao Lee, Fu-En Yang, and Yu-Chiang Frank Wang. A pixel-level meta-learner for weakly supervised few-shot semantic segmentation. In *WACV*, 2022.
- [Li *et al.*, 2021a] Gen Li, Varun Jampani, Laura Sevilla-Lara, Deqing Sun, Jonghyun Kim, and Joongkyu Kim. Adaptive prototype learning and allocation for few-shot segmentation. In *CVPR*, 2021.
- [Li *et al.*, 2021b] Yiwen Li, Gratianus Wesley Putra Data, Yunguan Fu, Yipeng Hu, and Victor Adrian Prisacariu. Few-shot semantic segmentation with self-supervision from pseudo-classes. *arXiv:2110.11742*, 2021.
- [Li *et al.*, 2022] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and Rene Ranftl. Language-driven semantic segmentation. In *ICLR*, 2022.

- [Lin *et al.*, 2014] 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 *ECCV*, 2014.
- [Liu *et al.*, 2021] Binghao Liu, Yao Ding, Jianbin Jiao, Xiangyang Ji, and Qixiang Ye. Anti-aliasing semantic reconstruction for few-shot semantic segmentation. In *CVPR*, 2021.
- [Lu *et al.*, 2021] Zhihe Lu, Sen He, Xiatian Zhu, Li Zhang, Yi-Zhe Song, and Tao Xiang. Simpler is better: Few-shot semantic segmentation with classifier weight transformer. In *ICCV*, 2021.
- [Min *et al.*, 2021] Juhong Min, Dahyun Kang, and Minsu Cho. Hypercorrelation squeeze for few-shot segmentation. In *ICCV*, 2021.
- [Radford *et al.*, 2021] 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 *ICML*, 2021.
- [Rao *et al.*, 2021] Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. Denseclip: Language-guided dense prediction with context-aware prompting. *arXiv:2112.01518*, 2021.
- [Raza *et al.*, 2019] Hasnain Raza, Mahdyar Ravanbakhsh, Tassilo Klein, and Moin Nabi. Weakly supervised one shot segmentation. In *ICCVW*, 2019.
- [Selvaraju *et al.*, 2017] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, 2017.
- [Shaban *et al.*, 2017] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots. One-shot learning for semantic segmentation. In *BMVC*, 2017.
- [Siam *et al.*, 2020] Mennatullah Siam, Naren Doraiswamy, Boris N Oreshkin, Hengshuai Yao, and Martin Jagersand. Weakly supervised few-shot object segmentation with visual and semantic embeddings. In *IJCAI*, 2020.
- [Simonyan and Zisserman, 2015] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [Sun *et al.*, 2021] Guolei Sun, Yun Liu, Jingyun Liang, and Luc Van Gool. Boosting few-shot semantic segmentation with transformers. *arXiv:2108.02266*, 2021.
- [Tian *et al.*, 2020] Zhuotao Tian, Hengshuang Zhao, Michelle Shu, Zhicheng Yang, Ruiyu Li, and Jiaya Jia. Prior guided feature enrichment network for few-shot segmentation. *TPAMI*, 2020.
- [Wang *et al.*, 2019] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *ICCV*, 2019.
- [Wu *et al.*, 2021] Zhonghua Wu, Xiangxi Shi, Guosheng Lin, and Jianfei Cai. Learning meta-class memory for few-shot semantic segmentation. In *ICCV*, 2021.
- [Xie *et al.*, 2021] Guo-Sen Xie, Jie Liu, Huan Xiong, and Ling Shao. Scale-aware graph neural network for few-shot semantic segmentation. In *CVPR*, 2021.
- [Yang *et al.*, 2020] Boyu Yang, Chang Liu, Bohao Li, Jianbin Jiao, and Qixiang Ye. Prototype mixture models for few-shot semantic segmentation. In *ECCV*, 2020.
- [Yang *et al.*, 2021] Lihe Yang, Wei Zhuo, Lei Qi, Yinghuan Shi, and Yang Gao. Mining latent classes for few-shot segmentation. In *ICCV*, 2021.
- [Zhang *et al.*, 2019] Chi Zhang, Guosheng Lin, Fayao Liu, Rui Yao, and Chunhua Shen. Canet: Class-agnostic segmentation networks with iterative refinement and attentive few-shot learning. In *CVPR*, 2019.
- [Zhang *et al.*, 2021] Bingfeng Zhang, Jimin Xiao, and Terry Qin. Self-guided and cross-guided learning for few-shot segmentation. In *CVPR*, 2021.
- [Zhou *et al.*, 2021] Chong Zhou, Chen Change Loy, and Bo Dai. Denseclip: Extract free dense labels from clip. *arXiv:2112.01071*, 2021.