

# Self-Supervised Tuning for Few-Shot Segmentation

Kai Zhu, Wei Zhai and Yang Cao\*

University of Science and Technology of China

{zkzy, wzhai056}@mail.ustc.edu.cn, forrest@ustc.edu.cn

## Abstract

Few-shot segmentation aims at assigning a category label to each image pixel with few annotated samples. It is a challenging task since the dense prediction can only be achieved under the guidance of latent features defined by sparse annotations. Existing meta-learning method tends to fail in generating category-specificly discriminative descriptor when the visual features extracted from support images are marginalized in embedding space. To address this issue, this paper presents an adaptive tuning framework, in which the distribution of latent features across different episodes is dynamically adjusted based on a self-segmentation scheme, augmenting category-specific descriptors for label prediction. Specifically, a novel self-supervised inner-loop is firstly devised as the base learner to extract the underlying semantic features from the support image. Then, gradient maps are calculated by back-propagating self-supervised loss through the obtained features, and leveraged as guidance for augmenting the corresponding elements in embedding space. Finally, with the ability to continuously learn from different episodes, an optimization-based meta-learner is adopted as outer loop of our proposed framework to gradually refine the segmentation results. Extensive experiments on benchmark PASCAL-5<sup>i</sup> and COCO-20<sup>i</sup> datasets demonstrate the superiority of our proposed method over state-of-the-art.

## 1 Introduction

Recently, semantic segmentation models [Long *et al.*, 2015] have made great progress under full supervision, and some of them have even surpassed the level of human recognition. However, when the learned model is applied to a new segmentation task, it takes great cost to collect a large amount of full annotated data in pixel level. Furthermore, samples are not available in large quantities in some areas such as health care, security and so on. To address this problem, various few-shot segmentation methods are proposed.

\*Corresponding author




Figure 1: Comparison of prediction results with and without the self-supervised tuning framework. The self-supervised tuning process provides the category-specific semantic constraint to facilitate the features of the person and the bottle more discriminative, thereby improving the few-shot segmentation performance of corresponding categories. (a) Self-segmentation results. The support set acts as the supervision to segment the support image itself. Before tuning, the person is incorrectly identified as the bottle even in self-segmentation case. (b) Common one-shot segmentation. The query and support images are different objects belonging to the same category. After tuning by the self-supervised branch, the bottle regions are distinguished from the person.

One solution for solving few-shot segmentation [Shaban *et al.*, 2017] is meta-learning [Munkhdalai and Yu, 2017], whose general idea is to utilize a large number of episodes similar to target task to learn a meta learner that generate an initial segmentation model, and a base learner that quickly tunes the model with few samples. In most methods, a powerful feature extractor with good migration ability is provided for the meta learner to map the query and support images into a shared embedding space. And the base learner generates a category-specific descriptor with the support set. The similarity between the query’s feature maps and the descriptor is measured under a parametric or nonparametric metric, and leveraged as guidance for dense prediction of the query branch.

However, when the low-level visual features of foreground objects extracted from the support images are too marginalized in the embedding space [Zhang *et al.*, 2014], the generated descriptor by base learner is not category-specificly discriminative. In this case, the regions to be segmented in the query image may be ignored or even confused with other categories in the background. For example, as shown in Fig. 1, the bottle-specific descriptor generated from the support

image is unable to identify the bottle itself in the same image, and therefore inapplicable for bottles in other complicated scenarios to be segmented.

To address the issue, this paper presents an adaptive tuning framework for few-shot segmentation, in which the marginalized distributions of latent features are dynamically adjusted by a self-supervised scheme. The core of the proposed framework is the base learner driven by a self-segmentation task, which augments category-specific constraint on each new episode and facilitate successive labeling. Specifically, the base learner is designed as an inner-loop task, where the support images are segmented under the supervision of the presented masks of input support images. By back-propagating self-supervised loss through the support feature map, the corresponding gradients are calculated and used as guidance for augmenting each element in the support embedding space.

Moreover, the self-segmentation task can be considered as a special case of few-shot segmentation with the query and support image to be the same. Therefore, we also utilize the resulting loss of the task (we call it auxiliary loss in this paper) to promote the training process. Since auxiliary loss is equivalent to performing data enhancement to annotated samples, it can involve more information for training with the same number of iterations. The evaluation is performed on the prevailing public benchmark datasets PASCAL-5<sup>i</sup> and COCO-20<sup>i</sup>, and the results demonstrate above-par segmentation performance and generalization ability of our proposed method.

Our main contributions are summarized as follows:

1. An adaptive tuning framework is proposed for few-shot segmentation, in which the marginalized distributions of latent category features are dynamically adjusted by a self-supervised scheme.

2. A novel base learner driven by a self-segmentation task is proposed, which augments category-specific feature description on each new episode, resulting in better performance on label prediction.

3. Experimental results on two public benchmark datasets PASCAL-5<sup>i</sup> and COCO-20<sup>i</sup> demonstrate the superiority of our proposed method over SOTA.

## 2 Related Work

**Few-shot learning.** Few-shot learning has recently received a lot of attention and substantial progress has been made based on meta-learning. Generally, these methods can be divided into three categories. Metric-based methods focus on the similarity metric function over the embeddings [Snell *et al.*, 2017]. Model-based method mainly utilizes the internal architecture of the network (such as memory module [Santoro *et al.*, 2016], etc.) to realize the rapid parameter adaptation in new categories. The optimization-based method aims at learning an update scheme for base learner [Munkhdalai and Yu, 2017] in each episode. In the latest study, [Lee *et al.*, 2019] and [Bertinetto *et al.*, 2018] introduce machine learning methods such as SVM and ridge regression into the inner loop of the base learner, and [Rusu *et al.*, 2018] directly replaces the inner loop with an encoded-decode network. These methods have achieved state-of-the-art performance in few-shot classification task. Our model also takes inspiration of

them.

**Semantic segmentation.** Semantic segmentation is an important task in computer vision and FCNs [Long *et al.*, 2015] [Zhang *et al.*, 2019b] have greatly promoted the development of the field. After that, DeepLabV3 [Chen *et al.*, 2017] and PSPNet [Zhao *et al.*, 2017] propose different global contextual modules, which pay more attention to the scale change and global information [Chen *et al.*, 2019] in the segmentation process. [Zhu *et al.*, 2019] considers the full-image dependencies from all pixels based on Non-local Networks, which shows superior performance in terms of reasoning [Qiao *et al.*, 2019].

**Few-shot semantic segmentation.** While the work on few-shot learning is quite extensive, the research on few-shot segmentation [Hu *et al.*, 2019] has been presented only recently. [Shaban *et al.*, 2017] first proposes the definition and task of one-shot segmentation. Following this, [Rakelly *et al.*, 2018] solves the problem with sparse pixel-wise annotations, and then extends their method to interactive image segmentation and video object segmentation. [Dong and Xing, 2018] generalizes the few-shot semantic segmentation problem from 1-way (class) to N-way (classes). [Zhang *et al.*, 2019a] introduces an attention mechanism to effectively fuse information from multiple support examples and proposes an iterative optimization module to refine the predicted results. [Nguyen and Todorovic, 2019] and [Wang *et al.*, 2019] leverage the annotations of the support images as supervision in different ways. Different from [Tian *et al.*, 2019] which employs the base learner directly with linear classifier, our method devises a novel self-supervised base learner which is more intuitive and effective for few-shot segmentation. Compared to [Nguyen and Todorovic, 2019], the support images are used for supervision in both training and test stages of our framework.

## 3 Problem Description

Here we define an input triple  $Tr_i = (Q_i, S_i, T_S^i)$ , a label  $T_Q^i$  and a relation function  $F : A_i = F(Q_i, S_i, T_S^i; \theta)$ , where  $Q_i$  and  $S_i$  are the query and support images containing objects of the same class  $i$ , correspondingly.  $T_S^i$  and  $T_Q^i$  are the pixel-wise labels corresponding to the  $i^{th}$  class objects in  $S_i$  and  $Q_i$ .  $A_i$  is the actual segmentation result, and  $\theta$  is all parameters to be optimized in function  $F$ . Our task is to randomly sample triples from the dataset, train and optimize  $\theta$ , thus minimizing the loss function  $L$ :

$$\theta_* = \arg \min_{\theta} L(A_i, T_Q^i). \quad (1)$$

We expect that the relationship function  $F$  can segment object regions of the same class in another target image each time it sees few support images belonging to a new class. This is the embodiment of the meaning of few-shot segmentation. It should be mentioned that the classes sampled by the test set are not present in the training set, that is,  $U_{train} \cap U_{test} = \emptyset$ . The relation function  $F$  in this problem is implemented by the model detailed in Sec. 4.3.

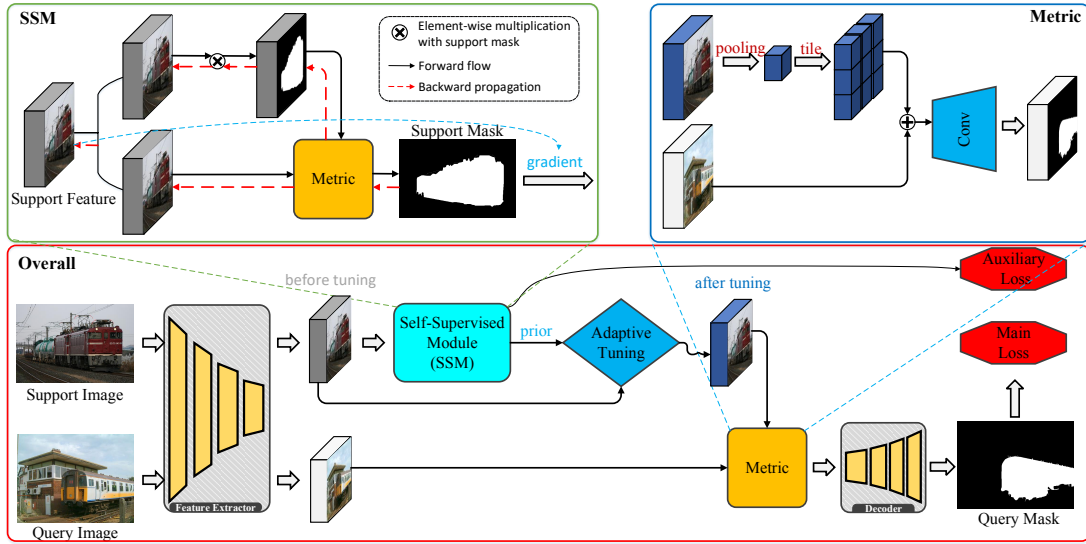


Figure 2: Overall architecture of our model. It mainly consists of an adaptive tuning mechanism, a self-supervised base learner and a deep non-linear metric.

## 4 Method

### 4.1 Model Overview

Different from the existing methods, this paper proposes a novel self-supervised tuning framework for few shot segmentation, which is mainly composed of an adaptive tuning mechanism, a self-supervised base learner and a meta learner. These three components will be illustrated in the next three subsections. In this subsection, we mainly present the description of the whole framework mathematically, and the symbolic representation is consistent with that in Sec. 3.

As shown in Fig. 2, a Siamese network  $f_e$  [Koch *et al.*, 2015] is first proposed to extract the features of input query and support images [Zhang *et al.*, 2012]. The mechanism of parameters sharing not only promotes optimization but also reduces the amount of calculation. In this step, the latent visual feature representations  $R_q$  and  $R_s$  are obtained as follows:

$$R_q = f_e(Q_i; \theta_e), \quad (2)$$

$$R_s = f_e(S_i; \theta_e). \quad (3)$$

Here  $\theta_e$  is the learnable parameter of the sharing encoder. To dynamically adjust the latent features, we firstly devise a novel self-supervised inner-loop as the base learner ( $f_b$ ) to exploit the underlying semantic information ( $\theta_s$ ).

$$\theta_s = f_b(R_s, T_S^i; \theta_b) \quad (4)$$

Then the distribution of low-level visual features is tuned ( $f_t$ ) according to the high-level category-specific cues obtained above:

$$R'_s = f_t(R_s; \theta_s) \quad (5)$$

Inspired by Relation Network [Sung *et al.*, 2018], a deep non-linear metric  $f_m$  is introduced into our meta learner. It measures the similarity between the feature map of query image and the tuned feature, and accordingly determines regions of

interest in query images. Finally, a segmentation decoder  $f_d$  is utilized to refine the response area to the original image size:

$$M = f_m(R_q, R'_s; \theta_m) \quad (6)$$

$$S = f_d(M; \theta_d) \quad (7)$$

Similarly,  $\theta_m$  and  $\theta_d$  stand for the parameters of metric and decoding part. With the ability to continuously learn from different episodes, our meta learner gradually optimizes the whole process above (outer loop) and improves the performance of base learner, metric and decoder.

### 4.2 Self-Supervised Tuning Scheme

The self-supervised tuning scheme is the core of our proposed method, which is implemented by a base learner driven by self-segmentation and an adaptive tuning module. Different from the existing base learner used in [Lee *et al.*, 2019] [Bertinetto *et al.*, 2018], our proposed base learner is generated under the supervision of the presented masks of input support images. This method originates from an intuitive idea, that is, the premise of identifying the regions with the same category as target objects is to identify the object itself first. Therefore, we duplicate the features of the support image as the two inputs of Eq. 6 in Sec. 4.1 and calculate the standard cross-entropy loss ( $L_{cross}$ ) with the corresponding support mask:

$$M_{sup} = f_m(R_s, R_s; \theta_m) \quad (8)$$

$$S_{sup} = f_d(M_{sup}; \theta_d) \quad (9)$$

$$L_{sup} = L_{cross}(S_{sup}, T_S^i) \quad (10)$$

To exploit the above information to better implement the segmentation task of the query branch, the marginalized distribution is adjusted in embedding space based on the category-specific semantic constraint. Specifically, a gradient map that is calculated by back-propagating self-supervised

+SSM	+SS Loss	i=0	i=1	i=2	i=3	mean
		49.6	62.6	48.7	48.0	52.2
✓		53.2	63.6	48.7	47.9	53.4
	✓	51.1	64.9	51.9	50.2	54.5
✓	✓	<b>54.4</b>	<b>66.4</b>	<b>57.1</b>	<b>52.5</b>	<b>57.6</b>

(a) Results for 1-shot segmentation.

+SSM	+SS Loss	i=0	i=1	i=2	i=3	mean
		53.2	66.5	55.5	51.4	56.7
✓		56.6	67.2	60.4	54.0	59.6
	✓	56.8	68.7	61.4	55.0	60.5
✓	✓	<b>58.6</b>	<b>68.7</b>	<b>63.1</b>	<b>55.3</b>	<b>61.4</b>

(b) Results for 5-shot segmentation.

 Table 1: Ablation study on PASCAL-5<sup>i</sup> dataset under the metric of mean-IoU. Bold fonts represent the best results.

loss through the support feature map is used as a guidance for augmenting each element in the support embedding space. Mathematically,

$$R'_s = R_s - \frac{\partial L_{sup}}{\partial R_s} \quad (11)$$

Note that only the feature representation is updated here and the network parameters are unchanged.

### 4.3 Deep Learnable Meta Learner

Inspired by the Relation Network, we apply a deep non-linear metric to measure the similarity between the feature map of query image and the descriptor generated by base learner. As in [Rakelly *et al.*, 2018], we also use the deep learnable metric with late fusion as the main component of meta learner. First, we multiply the features by the downsampling mask and aggregate it to obtain the latent features of the foreground. This feature is then tiled to the original spatial scale, so that each dimension of the query feature is aligned with the representative feature ( $R'_s$ ):

$$R_s^r = \text{tile}(\text{pool}(R'_s \cdot T_S^i)) \quad (12)$$

Through the Relation Network comparator, the response area of the query image is obtained:

$$M = \text{Relation}(R_q, R_s^r) \quad (13)$$

Finally, we feed it into the segmentation decoder, refining and restoring the original image size to get accurate segmentation results.

### 4.4 Loss and Generalization to 5-Shot Setting

In addition to the cross-entropy loss function of segmentation obtained by query-support set commonly used in other methods (main loss), we also include the cross-entropy loss of support-support segmentation from base learner (auxiliary loss) into the final training loss as an auxiliary. In our method, the auxiliary loss itself is created as an intermediate process, so there is no much extra computation required. It can be seen

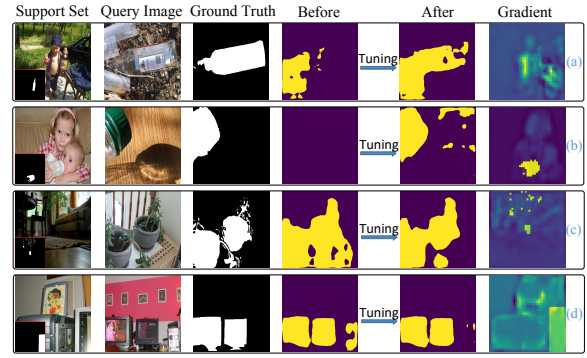


Figure 3: Visualization before and after applying the self-supervised module. From left to right in each row, they represent the support set, query image, the ground truth, two different segmentation results and the gradient information. The support mask is placed in the left-down corner of the support image.

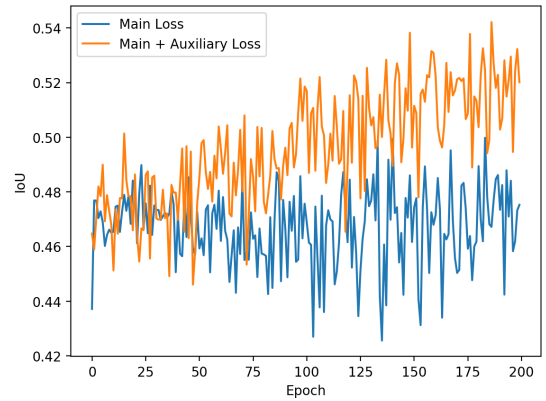


Figure 4: Curves of IoU results in test stage when different loss functions are applied.

from the experiment part that the auxiliary loss can accelerate the convergence and improve the performance.

When generalizing to 5-shot segmentation, the main difference is the base learner part. Considering that the number of samples for 5-shot segmentation is still small, if we calculate the gradient optimization together, it is easy to produce large discrepancy. Therefore, we calculate the gradient value separately, getting 5 separate response areas, and then take the weighted summation according to the self-supervised scores to get the final result, that is:

$$M_{weighted} = \sum_{i=1}^5 f_{IoU}(S_{sup}^i, T_S^i) \cdot f_m(R_q^i, R_s^{ii}; \theta_m) \quad (14)$$

where  $f_{IoU}$  represents the function which is used to calculate the IoU scores.

## 5 Experiment

### 5.1 Dataset and Settings

**Dataset.** To evaluate the performance of our model, we experiment on PASCAL-5<sup>i</sup> and COCO-20<sup>i</sup> datasets. The for-




Figure 5: Four sets of self-supervised results from different models trained with and without the auxiliary loss. The query image and support image (in the upper-right corner of the first image) are the same in each row.

Methods	Results (mean-IoU %)
Maximum	59.4
Average	61.2
Weighted	<b>61.4</b>

Table 2: Comparison among different fusion methods in 5-shot setting under the metric of mean-IoU.

mer is first proposed in [Shaban *et al.*, 2017] and is recognized as the standard dataset in the field of few-shot segmentation in subsequent work. That is, from the set of 20 classes on PASCAL dataset, we sample five and consider them as the test subsets  $U_{test_i} = \{4i + 1, \dots, 4i + 5\}$ , with  $i$  being the fold number ( $i = 0, 1, 2, 3$ ), and the remaining 15 classes form the training set  $U_{train_i}$ . COCO-20<sup>i</sup> dataset is proposed in recent work and the division is similar to PASCAL-5<sup>i</sup> dataset. In the test stage, we randomly sample 1000 pairs of images from the corresponding test subset.

**Settings.** The backbone of the existing methods are different, mainly VGG-16 [Simonyan and Zisserman, 2014] and ResNet-50 [He *et al.*, 2016]. To make a fair comparison, we separately train two models with different backbones for testing. As adopted in [Shaban *et al.*, 2017], we choose the per-class foreground Intersection-over-Union (IoU) and the average IoU over all classes (mean-IoU) as the main evaluation indicator of our task. While the foreground IoU and background IoU (FB-IoU) is a commonly used indicator in the field of binary segmentation, it is used by few papers of few-shot segmentation task. Because mean-IoU can better measure the overall performance of different classes and ignore the proportion of background pixels, we show the results of mean-IoU in all experiments. Our model uses the SGD optimizer during the training process. The initial learning rate




Figure 6: 5-shot segmentation results. The query image and the 5-shot results are placed in the first column, and the 1-shot results of 5 support images are in the second column. The green tick and red crosses represent right and wrong segmentation results, respectively.

is set to 0.0005 and the attenuation rate is set to 0.0005. The model stops training after 200 epochs. All images are resized to  $321 \times 321$  size and the batch size is set to 16.

## 5.2 Ablation Study

To prove the effectiveness of our architecture, We conduct several ablation experiments on PASCAL-5<sup>i</sup> dataset as shown in Table 1. To improve efficiency, we only choose ResNet50 as backbone for all models in the ablation study for fair comparison. The performance of our network is mainly attributed to two prominent components: SSM and auxiliary loss. Note that we progressively add addition components to the baseline, which enables us to gauge the performance improvement obtained by each of them. Due to the fact that the self-supervised tuning mechanism dynamically adjusts the marginalized distributions of latent features at each stage, SSM brings about a 1.2 and 2.9 percent mean-IoU increase in 1-shot and 5-shot settings, respectively. At the same time, we can see that auxiliary loss boosts the overall performance, resulting in a 2.3 and 3.8 percent improvement.

## 5.3 Analysis

As PASCAL-5<sup>i</sup> is the most commonly used dataset by all few-shot segmentation methods, the main analysis parts of our experiment are accomplished on PASCAL-5<sup>i</sup> dataset.

### Effect of the Self-Supervised Module

To clarify the actual function of the self-supervised module for few-shot segmentation task, we show the following visualization results. As shown in Fig. 3, we feed the feature representations generated before and after the tuning process to Relation Network to obtain two sets of segmentation results. Note that the original network often segments other objects by mistakes that are easily confused in the background. After adding the SSM module, the category-specific semantic constraint is introduced to help query images to correct the



Method	Backbone	i=0	i=1	i=2	i=3	mean
OSLSM	Vgg16	33.6	55.3	40.9	33.5	43.9
SG-One	Vgg16	40.2	58.4	48.4	38.4	46.3
PANet	Vgg16	42.3	58.0	51.1	41.2	48.1
FWBFS	Vgg16	47.0	59.6	52.6	48.3	51.9
Ours	Vgg16	<b>50.9</b>	<b>63.0</b>	<b>53.6</b>	<b>49.6</b>	<b>54.3</b>
CAnet	ResNet50	52.5	65.9	51.3	51.9	55.4
FWBFS	ResNet101	51.3	64.5	56.7	52.2	56.2
Ours	ResNet50	<b>54.4</b>	<b>66.4</b>	<b>57.1</b>	<b>52.5</b>	<b>57.6</b>

Table 3: Comparison with SOTA for 1-shot segmentation under the mean-IoU metric on PASCAL-5<sup>i</sup> dataset. Bold fonts represent the best results.

Method	Backbone	i=0	i=1	i=2	i=3	mean
OSLSM	Vgg16	35.9	58.1	42.7	39.1	43.8
SG-One	Vgg16	41.9	58.6	48.6	39.4	47.1
PANet	Vgg16	51.8	64.6	<b>59.8</b>	46.5	55.7
FWBFS	Vgg16	50.9	62.9	56.5	50.1	55.1
Ours	Vgg16	<b>52.5</b>	<b>64.8</b>	59.5	<b>51.3</b>	<b>57.0</b>
CAnet	ResNet50	55.5	67.8	51.9	53.2	57.1
FWBFS	ResNet101	54.8	67.4	62.2	55.3	59.9
Ours	ResNet50	<b>58.6</b>	<b>68.7</b>	<b>63.1</b>	<b>55.3</b>	<b>61.4</b>

Table 4: Comparison with SOTA for 5-shot segmentation under the mean-IoU metric on PASCAL-5<sup>i</sup> dataset. Bold fonts represent the best results.

prediction results. To demonstrate the meaning of this semantic constraint, the calculated gradient is visualized in the last column. We can see that it strengthens the focus on the regions of target categories.

### Importance of Auxiliary Loss

To verify the role of auxiliary loss in the training stage, we designed the following experiment. First, we show the mean-IoU curves generated with and without our auxiliary loss during test stage in Fig. 4. It can be clearly seen that the convergence is faster and the mean-IoU result is better after the loss is added. Then, we visualize the segmentation results of two models trained with and without auxiliary loss. As shown in Fig. 5, the model without auxiliary loss can not segment the support images themselves, therefore inapplicable for other images of the same category. The proposed auxiliary loss improves the self-supervised capability as well as the performance of few-shot segmentation.

### Comparison among Different 5-Shot Fusion Methods

To prove the superiority of our weighted fusion in 5-shot setting, we compare the 5-shot segmentation results with the average and maximum fusion methods, in which the average and maximum segmentation results of 5 support images are computed, respectively. It can be seen in Table 2 and Fig. 6 that our weighted fusion strategy achieves the best. The samples in the PASCAL-5<sup>i</sup> dataset are relatively simple, and most of the weights are close to one fifth, so the result of average fusion method is similar to weighted fusion in general.

Method	mean-IoU	
	1-shot	5-shot
PANet	20.9	29.7
FWBFS	20.0	22.6
Ours	<b>22.2</b>	<b>31.3</b>

Table 5: Comparison with SOTA under the mean-IoU metric on COCO-20<sup>i</sup> dataset. All models are implemented with VGG-16 backbone.

### 5.4 Comparison with SOTA

To better assess the overall performance of our network, we compare it to other methods (OSLSM [Shaban *et al.*, 2017], SG-One [Zhang *et al.*, 2018], PANet [Wang *et al.*, 2019], FWBFS [Nguyen and Todorovic, 2019] and CAnet [Zhang *et al.*, 2019a]) on PASCAL-5<sup>i</sup> and COCO-20<sup>i</sup> datasets.

#### PASCAL-5<sup>i</sup>

We train two types of SST (Self-Supervised Tuning) models with VGG-16 and ResNet-50 backbones (we call them SST-vgg and SST-res model) on PASCAL-5<sup>i</sup> dataset. It can be seen in Table 3 that our SST-vgg model surpasses the best existing method over two percentage points in 1-shot setting, and the SST-res model yields 2.2 points improvement, which is even 1.4 points higher than the method with ResNet-101 backbone. Under the setting of 5-shot, our SST-vgg and SST-res models significantly increase by 1.9 and 1.5 points as shown in Table 4, respectively. These comparisons indicate that our method boosts the recognition performance of few-shot segmentation.

#### COCO-20<sup>i</sup>

To prove that our method also has good generalization performance on larger datasets, we compare our method with others which recently report results on the COCO-20<sup>i</sup> dataset as shown in Table 5. Obviously, the average results of our method surpass other best methods by 1 and 1.6 points under 1-shot and 5-shot setting, respectively.

## 6 Conclusion

In this paper, a self-supervised tuning framework is proposed for few-shot segmentation. The category-specific semantic constraint is provided by the self-supervised inner loop and utilized to adjust the distribution of latent features across different episodes. The resulting auxiliary loss is also introduced into the outer loop of training process, achieving faster convergence and higher scores. Extensive experiments on benchmarks show that our model is superior in both performance and adaptability compared with existing methods.

## Acknowledgments

This work was supported by the National Key R&D Program of China under Grant 2017YFB130092, the National Natural Science Foundation of China (NSFC) under Grants 61872327, the Fundamental Research Funds for the Central Universities under Grant WK2380000001 as well as Huawei Technologies Co., Ltd.

## References

- [Bertinetto *et al.*, 2018] Luca Bertinetto, Joao F Henriques, Philip HS Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. *arXiv preprint arXiv:1805.08136*, 2018.
- [Chen *et al.*, 2017] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- [Chen *et al.*, 2019] Zhe Chen, Jing Zhang, and Dacheng Tao. Progressive lidar adaptation for road detection. *IEEE/CAA Journal of Automatica Sinica*, 6(3):693–702, 2019.
- [Dong and Xing, 2018] Nanqing Dong and Eric P Xing. Few-shot semantic segmentation with prototype learning. In *BMVC*, volume 1, page 6, 2018.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [Hu *et al.*, 2019] 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 *AAAI*, volume 33, pages 8441–8448, 2019.
- [Koch *et al.*, 2015] Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop*, volume 2. Lille, 2015.
- [Lee *et al.*, 2019] Kwonjoon Lee, Subhransu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with differentiable convex optimization. In *CVPR*, pages 10657–10665, 2019.
- [Long *et al.*, 2015] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, pages 3431–3440, 2015.
- [Munkhdalai and Yu, 2017] Tsendsuren Munkhdalai and Hong Yu. Meta networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2554–2563. JMLR. org, 2017.
- [Nguyen and Todorovic, 2019] Khoi Nguyen and Sinisa Todorovic. Feature weighting and boosting for few-shot segmentation. In *ICCV*, pages 622–631, 2019.
- [Qiao *et al.*, 2019] Tingting Qiao, Jing Zhang, Duanqing Xu, and Dacheng Tao. Mirrorgan: Learning text-to-image generation by redescription. In *CVPR*, pages 1505–1514, 2019.
- [Rakelly *et al.*, 2018] 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.
- [Rusu *et al.*, 2018] Andrei A Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. Meta-learning with latent embedding optimization. *arXiv preprint arXiv:1807.05960*, 2018.
- [Santoro *et al.*, 2016] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. One-shot learning with memory-augmented neural networks. *arXiv preprint arXiv:1605.06065*, 2016.
- [Shaban *et al.*, 2017] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots. One-shot learning for semantic segmentation. *arXiv preprint arXiv:1709.03410*, 2017.
- [Simonyan and Zisserman, 2014] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [Snell *et al.*, 2017] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pages 4077–4087, 2017.
- [Sung *et al.*, 2018] 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 *CVPR*, pages 1199–1208, 2018.
- [Tian *et al.*, 2019] Pinzhuo Tian, Zhangkai Wu, Lei Qi, Lei Wang, Yinghuan Shi, and Yang Gao. Differentiable meta-learning model for few-shot semantic segmentation. *arXiv preprint arXiv:1911.10371*, 2019.
- [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*, pages 9197–9206, 2019.
- [Zhang *et al.*, 2012] Hanwang Zhang, Zheng-Jun Zha, Shuicheng Yan, Jingwen Bian, and Tat-Seng Chua. Attribute feedback. In *MM*, pages 79–88, 2012.
- [Zhang *et al.*, 2014] Hanwang Zhang, Zheng-Jun Zha, Yang Yang, Shuicheng Yan, and Tat-Seng Chua. Robust (semi) nonnegative graph embedding. *IEEE transactions on image processing*, 23(7):2996–3012, 2014.
- [Zhang *et al.*, 2018] Xiaolin Zhang, Yunchao Wei, Yi Yang, and Thomas Huang. Sg-one: Similarity guidance network for one-shot semantic segmentation. *arXiv preprint arXiv:1810.09091*, 2018.
- [Zhang *et al.*, 2019a] 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*, pages 5217–5226, 2019.
- [Zhang *et al.*, 2019b] Qiming Zhang, Jing Zhang, Wei Liu, and Dacheng Tao. Category anchor-guided unsupervised domain adaptation for semantic segmentation. In *Advances in Neural Information Processing Systems*, pages 433–443, 2019.
- [Zhao *et al.*, 2017] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, pages 2881–2890, 2017.
- [Zhu *et al.*, 2019] Zhen Zhu, Mengde Xu, Song Bai, Tengting Huang, and Xiang Bai. Asymmetric non-local neural networks for semantic segmentation. In *ICCV*, pages 593–602, 2019.