

# FEW-SHOT SEMANTIC SEGMENTATION VIA PROTOTYPE AUGMENTATION WITH IMAGE-LEVEL ANNOTATIONS

Shuo Lei<sup>1</sup>, Xuchao Zhang<sup>1,2</sup>, Jianfeng He<sup>1\*</sup>, Fanglan Chen<sup>1</sup>, Chang-Tien Lu<sup>1</sup>

<sup>1</sup>Department of Computer Science, Virginia Tech, Falls Church, VA, USA

<sup>2</sup>NEC Laboratories America, Princeton, NJ, USA

{slei, xuczhang, jianfenghe, fanglanc, ctlu}@vt.edu

## ABSTRACT


Despite the great progress made by deep neural networks in the semantic segmentation task, traditional neural-network-based methods typically suffer from a shortage of large amounts of pixel-level annotations. Recent progress in few-shot semantic segmentation tackles the issue by only a few pixel-level annotated examples. However, these few-shot approaches cannot easily be applied to multi-way or weak annotation settings. In this paper, we advance the few-shot segmentation paradigm towards a scenario where image-level annotations are available to help the training process of a few pixel-level annotations. Our key idea is to learn a better prototype representation of the class by fusing the knowledge from the image-level labeled data. Specifically, we propose a new framework, called PAIA, to learn the class prototype representation in a metric space by integrating image-level annotations. Furthermore, by considering the uncertainty of pseudo-masks, a distilled soft masked average pooling strategy is designed to handle distractions in image-level annotations. Extensive empirical results on two datasets show superior performance of PAIA.

**Index Terms**— Few-shot learning, semantic segmentation

## 1. INTRODUCTION

Semantic segmentation, one of the most challenging tasks in computer vision, aims to assign a categorical label to each pixel of an image according to its enclosing object or region. In the past few years, a number of deep-neural-network-based approaches have been proposed for the semantic segmentation task. However, training these models typically requires large-scale pixel-level annotations, which are expensive to obtain. Some semi-/weakly-supervised segmentation models were proposed to reduce the dependence on pixel-level annotated data but still suffer from the issue of model generalization, which makes them hard to be applied to unseen categories.

\*Corresponding author.



**Fig. 1.** Illustration of a 1-way 1-shot with 2 image-level auxiliary images task. Our objective is to augment the class representation in the metric space by using image-level labeled data. Our model employs soft-mask average pooling strategy to handle the distraction in the image-level annotations and refine the related embedding features to refine the prototypes.

Recently, there has been increasing interest in the study of few-shot semantic segmentation, which learns to perform segmentation on novel classes with only a few pixel-level annotated images. However, existing studies suffer from three issues: 1) Hard to handle the multi-way segmentation task. One group of methods [1, 2, 3] perform segmentation by matching dense pair-wise features to query images. However, those methods mainly focus on one-way few-shot setting, and hence it is non-trivial to generalize them to the multi-way segmentation task. 2) Inaccurate class representations due to limited number of labeled data. The other group of studies [4, 5] are based on prototype learning, which learns a metric space to employ segmentation on new classes with similarity measurements. These prototype-based methods are highly dependent on the accuracy of prototype representations and provide a solution to the multi-way segmentation task. However, all those methods represent a semantic class based on a small support set, which are restrictive for capturing rich and accurate prototypes. 3) Poor robustness to the weak annotations. Recently, PPNNet [6] adopts extra image-level annotations and decomposes the holistic prototypes to enrich the prototype representations. Although matching the part-aware prototypes captures fine-grained object features, it cannot handle the weak annotations (like scribble or bounding box annotations) in the support set since part-aware prototypes require accurate pixel-

level annotations. In addition, it ignores the uncertainty of pseudo-labels when refining the prototypes. Concretely, it is non-trivial to select class-related parts in the image-level annotations due to lacking pixel-level annotations. Nevertheless, PPNet filters the irrelevant parts only by an unified pre-defined threshold, which may introduce some distractions to the prototypes.

To tackle the above issues, we first consider a new scenario to augment the prototype presentations, where the segmentation of novel classes are learned in the combination of pixel- and image-level annotations, which is shown in Fig.1. Specifically, in a  $N$ -way  $K$ -shot segmentation task, we aim to perform segmentation on the query images with  $K$  pixel-level and  $U$  image-level annotated images from each of the  $N$  classes. It is worth emphasizing that mixing strong (pixel-level) and weak (image-level) annotations is a widely used setting to improve the performance in the existing semantic segmentation works [7, 8]. But these methods still require large weak annotations to guide the training process. Instead, the number  $U$  of weak labeled data needed in our method is very small ( $U < 15$ ), which is more feasible for few-shot setting. Second, we propose a distilled soft-masked average pooling strategy to handle the distraction in image-level annotations. It considers the uncertainty of the pseudo-mask and applies the task-specific threshold to filter the class relevant parts in image-level annotations, which is more effective than an unified pre-defined hard threshold. Finally, we propose an Iterative Fusion Module (IFM) to refine the prototypes by integrating the prototype of image-level annotations into original one.

To sum up, our main contributions are as follows:

- Propose a class-prototype augmentation method to enrich the prototype representation by utilizing a few image-level annotations, achieving superior performance in one-/multi-way and weak annotation settings.
- Design a robust strategy with soft-masked average pooling to handle the noise in image-level annotations, which considers the prediction uncertainty and employs the task-specific threshold to mask the distraction.
- Conduct extensive experiments on two datasets for performance evaluations. Our method outperforms the state-of-the-arts with less image-level annotations and can achieve 8.2% and 6.8% improvements in mIoU score for one-shot settings with scribble and bounding box annotations in PASCAL-5<sup>i</sup>, respectively.

## 2. PROBLEM SETTING

Our purpose is that a model trained on a large labeled dataset  $\mathcal{D}_{train}$  can make a segmentation prediction on a testing dataset  $\mathcal{D}_{test}$  with a few annotated examples. The class set  $\mathcal{C}_{train}$  in  $\mathcal{D}_{train}$  has no overlap with  $\mathcal{D}_{test}$ , i.e.,  $\mathcal{C}_{test} \cap \mathcal{C}_{train} = \emptyset$ . Following previous works [4, 2],

we adopt an episodic paradigm in the few-shot segmentation task. Each episode  $e$  is composed by 1) a *support* set  $\mathcal{S}_e = \{(I_i^s, M_i^s)\}_{i=1}^{N \times K}$ , containing  $K$  (*image, mask*) pairs for each of the  $N$  categories in the foreground, where  $(I_i^s, M_i^s)$  represents the pair of support image  $I_i^s$  and its corresponding binary mask  $M_i^s$  for the foreground class, 2) a *query* set  $\mathcal{Q}_e = \{(I_j^q, M_j^q)\}_{j=1}^Q$ , which contains  $Q$  different query sample pairs  $(I_j^q, M_j^q)$  from the same  $N$  categories, and 3) an *auxiliary* set  $\mathcal{A}_e = \{I_t^a\}_{t=1}^{N \times U}$ , containing  $U$  image-level labeled images  $I_t^a$  for each of the same  $N$  categories, but no pixel-level annotation is available. The set of all target classes in the foreground for episode  $e$  denotes as  $\mathcal{C}_e$ , and  $|\mathcal{C}_e| = N$ . For each episode  $e$ , the model is supposed to segment images from  $\mathcal{Q}_e$  with the combination of  $\mathcal{S}_e$  and  $\mathcal{A}_e$ .

## 3. MODEL

### 3.1. Overall Architecture

We propose a new framework that can solve the few-shot segmentation problem with a combination of pixel- and image-level labeled data. The main idea of our model is to learn a better prototype representation of the class by fusing the knowledge from the image-level labeled data. Specifically, the original prototypes are first obtained on the support set and are used to segment image-level annotations. Class irrelevant parts in those images should be filtered and the most confidently related features are considered to augment the original prototypes. To this end, we propose a novel prototype fusion strategy that contains the distilled soft-masked average pooling method and iterative fusion module, as shown in Fig. 2.

### 3.2. Prototype Representation Learning


We represent each category of segmentation task as a prototype in the metric space. The original class-specific prototypes are obtained by employing masked average pooling over the support set, which averages the features of the pixels only belonging to the support classes, and each pixel of the query image is labeled by its nearest prototype in the metric space. Thus, the prototype  $p_c$  of the foreground class  $c$  is defined as follows:

$$p_c = \frac{1}{K} \left( \sum \text{MP}(I_{c,i}^s, M_{c,i}^s) \right), \quad (1)$$

We have  $I_{c,i}^s = \{x_j\}_{j=1}^{w \times h}$  and  $M_{c,i}^s = \{y_j\}_{j=1}^{w \times h}$ , where  $w$  and  $h$  denote the width and height of the image, respectively.  $\text{MP}(\cdot)$  denotes the masked average pooling function. Moreover, the background prototype is computed by averaging all the features of the pixels that do not belong to any foreground class in  $\mathcal{C}_e$ .

### 3.3. Soft-masked Average Pooling

Our model enhances the prototypes by extracting more class representation knowledge from the additional image-level an-



**Fig. 2.** Overview of the proposed network structure. The support, auxiliary, and query images are embedded into deep features via the feature encoder. Then the model employs masked average pooling over the support set to get the original prototypes. Iterative Fusion Module (IFM) augments the original prototypes by re-feeding the related auxiliary features with the proposed distilled soft-masked average pooling method. The fused prototypes are obtained via the IFM and are used to segment query images.  $\mathcal{L}_{sup}$  and  $\mathcal{L}_{qry}$  denote the standard cross entropy loss between the segmentation results and the ground truth of the support set and query set, respectively.

notations. The most intuitive way to incorporate those image-level annotations is to obtain their pseudo masks and directly add them into the support set. However, this process may introduce some noise into the support set. Since the original prototypes can be biased due to support data scarcity, it may lead to inaccurate prediction results of image-level labeled data. To tackle this issue, we propose a soft-masked pooling method (SMP). Instead of assigning the same weight to each pixel belonging to the support class, we give them a partial assignment based on their probability of falling into the class. Pixels with lower predicted confidence would get lower weights preventing them from distracting the original prototypes. Specifically, for each foreground class  $c$ , we first compute the predicted probability map  $Y_{c,i}^a = \{\hat{y}_j^c\}_{j=1}^{w \times h}$  and pseudo binary mask  $\hat{M}_{c,i}^a = \{z_j\}_{j=1}^{w \times h}$  of  $I_{c,i}^a$ , where the indicator  $z_j$  is set to 1 if  $\hat{y}_j = c$ . Then, we compute the representative vector by averaging the pixels within the object regions on the feature map. Thus, the soft-masked average pooling can be formed as:

$$\text{SMP}(I_{c,i}^a, Y_{c,i}^a, \hat{M}_{c,i}^a) = \frac{\sum_j f(x_j) \hat{y}_j z_j}{\sum_j \hat{y}_j z_j}, \quad (2)$$

where  $f(\cdot)$  is defined as a feature encoder function. In this way, the original prototypes can be enhanced by incorporating part of image-level labeled samples. The fused prototypes can be computed as follows:

$$\tilde{p}_c = \frac{1}{K+U} \left( \sum_i \text{MP}(I_{c,i}^s, M_{c,i}^s) + \sum_j \text{SMP}(I_{c,j}^a, Y_{c,j}^a, \hat{M}_{c,j}^a) \right). \quad (3)$$

### 3.4. Distilled Soft-masked Average Pooling

In our problem setting, each image contains at least one foreground class in an episode. Therefore, the categories of referred segmentation mask belong to at least two of the  $N+1$  classes (including the background class). However, when computing the prototype of background class, we treat all pixels not belonging to foreground classes as the same category.

That means we cannot guarantee two support images have similar background class representation even if their foreground classes are the same. Moreover, the image-level annotated image may contain an unseen object, which does not show up in the background or does not belong to any foreground classes in the support set.

Under this circumstance, those unlabeled pixels could still get pseudo labels with higher confidence even if they are far away from all prototypes in the metric space. So the uncertainty of these unseen objects in images may reduce the accuracy of fused prototypes. To alleviate the issue, we use a filter strategy for each prototype when applying soft-masked average pooling over those unlabeled images, which is called distilled soft-masked average pooling. Inspired by Ren et al. [9], we try to compute a threshold  $\gamma_c$  based on the statistics of the distances between pixels and the prototypes. Specifically, we first compute the distance  $d_j^c$  between the prototype  $p_c$  and the pixel  $j$ , and obtain distance matrix  $D_i^c$  of image  $I_i^a$  for the prototype  $p_c$ , i.e.  $D_i^c = \{d_j^c\}_{j=1}^{w \times h}$ . Then, normalized distance set  $\tilde{D}^c$  is obtained by normalizing each distance from the distance set  $D^c = \{D_i^c\}_{i=1}^{N \times U}$ . Finally, the filter threshold  $\gamma_c$  for the prototype  $p_c$  in each episode  $e$  is defined as follows:

$$\gamma_c = \text{MLP}([\min_i(\tilde{D}_i^c), \max_i(\tilde{D}_i^c), \text{var}_i(\tilde{D}_i^c), \text{skew}_i(\tilde{D}_i^c), \text{kurt}_i(\tilde{D}_i^c)]). \quad (4)$$

For each foreground class  $c$ , the distraction indicator  $\tau_i$  of pixel  $i$  can be computed as  $\mathbb{1}(d(f(x_i), p_c) < \gamma_c)$ , where  $\mathbb{1}(\cdot)$  is an indicator function, outputting value 1 if the argument is true or 0 otherwise. Then, the indicator  $H_{c,i}^a = \{\tau_j\}_{j=1}^{w \times h}$  of  $I_{c,i}^a$  is applied to filter the pixels that are not worth considering. In this way, the model is forced to only extract objective class-related pixels instead of considering the whole image which may contain novel object classes in the background. Therefore, Eq.3 for the fused prototype computation can be updated as follows:

$$\bar{p}_c = \frac{1}{K+U} \left( \sum_i \text{MP}(I_{c,i}^s, M_{c,i}^s) + \sum_j \text{SMP}(I_{c,j}^a, Y_{c,j}^a, H_{c,j}^a \odot \hat{M}_{c,j}^a) \right), \quad (5)$$

where  $\odot$  is the element-wise product.

### 3.5. Iterative Fusion Module

Intuitively, if the knowledge extracted from the image-level annotations in the auxiliary set can improve the performance of our model, we can also utilize the image-level annotated images from the query set. As the original prototypes are inevitably biased due to data scarcity, the confidence of the initial probability maps of those images may not be high enough to be considered. Therefore, we iteratively repeat the refinement for several steps to optimize the fusion prototypes in the Iterative Fusion Module (IFM). This process is shown in Figure 2. In particular, we first compute the probability maps via the original prototypes and re-feed the embedding features with distilled soft-masked average pooling to the IFM. Then we alternatively use fused prototypes in the last epoch to recompute the probability maps. In this way, the bias of original prototypes can be reduced by adding more class-related features from image-level annotations. The more accurate prototypes are, the higher confidence can be obtained and the more class-related features can be considered.

## 4. EXPERIMENTS

### 4.1. Experimental Settings

**Datasets.** We evaluated the performance of our model on two common few-shot segmentation datasets: PASCAL-5<sup>i</sup> and COCO-20<sup>i</sup>. PASCAL-5<sup>i</sup> dataset is proposed by Shaban et.al [1] and is created from PASCAL VOC 2012 [10] with SBD [11] augmentation. The 20 categories in PASCAL VOC are evenly divided into 4 splits, each containing 5 categories. We used the rest of the images that do not have segmentation labels but have category information in PASCAL VOC 2012 as the auxiliary set. Similarly, COCO-20<sup>i</sup> is built from MS COCO [12] and 80 categories are split into 4 folds. As each image in MS COCO has its corresponding segmentation label, we used images in the validation folder as the auxiliary set. Models were trained on 3 splits and evaluated on the rest one in a cross-validation for both datasets. Following the same scheme for testing [4], we averaged the results from 5 runs with different random seeds, each run containing 1,000 episodes to get stable results.  $N_{query} = 1$  is used for all experiments.

**Implementation details.** We adopted a VGG-16 [13] and ResNet-50 [14] network as the feature extractor following conventions. For the MLP used in the distilled soft-masked pooling, we used a single hidden layer with 20 hidden units with a Tanh activation. For implementation, we used PyTorch [15]. Following previous works [16, 4, 17], we pre-trained the CNN on ImageNet [18]. All images were resized

to  $417 \times 417$  and augmented by random horizontal flipping. The network was trained end-to-end by SGD with a learning rate of 1e-3, momentum of 0.9 and weight decay of 5e-4. We trained the model in 20,000 iterations and the batch size is 1. The learning rate was reduced by 0.1 after 10,000 iterations. We adopted mean-IoU as the metric method to evaluate the model performance [4, 2]. PANet\* denotes the baseline that taking ResNet-50 [14] as feature extractor in PANet [4].

### 4.2. Comparison with the Competing Methods

We first compared our PAIA model with the state-of-the-art methods on PASCAL-5<sup>i</sup> dataset in 1-way segmentation task. Table 1 shows the results in mean-IoU metric. Specifically, compared with PANet [4], our model achieves an improvement of 4.93% in the 1-way 1-shot task and 4.78% in the 5-shot task, which means the combination of both pixel-level and image-level annotations can improve the performance in the few-shot segmentation task. Moreover, our model employs 2 image-level annotations and surpasses PPNet [6] by 1.19% for 1-shot and 1.07% for 5-shot while the PPNet method even uses 3 times more image-level annotations. This indicates that our soft-masked average pooling strategy can enhance the utilization of the image-level annotations. Table 3 shows the evaluation results on COCO-20<sup>i</sup>. Compared to PASCAL VOC dataset, MS COCO has more object categories, which makes it more difficult than PASCAL VOC. Our model can outperform the previous methods due to the capability of extracting class-related knowledge from the image-level annotations even though more unseen objects are included.



**Fig. 3.** Qualitative examples of 2-way 1-shot segmentation on the PASCAL-5<sup>i</sup>.

### 4.3. Results on Weak Annotations

To validate the robustness of our model, we evaluated PAIA with two types of weak annotations: scribble and bounding box. The pixel-level annotations of the support set were replaced by scribbles or bounding boxes. For fair comparison, we adopted VGG-16 as feature extractor and used the same annotation settings in Wang et al. [4]: scribbles are generated from the dense segmentation masks automatically and bounding box is randomly chosen from instance mask.

**Table 1.** Mean-IoU of 1-way 1-shot and 5-shot segmentation on PASCAL-5<sup>i</sup>.  $S^u$  denotes the image-level annotations. CANet reports multi-scale test performance. The single-scale test performance is reported by [19].

Methods	$S^u$	Backbone	1-shot					5-shot				
			split-1	split-2	split-3	split-4	mean	split-1	split-2	split-3	split-4	mean
OSLSM [1]	×	VGG16	33.60	55.30	44.90	33.50	40.80	35.90	58.10	42.70	39.10	43.90
SG-One [16]	×	VGG16	40.20	58.40	48.40	38.40	46.30	41.90	58.60	48.60	39.40	47.10
AMP [20]	×	VGG16	36.80	51.60	46.90	36.00	42.80	44.60	58.00	53.30	42.10	49.50
FWB [17]	×	VGG16	47.04	59.64	52.61	48.27	51.90	50.87	62.86	56.48	50.09	55.08
PANet [4]	×	VGG16	42.30	58.00	51.10	41.20	48.10	51.80	64.60	59.80	46.50	55.70
PANet* [4]	×	RN50	44.03	57.52	50.84	44.03	49.10	55.31	67.22	61.28	53.21	59.26
CANet [2]	×	RN50	49.56	64.97	49.83	<b>51.49</b>	53.96	-	-	-	-	55.80
PMMs [19]	×	RN50	<b>51.98</b>	<b>67.54</b>	51.54	49.81	<b>55.22</b>	55.03	68.22	52.89	51.11	56.81
PPNet [6]	✓	RN50	48.58	60.58	55.71	46.47	52.84	58.85	68.28	66.77	57.98	62.97
<b>PAIA (ours)</b>	✓	VGG16	49.50	61.64	<b>56.03</b>	45.61	53.20	54.14	63.82	62.30	50.53	57.70
<b>PAIA (ours)</b>	✓	RN50	50.31	62.10	55.97	47.72	54.03	<b>59.70</b>	<b>69.82</b>	<b>66.92</b>	<b>59.73</b>	<b>64.04</b>

**Table 2.** Mean-IoU of 2-way 1-shot and 5-shot segmentation on PASCAL-5<sup>i</sup>.  $S^u$  denotes the image-level annotations.

Methods	$S^u$	Backbone	1-shot					5-shot				
			split-1	split-2	split-3	split-4	mean	split-1	split-2	split-3	split-4	mean
SG-One [16]	×	VGG16	-	-	-	-	-	-	-	-	-	29.40
PANet [4]	×	VGG16	-	-	-	-	45.10	-	-	-	-	53.10
PPNet [6]	✓	RN50	47.36	58.34	52.71	<b>48.18</b>	51.65	<b>55.54</b>	67.26	64.36	58.02	61.30
<b>PAIA (ours)</b>	✓	VGG16	47.71	<b>60.04</b>	53.30	46.02	51.76	49.10	60.72	58.64	48.22	54.17
<b>PAIA (ours)</b>	✓	RN50	<b>48.94</b>	59.95	<b>54.11</b>	47.79	<b>52.70</b>	55.21	<b>68.50</b>	<b>64.97</b>	<b>59.76</b>	<b>62.11</b>


**Table 3.** Mean-IoU of 1-way segmentation on COCO-20<sup>i</sup>.  $S^u$  denotes the image-level annotations.

Method	$S^u$	Backbone	mean-IoU	
			1-shot	5-shot
PANet* [4]	×	RN50	22.95	33.80
PPNet [6]	✓	RN50	27.16	36.73
<b>PAIA (ours)</b>	✓	RN50	<b>28.12</b>	<b>37.63</b>

**Table 4.** Mean-IoU of using different types of annotations on PASCAL-5<sup>i</sup>.  $S^u$  denotes the image-level annotations.

Method	$S^u$	Scribble		Bounding Box		Densed	
		1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
PANet [4]	×	44.8	54.6	45.1	52.8	48.1	55.7
<b>PAIA</b>	✓	<b>53.0</b>	<b>57.3</b>	<b>51.9</b>	<b>56.2</b>	<b>53.2</b>	<b>57.7</b>

As shown in Table 4, for scribble annotations, our model achieves significant improvements of 8.2% and 2.7% in 1-shot and 5-shot tasks, respectively. This performance is comparable to the result with an expensive pixel-level annotated support set, which means our model works very well with sparse annotations. In addition, with bounding box annotations, our model significantly outperforms the state-of-the-art methods by 6.8% for the 1-shot task and 3.4% for the 5-shot task. This demonstrates that our model has a greater ability to withstand the noise introduced by the background area within the bounding box. Furthermore, the improvements in weak



**Fig. 4.** Qualitative examples of 1-way 1-shot segmentation. The left two columns are on PASCAL-5<sup>i</sup> and the right two columns are on COCO-20<sup>i</sup>.

annotations validate the robustness of our model.

#### 4.4. Ablation Study

We implemented extensive ablation experiments on the PASCAL-5<sup>i</sup> dataset to evaluate the effectiveness of different components in our network by using the mean-IoU metric in the 1-way 1-shot task. In Table 5, we compare our model with two baseline models. The first one does not adopt the distilled strategy when applying soft-masked pooling (DSMP), which

is denoted as PAIA-Smp. The second one does not employ an additional iterative fusion module for the fused prototypes, i.e., the initial prediction from PAIA (PAIA-Init). As shown in Table 5, the distilled soft-masked pooling method achieves a 2.2% improvement over the soft-masked pooling method. In addition, the iterative fusion module yields an improvement of 1.1% over the initial prediction. The combination of both modules achieves the best performance.

**Table 5.** Ablation study on the choice of proposed module on 1-way 1-shot segmentation task on PASCAL-5<sup>i</sup>.

	Backbone	DSMP	IFM	mean-IoU
PAIA-Smp	VGG16	✓		51.0
PAIA-Init	VGG16		✓	52.1
PAIA	VGG16	✓	✓	<b>53.2</b>

## 5. CONCLUSION

In this paper, a novel weak-annotation-augmented few-shot segmentation model is proposed to learn an augmented prototype based on both pixel-level and image-level annotations. To achieve this, we design a robust strategy with soft-masked average pooling to handle the noise in image-level annotations. It considers the prediction uncertainty of the image-level annotations and employs the task-specific threshold to filter the distraction. Our evaluation results demonstrated the superiority of the proposed method over existing few-shot segmentation models by a sizeable margin.

## 6. REFERENCES

- [1] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots, "One-shot learning for semantic segmentation," *arXiv preprint arXiv:1709.03410*, 2017.
- [2] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5217–5226.
- [3] 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 International Conference on Computer Vision*, 2019, pp. 9587–9595.
- [4] Kaixing 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 International Conference on Computer Vision*, 2019, pp. 9197–9206.
- [5] Jinlu Liu and Yongqiang Qin, "Prototype refinement network for few-shot segmentation," *arXiv preprint arXiv:2002.03579*, 2020.
- [6] Yongfei Liu, Xiangyi Zhang, Songyang Zhang, and Xuming He, "Part-aware prototype network for few-shot semantic segmentation," in *European Conference on Computer Vision*. Springer, 2020, pp. 142–158.
- [7] Ping Luo, Guangrun Wang, Liang Lin, and Xiaogang Wang, "Deep dual learning for semantic image segmentation," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 2718–2726.
- [8] Nasim Souly, Concetto Spampinato, and Mubarak Shah, "Semi supervised semantic segmentation using generative adversarial network," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 5688–5696.
- [9] Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B Tenenbaum, Hugo Larochelle, and Richard S Zemel, "Meta-learning for semi-supervised few-shot classification," *arXiv preprint arXiv:1803.00676*, 2018.
- [10] 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*, vol. 88, no. 2, pp. 303–338, 2010.
- [11] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik, "Semantic contours from inverse detectors," in *Proceedings of the IEEE International Conference on Computer Vision*, 2011, pp. 991–998.
- [12] 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 *European conference on computer vision*. Springer, 2014, pp. 740–755.
- [13] Karen Simonyan and Andrew Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [15] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer, "Automatic differentiation in pytorch," 2017.
- [16] Xiaolin Zhang, Yunchao Wei, Yi Yang, and Thomas S Huang, "Sg-one: Similarity guidance network for one-shot semantic segmentation," *IEEE Transactions on Cybernetics*, vol. 50, no. 9, pp. 3855–3865, 2020.
- [17] Khoi Nguyen et al., "Feature weighting and boosting for few-shot segmentation," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 622–631.
- [18] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Sathesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al., "Imagenet large scale visual recognition challenge," *International journal of computer vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [19] Boyu Yang, Chang Liu, Bohao Li, Jianbin Jiao, and Qixiang Ye, "Prototype mixture models for few-shot semantic segmentation," in *European Conference on Computer Vision*. Springer, 2020, pp. 763–778.
- [20] Mennatullah Siam et al., "Amp: Adaptive masked proxies for few-shot segmentation," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 5249–5258.