

Relevant Intrinsic Feature Enhancement Network for Few-Shot Semantic Segmentation

Xiaoyi Bao,^{*1,2,3} Jie Qin,^{*1,2} Siyang Sun,³ Xingang Wang,^{2†} Yun Zheng³

¹School of Artificial Intelligence, University of Chinese Academy of Sciences

²Institute of Automation, Chinese Academy of Sciences

³Alibaba group

{baoxiaoyi2021, qinjie2019}@ia.ac.cn, siyang.ssy@alibaba-inc.com,
xingang.wang@ia.ac.cn, zhengyun.zy@alibaba-inc.com

Abstract

For few-shot semantic segmentation, the primary task is to extract class-specific intrinsic information from limited labeled data. However, the semantic ambiguity and inter-class similarity of previous methods limit the accuracy of pixel-level foreground-background classification. To alleviate these issues, we propose the Relevant Intrinsic Feature Enhancement Network (RiFeNet). To improve the semantic consistency of foreground instances, we propose an unlabeled branch as an efficient data utilization method, which teaches the model how to extract intrinsic features robust to intra-class differences. Notably, during testing, the proposed unlabeled branch is excluded without extra unlabeled data and computation. Furthermore, we extend the inter-class variability between foreground and background by proposing a novel multi-level prototype generation and interaction module. The different-grained complementarity between global and local prototypes allows for better distinction between similar categories. The qualitative and quantitative performance of RiFeNet surpasses the state-of-the-art methods on *PASCAL-5ⁱ* and *COCO* benchmarks.

1 Introduction

As a fundamental and crucial task in the field of computer vision, semantic segmentation is widely applied in visual tasks such as medical image understanding, industrial defect inspection, virtual reality games, autonomous driving, *etc.* With the development of convolutional neural networks, fully-supervised semantic segmentation has achieved remarkable success (Long, Shelhamer, and Darrell 2015). Subsequent transformer-based methods also greatly improve the segmentation performance (Zhu et al. 2020; Xie et al. 2021; Yuan et al. 2021). However, it is still arduous to acquire pixel-level annotations that require a huge amount of manual effort and costs. To alleviate such expensive and unwieldy annotations, many works tend to resort to a few-shot semantic segmentation paradigm. In this setting, the model learns to segment with a few limited labeled data as the support set and then transferred to test the query inputs.

*These authors contributed equally.

†Corresponding author.

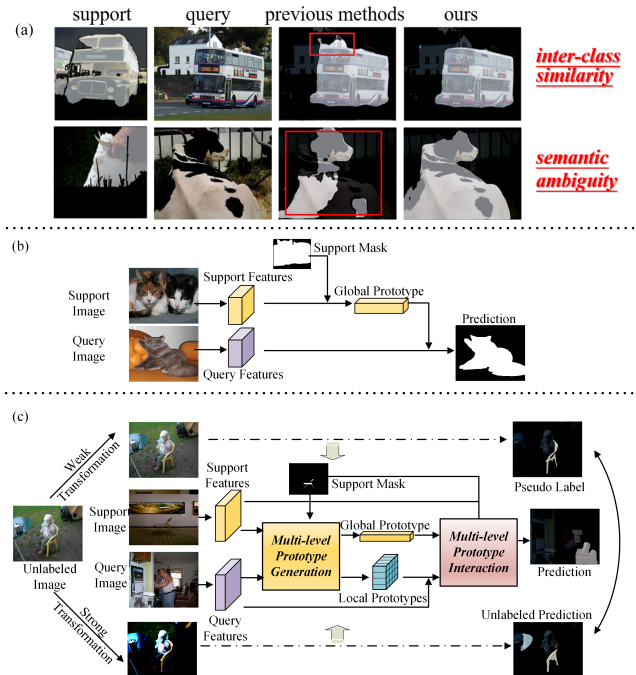


Figure 1: Comparison of RiFeNet and other works. (a) Two foreground-related issues that limit the effectiveness of previous research. (b) Schedule of previous prototype-based methods. The two-branch structure uses a single global prototype as the only medium for information interaction. (c) Framework of RiFeNet. An additional unlabeled branch is used for feature mining during training, with multi-granularity prototypes extracted from different branches.

Previous approaches (Shaban et al. 2017; Liu et al. 2020) focus on making full use of limited support data. The training data is divided into a support set and a query set, which are processed using two different branches. Relationship prompt and information interaction between the two branches are achieved by proposing superior matching strategies (Kang and Cho 2022) or generating representative prototypes (Tian et al. 2020; Wang et al. 2019; Zhang et al. 2019a, 2020). The former class of methods is represented

by HSNet (Min, Kang, and Cho 2021), which usually uses a cost volume method to aggregate dense correlation scores in high-dimensional space by 4D convolution. Subsequent researches improve the network using a swin transformer (Hong et al. 2022) or a doubly deformable transformer (Xiong, Li, and Zhu 2022) in 4-dimension space. With much lower computational costs, prototype-based methods also achieve good segmentation results. As illustrated in Fig. 1 (a), the single or multiple prototypes integrating semantic class-specific features are extracted from the support branch and used to activate query features. Researchers have proposed different activation methods in PpNet (Liu et al. 2020), PFENet (Tian et al. 2020), CyCTR (Zhang et al. 2021), *etc.*, to fully exploit the limited category features.

The object to be segmented is called the foreground while the background refers to the rest of the stuff and things. Despite the prevailing success of these methods, there are still two main issues affecting the segmentation effect of previous methods, *i.e.* **semantic ambiguity** and **inter-class similarity**, as shown in Fig.1 (a). For the foreground itself, semantic ambiguity appears for different instances of the same class. The intra-class differences between the two branches lead to semantic mistakes on query images. Further to the problem of distinguishing between foreground and background, inter-class similarity is specifically the difficulty in pixel-level binary classification. When objects of different classes with similar textures appear simultaneously, the local features of foreground and background become confusing.

The semantic ambiguity of the foreground is caused by the poor intra-class generalization of the model. Previous models mainly explore how to match the support set and the query set. With little labeled data in few-shot tasks, their models are prone to extract semantically ambiguous shallow features such as shape and pose, and thus unable to learn robust class representations for highly diverse query appearances. This distribution discrepancy makes the model identify the semantic parts of the foreground objects incompletely and segment them inappropriately. As for the inter-class similarity between foreground and background, the lack of discriminative features for query data is to blame for this problem. Specifically, previous methods extract global prototypes or prototypes with a single level of granularity. The information they carry is inadequate and monolithic, thus leading to pixel misidentify.

To address the above problems, we propose a novel relevant intrinsic feature enhancement network, as shown in Fig. 1 (c). For foreground objects, we enhance the semantic consistency of the same class, thus improving the intra-class generalization of features. We incorporate a novel unlabeled branch into the support-query pair branches to teach the model how to mine intrinsic robustness behind appearance variability. The additional unlabeled data branch serves as a regularization term during training that enhances the semantic consistency of few-shot segmentation models.

We go one step further in achieving an overall enhancement of the distinction between foreground and background. A novel multi-level prototype generation and interaction module is proposed to ensure inter-class variability. The multi-level prototype features contain a global prototype

from the support data and a local prototype from the query data. The former represents the high-level semantic abstraction about the entire structure of the target category, while the latter captures the fine-grained appearance concepts providing details for category discrimination. The global and local prototype interaction can complement each other to ensure that the network extracts the corresponding category features. This multi-level prototype feature interaction module can greatly widen the inter-class differences, benefiting the identification between foreground and background.

Our contributions can be summarized as follows:

- We propose a relevant intrinsic feature enhancement network. By alleviating semantic ambiguity and inter-class similarity, our model improves the forehead segmentation performance in this few-shot task.
- To maintain the foreground semantic consistency, we propose an unlabeled branch as an efficient data utilization method, which teaches the model how to extract intrinsic features robust to intra-class differences.
- To further achieve an overall effective distinction between foreground and background objects, we propose a novel multi-level prototype generation and interaction module to extend the inter-class variability. With different-grained semantic information from different sources, the mutual complementarity of information is facilitated and the discriminative representation ability of confusing foreground and background objects is boosted.
- Extensive experiments demonstrate the effectiveness of our proposed method, which achieves state-of-the-art accuracy on both PASCAL-5ⁱ and COCO benchmarks.

2 Related Work

2.1 Semantic Segmentation

Rapid progress has been made on semantic segmentation tasks since fully convolutional networks (FCN) transformed them into pixel-level classification (Long, Shelhamer, and Darrell 2015). After that, the corresponding encoder-decoder architecture has been widely used (Qin et al. 2022b,a, 2023). Recent research mainly focuses on multi-scale feature fusion (Zhao et al. 2017; Chen et al. 2018; Yang et al. 2018; He et al. 2019), insertion of attention modules (Fu et al. 2019; Yuan et al. 2018; Li et al. 2019; Tao, Sapra, and Catanzaro 2020; Zhu et al. 2019; Zhang et al. 2019b), and context prior (Lin et al. 2017; Zhang et al. 2018; Yu et al. 2020; Jin et al. 2021).

Inspired by the success of the vision transformer (Dosovitskiy et al. 2020), researchers have carried out attempts to apply the transformer structure to segmentation (Yuan et al. 2021; Wang et al. 2021; Lee et al. 2022; Xie et al. 2021; Strudel et al. 2021). These works perform well under the task setting of semantic segmentation. However, in practical application scenarios, they are unable to cope with sparse training data and thus often fail on categories that they have not seen during the training process.

2.2 Few-Shot Segmentation

A two-branch structure is widely used in few-shot segmentation, *i.e.*, a support branch and a query branch. Mainstream

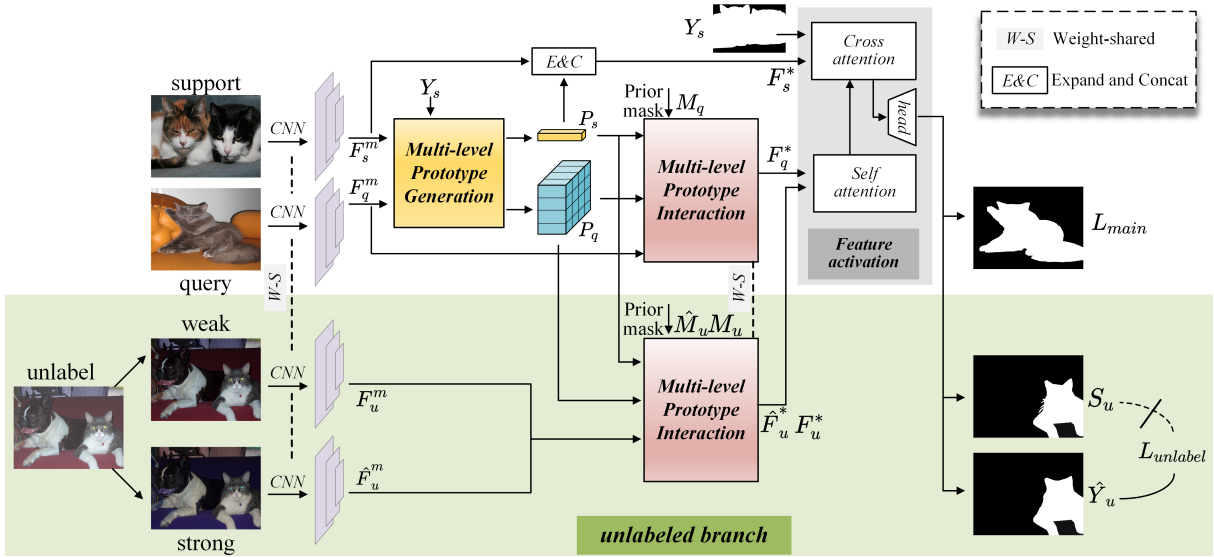


Figure 2: Overall architecture of RiFeNet. An additional branch using unlabeled inputs is attached to the traditional two-branch structure, as shown in green. The forward process of RiFeNet consists of three main blocks. The multi-level prototype generation block generates global and local prototypes. The multi-level prototype interaction module allows for interaction and integration between prototypes. A feature activation block is used consequently for obtaining final segmentation results.

few-shot methods are divided into two categories: prototype extraction and spatial correlation. With the development of image-matching, spatial correlation-based methods have regained attention in few-shot segmentation. HSNet (Min, Kang, and Cho 2021) applies 4D convolution to 4D correlation tensors in a pyramid structure, and the following ASNet (Kang and Cho 2022) squeezes semantic correlation into the foreground map. VAT (Hong et al. 2022) and DACM (Xiong, Li, and Zhu 2022) use cost volume aggregation-based models with 4D transformers.

The above methods benefit from the retained spatial structure. However, their high computational complexity and the excessive number of parameters are the drawbacks, making training slow and failing to generalize. Since PL (Dong and Xing 2018) introduced prototype learning to few-shot segmentation, most of the following research focus on prototype-based methods (Zhang et al. 2019a; Wang et al. 2019). SG-ONE (Zhang et al. 2020) proposes to use masked average pooling to extract a prototype from support features carrying the category information. Dense matching computation is performed between query features and this prototype. Tian proposes the PFENet (Tian et al. 2020), adding a train-free prior mask to the feature enrichment process. Multi-prototypes are generated in PpNet (Liu et al. 2020), ASGNet (Li et al. 2021), and RPMMs (Yang et al. 2020) to include more local information with different generating methods, while the use of multiple prototypes in these networks does not bring them more competitive results. IPMT (Liu et al. 2022) proposes an intermediate prototype and uses the transformer to update the prototype iteratively.

SSP (Fan et al. 2022) also proposes the idea of the self-support prototype, but their single spacial-agnostic prototype is generated for self-matching. Although also generated

from query branches, our multiple local prototypes carry discriminative and comprehensive spatial information that helps to provide details for intra-class discrimination.

2.3 Unlabeled Usage in Few-Shot Segmentation

Very few researchers have explored the leverage of unlabeled data in few-shot tasks. PpNet (Liu et al. 2020) is one of the cases, which incorporates 100 unlabeled data into every support input, with the use of a graph neural network (GNN). Based on it, Soopil proposes an uncertainty-aware model to make more adequate use of unlabeled data without using graph neural networks (Kim et al. 2023). However, the above two methods also require additional unlabeled data of the current novel class, which is inconsistent with the original few-shot task setting but relevant to the semi-supervised paradigm. And the number of unlabeled images used in every meta-training process far exceeds that of support and query images. In contrast, our work removes the dependence of the testing process on unlabeled data. During training, a tiny amount of unlabeled data is used to constrain the foreground semantic similarity.

3 Method

3.1 Problem Definition

The few-shot semantic segmentation task is defined as the segmentation of novel-class objects, based on a very small number of images with pixel-level annotations. According to this definition, the dataset for training D_{train} and the one for testing D_{test} have no overlap in their categories.

In each meta-training process of K -shot setting, conventional methods randomly sample $K + 1$ image pairs of the same class j : $\{(I^i, Y^i), i \in \{1, 2, \dots, K + 1\}\} | (I^i, Y^i) \in$

$D_{train}, C^i = j\}$. I^i, Y^i , and C^i refer to the RGB image, its segmentation mask, and the class of the segmented object. Among them, K pairs are sent to the support branch, and the last is used as the query input and ground truth. The difference in our network is that we take $K + 1 + M$ pairs of images from D_{train} at a time, and I^i of the extra M pairs are used as input to the unlabeled branch.

In each testing process, the $K + 1$ pairs from the j -th novel classes $\{(I^i, Y^i), i \in \{1, 2, \dots, K + 1\} | (I^i, Y^i) \in D_{test}, C^i = j\}$ are sampled and assigned to support and query branch as before.

3.2 Relevant Intrinsic Feature Enhancement

As a pixel-level binary classification task, semantic segmentation needs to identify the foreground, *i.e.* object to be segmented, and the background, *i.e.* the rest of the stuff and things. To alleviate the semantic ambiguity and inter-class similarity in Fig.1, we propose this relevant intrinsic feature enhancement network. We present the complete framework under the 1-shot setting in Fig. 2. The proposed RiFeNet consists of three branches with a shared backbone. The extra unlabeled branch helps the traditional support-query framework learn how to assure semantic consistency.

The forward process consists of three main modules: the multi-level prototype generation module, the multi-level prototype interaction module, and the feature activation module. The first two modules are used to provide multi-grained evidence for better inter-class discrimination. The feature activation module activates pixels containing objects of the target class and deactivates the others, providing the final segmentation result.

3.3 Feature Enhancement with Unlabeled Data

There are large intra-class distributional differences in the foreground objects that need to be segmented. Therefore, the training purpose is to ensure the semantic consistency of features extracted from different instances of the foreground category, rather than focusing only on semantically ambiguous appearance features.

Therefore, we introduce an auxiliary unlabeled branch as an effective data utilization method to aid model learning. Without adding training data, we re-sample a subset of training samples as unlabeled data, with the same segmentation loss applied. By augmenting the sample diversity, it teaches the model to avoid learning sample-specific biases of labeled inputs, even in the absence of unlabeled data during testing.

Our unlabeled branch shares parameters with the query branch to align features and enhance relevance. It is trained with pseudo-labels generated from each other. Specifically, the initial unlabeled input goes through two versions of data augmentation, different in their kinds and intensities. We use I_u and \hat{I}_u to refer to the weakly and strongly transformed unlabeled input. Both are sent to the backbone simultaneously with an identical forward process, thus we use symbols with a hat in all subsequent equations for simplicity.

$$F_u^m = ReLU(Conv(CAT(F^1(I_u), F^2(I_u)))). \quad (1)$$

F_u^m refers to merged unlabeled feature maps. For the sake of simplicity, we omit the above processing in Fig.2. The sup-




Figure 3: Visual illustration of the multi-level prototype generation block that extracts global and local prototypes.




Figure 4: Schedule of the multi-level prototype interaction merging relevant intrinsic information for enhancement.

port and query branches are handled in the same way, and all parameters in Eq.1 are shared with query branches. $F^j(I_u)$ corresponds to the feature map from the j -th backbone layer.

As for the prototype generation block, if we follow the query paradigm to generate local prototypes from the unlabeled branch, the self-supervised training may be off-target in a few extreme cases. In complex environments with multiple objects, the segmented class in the generated pseudo mask may not be inconsistent. This off-track pseudo-truth may worsen the expression capability of the model.

We introduce query prototypes here to provide class-specific prior. Only intrinsic features with intra-class relevance are extracted when processing unlabeled data.

The subsequent prototype interaction is the same as the query one. Augmented unlabeled features are obtained and sent to the activation block for prediction. S^u and \hat{Y}^u represent the predicted results of weakly and strongly transformed inputs, respectively. The former output S^u serves as a pseudo label for the latter one. The loss of this self-supervised process is calculated as follows, with dice loss (Milletari, Navab, and Ahmadi) as its training objective.

$$\mathcal{L}_{unlabel} = DICE(S^u, \hat{Y}^u). \quad (2)$$

3.4 Multi-Level Prototype Processing

While the unlabeled branch ensures intra-class semantic consistency of the foreground, the background noise it introduces may worsen existing inter-class similarity between foreground and background. Therefore, we propose the multi-level prototype generation and interaction blocks aim to capture representative information in different granularity and augment the discrimination of foreground and background in confusing areas.

Global Support Prototype Generation. To capture category information at the high-dimensional category level, we extract global prototypes from support features.

As shown in Fig.3, $P_s \in \mathbb{R}^{1 \times 1 \times C}$ is obtained through masked average pooling on the global feature. w and h are the width and height of F_s^m . \odot means dot production.

$$P_s = \frac{\sum_{x=1, y=1}^{w, h} (F_s^m \odot Y_s)^{x, y}}{\sum_{x=1, y=1}^{w, h} Y_s^{x, y}}. \quad (3)$$

Local Query Prototype Generation. In some cases, the similarities between different classes may be more significant than the differences within classes (Fan et al. 2022). Meanwhile, the prototypes obtained by global average pooling completely discard spatial structure information, making the decoupling of local components unrealistic.

To provide discriminative details, we additionally extract local prototypes from the query branch. These prototypes facilitate binary classification by prompting fine-grained information. Fig.3 illustrates their generation process.

$$P_q^* = \text{Pooling}_{avg}(F_q^m \odot M(F^3(I_q))). \quad (4)$$

$M(F^3(I_q))$ is the prior mask proposed in PFENet (Tian et al. 2020), as a form of similarity measure between high-level feature maps of support and query branches. Local average pooling is done on the multiplication of prior masks and query feature maps.

To squeeze the channel and enhance the representation of local prototypes, 1×1 convolution and channel-wise attention are used. The refined local prototypes $P_q \in \mathbb{R}^{m \times m \times C'}$ from the query branch itself are obtained as follows. The symbol Att_{chn} refers to channel-wise attention.

$$P_q = Att_{chn}(\text{Conv}^{1 \times 1}(P_q^*)). \quad (5)$$

For the generation of prior masks, fixed high-level feature maps F_s^4 and F_q^4 are more suitable. But when generating local query prototypes, superficial partly appearance information counts a lot. With the need for pixel-level correspondence, we use feature map F_q^1 from the first layer to extract locally consistent appearance information in the local area, as shown in Fig.3.

Multi-Level Prototype Interaction. To emphasize the intrinsic class information, we build interactions between different-grained prototypes, thus enhancing the feature-mining ability for identification.

Fig.4 shows detailed steps of our interaction block. The generated global and local prototypes are expanded to the size of feature maps and then concatenated with query features and the prior mask. After a 1×1 convolution with activation, augmented query features F_q^* are obtained.

$$F_q^* = \text{ReLU}(\text{Conv}(\text{CAT}(P_s, P_q, F_q^m, M(F^3(I_q)))). \quad (6)$$

3.5 Feature Activation

For deeper information interaction between the query and the support branch, we use an n-layer transformer encoder for feature activation. In this process, information with relevance is activated, enhancing the total feature map.

Augmented query features F_q^* are processed by self-attention and then cross-attention with support features F_s^* . Each block is made up of a multi-head attention and a feed-forward network. The setting of multi-head attention for cross-attention follows the one in CyCTR (Zhang et al. 2021), with the segmentation mask of the support image Y_s used to maintain the cycle consistency.

$$F_q^{final} = Att_{cross}(Att_{self}(F_q^*), F_s^*, Y_s). \quad (7)$$

Att_{cross} and Att_{self} represent the cross-attention and self-attention block, respectively. The output of the transformer F_q^{final} is resized and passed to the classification head to obtain the final pixel-by-pixel segmentation result, as shown in the following equation.

$$Y_q = \text{CLS}(F_q^{final} + \text{ReLU}(\text{Conv}(\text{CAT}(F_q^{final}))). \quad (8)$$

Y_q represents the predicted result of the query input and CLS refers to the segmentation head.

The loss of RiFeNet is calculated between predictions of query inputs and their corresponding ground truth. We choose dice loss as the loss function, thus the loss for each meta-learning task can be represented as:

$$\mathcal{L}_{main} = \text{DICE}(\hat{Y}^q, Y^q) + \mathcal{L}_{aux}. \quad (9)$$

where \hat{Y}^q and Y^q are the prediction and ground truth of query images. \mathcal{L}_{aux} refers to an auxiliary loss attached for feature alignment, following the common setting (Milletari, Navab, and Ahmadi 2016).

$$\mathcal{L}_{final} = \mathcal{L}_{main} + \beta \mathcal{L}_{unlabel}. \quad (10)$$

As the above equation shows, the final loss is a weighted sum of the main loss and the self-supervised loss $\mathcal{L}_{unlabel}$. The weight β is set to 0.5 empirically.

4 Experiments

We conduct experiments on the *PASCAL-5ⁱ* and *COCO* datasets. The primary evaluation metric is mean intersection-over-union (mIoU). The foreground-background IoU (FB-IoU) is used as an auxiliary indicator.

4.1 Process of Training and Testing

The number of support: query: unlabeled images is 1 : 1 : 2 or 5 : 1 : 2 for each meta-training task in the one-shot and five-shot settings. The unlabeled input branch comes from the same training dataset. We resample M extra images of the same class without mask truth, rather than adding another dataset. For a fair comparison, we exclude the unlabeled images during testing.

Pre-trained ResNet50 and ResNet101 (He et al. 2016) are used as the backbone of the model, frozen without parameter update. The baseline method refers to the CyCTR. Local query prototypes are generated by the multi-level prototype generation. Concatenated with one global support prototype, m^2 local query prototype, and the prior mask, query features are sent to the transformer after a merging and flattening operation. We set $m = 4$ in the main experiments. All data in the following tables are the average results of five experiments with the same settings. Specific implementation details are presented in the Appendix.

Method	Backbone	1-shot					5-shot					#learnable params
		split0	split1	split2	split3	mean	split0	split1	split2	split3	mean	
PPNet 2020	Res-50	52.7	62.8	57.4	47.7	55.2	60.3	70.0	69.4	60.7	65.1	31.5M
PMM 2020		52.0	67.5	51.5	49.8	55.2	55.0	68.2	52.9	51.1	56.8	-
PFENet 2020		61.7	69.5	55.4	56.3	60.8	63.1	70.7	55.8	57.9	61.9	10.8M
CyCTR 2021		65.7	71.0	59.5	<u>59.7</u>	64.0	69.3	73.5	63.8	63.5	67.5	7.4M
HSNet 2020		64.3	70.7	60.3	60.5	64.0	<u>70.3</u>	73.2	67.4	67.1	<u>69.5</u>	2.6M
ASGNet 2021		58.8	67.9	56.8	53.7	59.3	63.7	70.6	64.1	57.4	63.9	10.4M
SSP 2022		61.4	67.2	<u>65.4</u>	49.7	60.9	68.0	72.0	74.8	60.2	68.8	-
DCAMA 2022		<u>67.5</u>	<u>72.3</u>	59.6	59.0	<u>64.6</u>	70.5	<u>73.9</u>	63.7	<u>65.8</u>	68.5	-
RiFeNet (Ours)		68.4	73.5	67.1	59.4	67.1	70.0	74.7	<u>69.4</u>	64.2	69.6	7.7M
DAN 2020		Res-101	54.7	68.6	57.8	51.6	58.2	57.9	69.0	60.1	54.9	60.5
PMM 2020	54.7		68.6	57.8	51.6	58.2	57.9	69.0	60.1	54.9	60.5	-
PFENet 2020	60.5		69.4	54.4	55.9	60.1	62.8	70.4	54.9	57.6	61.4	10.8M
CyCTR 2021	67.2		71.1	57.6	59.0	63.7	<u>71.0</u>	<u>75.0</u>	58.5	65.0	67.4	7.4M
HSNet 2020	<u>67.3</u>		<u>72.3</u>	62.0	63.1	<u>66.2</u>	71.8	74.4	67.0	68.3	70.4	2.6M
ASGNet 2021	59.8		67.4	55.6	54.4	59.3	64.6	71.3	64.2	57.3	64.4	10.4M
SSP 2022	63.7		70.1	66.7	55.4	64.0	70.3	76.3	77.8	<u>65.5</u>	72.5	-
DCAMA 2022	65.4		71.4	63.2	58.3	64.6	70.7	73.7	66.8	61.9	68.3	-
RiFeNet (Ours)	68.9		73.8	<u>66.2</u>	<u>60.3</u>	67.3	70.4	74.5	<u>68.3</u>	63.4	69.2	7.7M

Table 1: Performance comparison on PASCAL-5ⁱ in terms of mIoU (%).

Method	Backbone	1-shot					5-shot				
		split0	split1	split2	split3	mean	split0	split1	split2	split3	mean
PPNet 2020	Res-50	28.1	30.8	29.5	27.7	29.0	39.0	40.8	37.1	37.3	38.5
PMM 2020		29.3	34.8	27.1	27.3	29.6	33.0	40.6	30.3	33.3	34.3
RPMMs 2020		29.5	36.8	28.9	27.0	30.6	33.8	42.0	33.0	33.3	35.5
CyCTR 2021		38.9	43.0	39.6	39.8	40.3	41.1	48.9	45.2	<u>47.0</u>	45.6
HSNet 2020		36.3	43.1	38.7	38.7	39.2	43.3	<u>51.3</u>	48.2	45.0	46.9
SSP 2022		46.4	35.2	27.3	25.4	33.6	53.8	41.5	36.0	33.7	41.3
DCAMA 2022		<u>41.9</u>	<u>45.1</u>	<u>44.4</u>	<u>41.7</u>	<u>43.3</u>	45.9	50.5	50.7	46.0	<u>48.3</u>
RiFeNet (Ours)		39.1	47.2	44.6	45.4	44.1	<u>44.3</u>	52.4	<u>49.3</u>	48.4	48.6

Table 2: Performance comparison on COCO in terms of mIoU (%).

4.2 Comparison with State-of-the-Arts

We compare the performance of RiFeNet with other classical or effective methods on $PASCAL - 5^i$, as is shown in Tab.1 and Tab. A of the Appendix. RiFeNet outperforms the best method under most of the experimental scenarios. RiFeNet outperforms CyCTR by about 3.5% under the 1-shot setting and about 2% for the 5-shot one. Compared with the existing state-of-the-art DCAMA in the 1-shot setting with ResNet50 backbone, it surpasses by 2.5%, rising to 2.7% with the use of ResNet101. As for the larger gain in the 1-shot than in the 5-shot setting, we attribute this to the decreasing impact weight due to the constant amount of unlabeled data, as the ratio of unlabeled to labeled images decreases from 2 to 0.4. As labeled images increase, the positive effects of the unlabeled branch decrease, with a proportional decline in performance gain in 5-shot.

Similar experiments on COCO support the above conclusion in Tab.2. Faced with a scenario with multiple objects in this dataset and a complex environment, RiFeNet still outperforms the current best DCAMA by 0.8% for almost all splits in the 1-shot setting. The comparison results demonstrate the benefits of RiFeNet. The unlabeled branch provides RiFeNet with richer relevant information, which in

turn improves the performance of the model.

Qualitative results also prove the effectiveness of RiFeNet. In Fig.5 and Fig.A of the Appendix, the foreground objects in support and query images vary a lot, with inconsistent postures, appearances, and angles of photography. Despite this large intra-class variability, RiFeNet achieves significant improvement in maintaining foreground semantic consistency compared with the baseline. As for the similarity of background and foreground, the model deals with this binary identification much better even in cases with neighboring objects of similar appearance, with foreground occlusion, and with multiple classes of objects. Looking back to Fig.1, our extracted features are essential to maintain foreground semantic consistency and provide inter-class distinction for binary classification.

4.3 Ablation Studies

RiFeNet improves the pixel-level binary classification. To demonstrate the effectiveness of our proposed unlabeled enhancement and multi-level prototypes in RiFeNet, we conduct diagnostic experiments in Tab.3. All comparisons are set under a 1-shot setting, with ResNet50 as the backbone. Using either the unlabeled branch or multi-level prototype

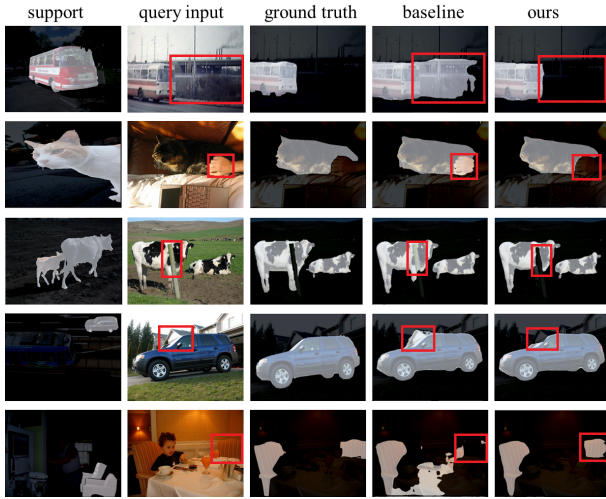


Figure 5: Qualitative segmentation results on novel classes on PASCAL-5ⁱ. From left to right: support image with mask, query input, query ground-truth mask, query prediction of the baseline, and prediction of RiFeNet.

Un	MP	split0	split1	split2	split3	mIoU
		65.7	71.0	59.5	59.7	64.0
✓		67.3	71.8	66.2	59.2	66.1
	✓	66.0	72.1	66.2	60.4	66.2
✓	✓	68.4	73.5	67.1	59.4	67.1

Table 3: Ablation studies on the key components of RiFeNet. “Un” and “MP” denote the use of the unlabeled branch and the multi-level prototypes, respectively.

interaction results in a boost of approximately 2%. When two strategies work together, RiFeNet improves by 3.1% on top of the baseline.

Different design choices of multi-level prototypes. We conduct ablation experiments on the model design details mentioned for the multi-level prototype, as is shown in Tab.4. Consistent with the theoretical analysis in Sec.3.4, it proves that our practices such as adding guidance to unlabeled branches are reasonable and reliable.

Different design choices of unlabeled branches. We conduct experiments with different designed unlabeled branches to further explore their effect. As shown in Tab.5, the unlabeled branch without guided query prototypes results in even worse performance than the baseline, which is consistent with our analysis in Sec.3.3. On the other hand, because the unlabeled inputs come from resampling the training dataset, we double the training iterations of the baseline for a fair comparison. Increased training iterations have little effect on the baseline due to early convergence. This proves that the effectiveness of our method is not from the multiple sampling of data but from the learned discriminative and semantic features.

Different hyper-parameters. We first look into the effect of different numbers of unlabeled input in a single meta-

components	split0	split1	split2	split3	mIoU
gp (support-only)	67.3	71.8	66.2	59.2	66.1
gp+gp	67.5	73.1	66.2	58.4	66.3
gp+lp (w/o CA)	68.1	73.2	66.7	59.1	66.8
gp+lp (w/ CA)	68.4	73.5	67.1	59.4	67.1

Table 4: Ablation studies on multi-level prototypes. “gp” and “lp” denote global and local prototypes, respectively. That is, “gp+gp” means extracting both query and support prototypes globally. “CA” refers to channel-wise attention.

components	epoch	split0	split1	split2	split3	mIoU
w/o unlabel	200	66.0	72.1	66.2	60.4	66.2
w/o unlabel	400	66.5	72.4	65.5	59.5	66.0
un (w/o guide)	200	66.9	72.2	65.9	58.3	65.8
un (w/ guide)	200	68.4	73.5	67.1	59.4	67.1

Table 5: Ablation studies on the unlabeled branch. “w/ guide” refers to the use of query local prototypes in the unlabeled branch for guidance, while “w/o guide” means using prototypes generated from the unlabeled branch itself.

num	split0	split1	split2	split3	mIoU
0	66.0	72.1	66.2	60.4	66.2
1	66.8	72.8	66.9	59.8	66.6
2	68.4	73.5	67.1	59.4	67.1
3	65.9	72.6	66.9	59.8	66.0

Table 6: Ablation studies of different numbers of unlabeled images in the single meta-training process.

training process. Tab.6 shows the results on PASCAL-5ⁱ under a 1-shot setting, with ResNet50 as its backbone. The best results are obtained when the number of unlabeled images is set to 2. Initially, the segmentation effect of the model increased as the number of unlabeled images increased. When the number continues to increase, the accuracy decreases instead. We deem the reason is that when the effect of unlabeled enhancement counts much more than the query branch itself, the attention of feature mining may turn to the unlabeled branch, thus disturbing the query prediction. The segmentation accuracy decreases after the features are blurred. We also conduct detailed ablation experiments with other parameters, which are included in the Appendix.

5 Conclusion

In few-shot segmentation, traditional methods suffer from semantic ambiguity and inter-class similarity. Thus from the perspective of pixel-level binary classification, we propose RiFeNet, an effective model with an unlabeled branch constraining foreground semantic consistency. Without extra data, this unlabeled branch improves the in-class generalization of the foreground. Moreover, we propose to further enhance the discrimination of background and foreground by a multi-level prototype generation and interaction module.

References

- Chen, L.-C.; Zhu, Y.; Papandreou, G.; Schroff, F.; and Adam, H. 2018. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European Conference on Computer Vision*.
- Dong, N.; and Xing, E. P. 2018. Few-shot semantic segmentation with prototype learning. In *BMVC*, volume 3.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Fan, Q.; Pei, W.; Tai, Y.-W.; and Tang, C.-K. 2022. Self-support few-shot semantic segmentation. In *Proceedings of the European Conference on Computer Vision*. Springer.
- Fu, J.; Liu, J.; Tian, H.; Li, Y.; Bao, Y.; Fang, Z.; and Lu, H. 2019. Dual attention network for scene segmentation. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 3146–3154.
- He, J.; Deng, Z.; Zhou, L.; Wang, Y.; and Qiao, Y. 2019. Adaptive pyramid context network for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7519–7528.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 770–778.
- Hong, S.; Cho, S.; Nam, J.; Lin, S.; and Kim, S. 2022. Cost aggregation with 4d convolutional swin transformer for few-shot segmentation. In *Proceedings of the European Conference on Computer Vision*. Springer.
- Jin, Z.; Gong, T.; Yu, D.; Chu, Q.; Wang, J.; Wang, C.; and Shao, J. 2021. Mining contextual information beyond image for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7231–7241.
- Kang, D.; and Cho, M. 2022. Integrative few-shot learning for classification and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9979–9990.
- Kim, S.; Chikontwe, P.; An, S.; and Park, S. H. 2023. Uncertainty-aware semi-supervised few shot segmentation. *Pattern Recognition*, 109292.
- Lee, Y.; Kim, J.; Willette, J.; and Hwang, S. J. 2022. Mpvit: Multi-path vision transformer for dense prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7287–7296.
- Li, G.; Jampani, V.; Sevilla-Lara, L.; Sun, D.; Kim, J.; and Kim, J. 2021. Adaptive prototype learning and allocation for few-shot segmentation. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 8334–8343.
- Li, X.; Zhong, Z.; Wu, J.; Yang, Y.; Lin, Z.; and Liu, H. 2019. Expectation-maximization attention networks for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9167–9176.
- Lin, G.; Milan, A.; Shen, C.; and Reid, I. 2017. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 1925–1934.
- Liu, Y.; Liu, N.; Yao, X.; and Han, J. 2022. Intermediate prototype mining transformer for few-shot semantic segmentation. *arXiv preprint arXiv:2210.06780*.
- Liu, Y.; Zhang, X.; Zhang, S.; and He, X. 2020. Part-aware prototype network for few-shot semantic segmentation. In *Proceedings of the European Conference on Computer Vision*. Springer.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 3431–3440.
- Milletari, F.; Navab, N.; and Ahmadi, S.-A. 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision*, 565–571. Ieee.
- Min, J.; Kang, D.; and Cho, M. 2021. Hypercorrelation squeeze for few-shot segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6941–6952.
- Qin, J.; Wu, J.; Li, M.; Xiao, X.; Zheng, M.; and Wang, X. 2022a. Multi-granularity distillation scheme towards lightweight semi-supervised semantic segmentation. In *European Conference on Computer Vision*, 481–498. Springer.
- Qin, J.; Wu, J.; Xiao, X.; Li, L.; and Wang, X. 2022b. Activation modulation and recalibration scheme for weakly supervised semantic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 2117–2125.
- Qin, J.; Wu, J.; Yan, P.; Li, M.; Yuxi, R.; Xiao, X.; Wang, Y.; Wang, R.; Wen, S.; Pan, X.; et al. 2023. FreeSeg: Unified, Universal and Open-Vocabulary Image Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19446–19455.
- Shaban, A.; Bansal, S.; Liu, Z.; Essa, I.; and Boots, B. 2017. One-shot learning for semantic segmentation. *arXiv preprint arXiv:1709.03410*.
- Shi, X.; Wei, D.; Zhang, Y.; Lu, D.; Ning, M.; Chen, J.; Ma, K.; and Zheng, Y. 2022. Dense cross-query-and-support attention weighted mask aggregation for few-shot segmentation. In *Proceedings of the European Conference on Computer Vision*. Springer.
- Strudel, R.; Garcia, R.; Laptev, I.; and Schmid, C. 2021. Segmenter: Transformer for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7262–7272.
- Tao, A.; Sapra, K.; and Catanzaro, B. 2020. Hierarchical multi-scale attention for semantic segmentation. *arXiv preprint arXiv:2005.10821*.
- Tian, Z.; Zhao, H.; Shu, M.; Yang, Z.; Li, R.; and Jia, J. 2020. Prior guided feature enrichment network for few-shot segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 44(2): 1050–1065.

- Wang, H.; Zhang, X.; Hu, Y.; Yang, Y.; Cao, X.; and Zhen, X. 2020. Few-shot semantic segmentation with democratic attention networks. In *Proceedings of the European Conference on Computer Vision*. Springer.
- Wang, K.; Liew, J. H.; Zou, Y.; Zhou, D.; and Feng, J. 2019. Panet: Few-shot image semantic segmentation with prototype alignment. In *proceedings of the IEEE/CVF International Conference on Computer Vision*, 9197–9206.
- Wang, W.; Xie, E.; Li, X.; Fan, D.-P.; Song, K.; Liang, D.; Lu, T.; Luo, P.; and Shao, L. 2021. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 568–578.
- Xie, E.; Wang, W.; Yu, Z.; Anandkumar, A.; Alvarez, J. M.; and Luo, P. 2021. SegFormer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34: 12077–12090.
- Xiong, Z.; Li, H.; and Zhu, X. X. 2022. Doubly Deformable Aggregation of Covariance Matrices for Few-Shot Segmentation. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XX*, 133–150. Springer.
- Yang, B.; Liu, C.; Li, B.; Jiao, J.; and Ye, Q. 2020. Prototype mixture models for few-shot semantic segmentation. In *Proceedings of the European Conference on Computer Vision*. Springer.
- Yang, M.; Yu, K.; Zhang, C.; Li, Z.; and Yang, K. 2018. Denseaspp for semantic segmentation in street scenes. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 3684–3692.
- Yu, C.; Wang, J.; Gao, C.; Yu, G.; Shen, C.; and Sang, N. 2020. Context prior for scene segmentation. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 12416–12425.
- Yuan, Y.; Fu, R.; Huang, L.; Lin, W.; Zhang, C.; Chen, X.; and Wang, J. 2021. Hrformer: High-resolution vision transformer for dense predict. *Advances in Neural Information Processing Systems*, 34: 7281–7293.
- Yuan, Y.; Huang, L.; Guo, J.; Zhang, C.; Chen, X.; and Wang, J. 2018. Ocnet: Object context network for scene parsing. *arXiv preprint arXiv:1809.00916*.
- Zhang, C.; Lin, G.; Liu, F.; Yao, R.; and Shen, C. 2019a. Canet: Class-agnostic segmentation networks with iterative refinement and attentive few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5217–5226.
- Zhang, F.; Chen, Y.; Li, Z.; Hong, Z.; Liu, J.; Ma, F.; Han, J.; and Ding, E. 2019b. Acfnets: Attentional class feature network for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6798–6807.
- Zhang, G.; Kang, G.; Yang, Y.; and Wei, Y. 2021. Few-shot segmentation via cycle-consistent transformer. *Advances in Neural Information Processing Systems*, 34: 21984–21996.
- Zhang, H.; Dana, K.; Shi, J.; Zhang, Z.; Wang, X.; Tyagi, A.; and Agrawal, A. 2018. Context encoding for semantic segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 7151–7160.
- Zhang, X.; Wei, Y.; Yang, Y.; and Huang, T. S. 2020. Sg-one: Similarity guidance network for one-shot semantic segmentation. *IEEE transactions on cybernetics*, 50(9): 3855–3865.
- Zhao, H.; Shi, J.; Qi, X.; Wang, X.; and Jia, J. 2017. Pyramid scene parsing network. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2881–2890.
- Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2020. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*.
- Zhu, Z.; Xu, M.; Bai, S.; Huang, T.; and Bai, X. 2019. Asymmetric non-local neural networks for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 593–602.