

# Prototype Mixture Models for Few-shot Semantic Segmentation

Boyu Yang, Chang Liu, Bohao Li, Jianbin Jiao, and Qixiang Ye\*

University of Chinese Academy of Sciences, Beijing, China  
{yangboyu18, liuchang615}@mailsucas.ac.cn, libohao1998@gmail.com  
{jiaojb, qxye}@ucas.ac.cn

**Abstract.** Few-shot segmentation is challenging because objects within the support and query images could significantly differ in appearance and pose. Using a single prototype acquired directly from the support image to segment the query image causes semantic ambiguity. In this paper, we propose prototype mixture models (PMMs), which correlate diverse image regions with multiple prototypes to enforce the prototype-based semantic representation. Estimated by an Expectation-Maximization algorithm, PMMs incorporate rich channel-wised and spatial semantics from limited support images. Utilized as representations as well as classifiers, PMMs fully leverage the semantics to activate objects in the query image while depressing background regions in a duplex manner. Extensive experiments on Pascal VOC and MS-COCO datasets show that PMMs significantly improve upon state-of-the-arts. Particularly, PMMs improve 5-shot segmentation performance on MS-COCO by up to 5.82% with only a moderate cost for model size and inference speed.<sup>1</sup>

**Keywords:** Semantic Segmentation, Few-shot Segmentation, Few-shot Learning, Mixture Models

## 1 Introduction

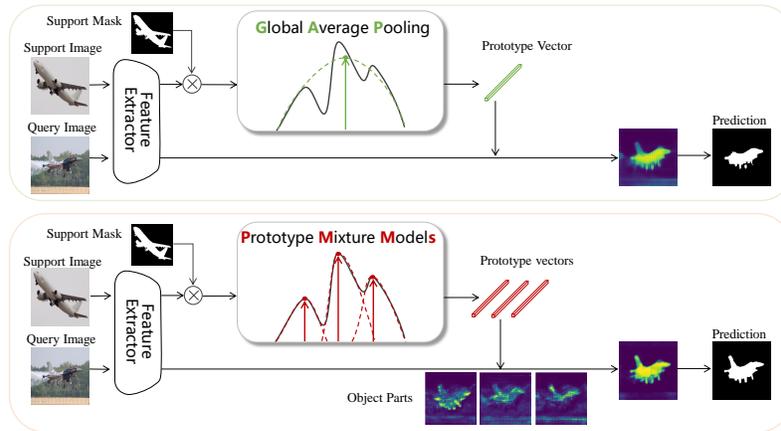
Substantial progress has been made in semantic segmentation [1,2,3,4,5,6,7,8]. This has been broadly attributed to the availability of large datasets with mask annotations and convolutional neural networks (CNNs) capable of absorbing the annotation. However, annotating object masks for large-scale datasets is laborious, expensive, and can be impractical [9,10,11]. It is also not consistent with cognitive learning, which builds a model upon few-shot supervision [12].

Given a few examples, termed *support* images, and the related segmentation masks [13], few-shot segmentation aims to segment the *query* images based on a feature representation learned on training images. It remains a challenging problem when we consider that target category is not included in the training

---

<sup>1</sup> Code is available at [github.com/Yang-Bob/PMMs](https://github.com/Yang-Bob/PMMs).

\*Qixiang Ye is the corresponding author.



**Fig. 1.** The single prototype model (upper) based on global average pooling causes semantic ambiguity about object parts. In contrast, prototype mixture models (lower) correlate diverse image regions, *e.g.*, object parts, with multiple prototypes to enhance few-shot segmentation model. (Best viewed in color)

data while objects within the support and query image significantly differ in appearance and pose.

By introducing the metric learning framework, Shaban *et al.* [14], Zhang *et al.* [15], and Dong *et al.* [16] contributed early few-shot semantic segmentation methods. They also introduced the concept of “prototype” which refers to a weight vector calculated with global average pooling guided by ground-truth masks embedded in feature maps. Such a vector squeezing discriminative information across feature channels is used to guide the feature comparison between support image(s) and query images for semantic segmentation.

Despite clear progress, we argue that the commonly used prototype model is problematic when the spatial layout of objects is completely dropped by global average pooling, Fig. 1(upper). A single prototype causes semantic ambiguity around various object parts and deteriorates the distribution of features [17]. Recent approaches have alleviated this issue by prototype alignment [18], feature boosting [13], and iterative mask refinement [19]. However, the semantic ambiguity problem caused by global average pooling remains unsolved.

In this paper, we propose prototype mixture models (PMMs) and focus on solving the semantic ambiguity problem in a systematic manner. During the training procedure, the prototypes are estimated using an Expectation-Maximization (EM) algorithm, which treats each deep pixel (a feature vector) within the mask region as a positive sample. PMMs are primarily concerned with representing the diverse foreground regions by estimating mixed prototypes for various object parts, Fig. 1(lower). They also enhance the discriminative capacity of features by modeling background regions.

The few-shot segmentation procedure is implemented in a metric learning framework with two network branches (a support branch and a query branch), Fig. 2. In the framework, PMMs are utilized in a duplex manner to segment a query image. On the one hand, they are regarded as spatially squeezed representation, which match (P-Match) with query features to activate feature channels related to the object class. On the other hand, each vector is regarded as a  $C$ -dimensional linear classifier, which multiplies (P-Conv) with the query features in an element-wised manner to produce a probability map. In this way, the channel-wised and spatial semantic information of PMMs is fully explored to segment the query image.

The contributions of our work are summarized as follows:

- We propose prototype mixture models (PMMs), with the target to enhance few-shot segmentation by fully leveraging semantics of limited support image(s). PMMs are estimated using an EM algorithm, which is integrated with feature learning by a plug-and-play manner.
- We propose a duplex strategy, which treats PMMs as representations and classifiers, to activate spatial and channel-wised semantics for segmentation.
- We assemble PMMs to RPMMs using a residual structure and significantly improve upon the state-of-the-arts.

## 2 Related work

**Semantic Segmentation.** Semantic segmentation, which performs per-pixel classification of a class of objects, has been extensively investigated. State-of-the-art methods, such as UNet [2], PSPNet [1], DeepLab [3,4,5], are based on fully convolutional networks (FCNs) [20]. Semantic segmentation has been updated to instance segmentation [8] and panoptic segmentation [21], which shared useful modules, *e.g.*, Atrous Spatial Pyramid Pooling (ASPP) [4] and multi-scale feature aggregation [1], with few-shot segmentation. The clustering method used in SegSort [7], which partitioned objects into parts using a divide-and-conquer strategy, provides an insight for this study.

**Few-shot Learning.** Existing methods can be broadly categorized as either: metric learning [22,23,17], meta-learning [24,25,26,27], or data argumentation. Metric learning based methods train networks to predict whether two images/regions belong to the same category. Meta-learning based approaches specify optimization or loss functions which force faster adaptation of the parameters to new categories with few examples. The data argumentation methods learn to generate additional examples for unseen categories [28,29].

In the metric learning framework, the effect of prototypes for few-shot learning has been demonstrated. With a simple prototype, *e.g.*, a linear layer learned on top of a frozen CNN [30], state-of-the-art results can be achieved based on a simple baseline. This provides reason for applying prototypes to capture representative and discriminative features.

**Few-shot Segmentation.** Existing few-shot segmentation approaches largely followed the metric learning framework, *e.g.*, learning knowledge using a prototype vector, from a set of support images, and then feed learned knowledge to a metric module to segment query images [18].

In OSLSM [14], a two-branch network consisting of a support branch and a query branch was proposed for few-shot segmentation. The support branch is devoted to generating a model from the support set, which is then used to tune the segmentation process of an image in the query branch. In PL [16], the idea of prototypical networks was employed to tackle few-shot segmentation using metric learning. SG-One [15] also used a prototype vector to guide semantic segmentation procedure. To obtain the squeezed representation of the support image, a masked average pooling strategy is designed to produce the prototype vector. A cosine similarity metric is then applied to build the relationship between the guidance features and features of pixels from the query image. PANet [18] further introduced a prototype alignment regularization between support and query branches to fully exploit knowledge from support images for better generalization. CANet [19] introduced a dense comparison module, which effectively exploits multiple levels of feature discriminativeness from CNNs to make dense feature comparison. With this approach comes an iterative optimization module which refines segmentation masks. The FWB approach [13] focused on discriminativeness of prototype vectors (support features) by leveraging foreground-background feature differences of support images. It also used an ensemble of prototypes and similarity maps to handle the diversity of object appearances.

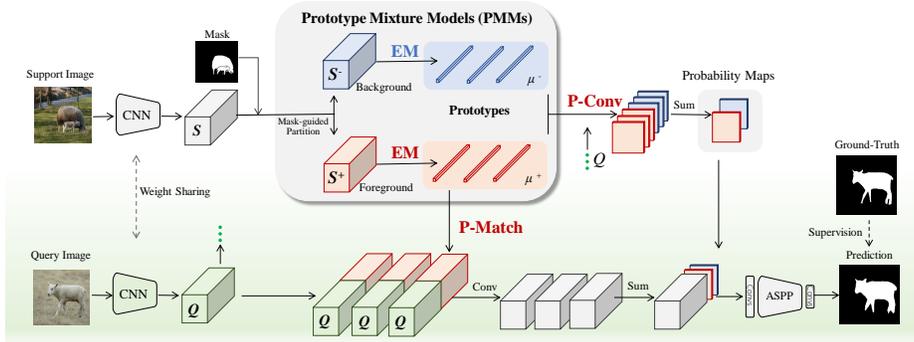
As a core of metric learning in few-shot segmentation, the prototype vector was commonly calculated by global average pooling. However, such a strategy typically disregards the spatial extent of objects, which tends to mix semantics from various parts. This unintended mixing seriously deteriorates the diversity of prototype vectors and feature representation capacity. Recent approaches alleviated this problem using iterative mask refinement [19] or model ensemble [13]. However, issues remain when using single prototypes to represent object regions and the semantic ambiguity problem remains unsolved.

Our research is inspired by the prototypical network [31], which learns a metric space where classification is performed using distances to the prototype of each class. The essential differences are twofold: (1) A prototype in prototypical network [31] represents a class of samples while a prototype in our approach represents an object part; (2) The prototypical network does not involve mixing prototypes for a single sample or a class of samples.

### 3 The Proposed Approach

#### 3.1 Overview

The few-shot segmentation task is to classify pixels in query images into foreground objects or backgrounds by solely referring to few labeled support images containing objects of the same categories. The goal of the training procedure is



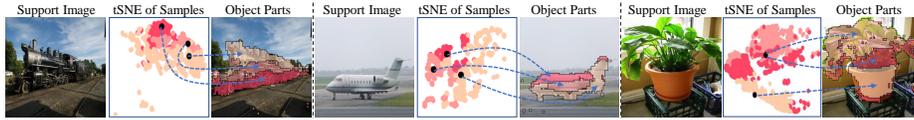
**Fig. 2.** The proposed approach consists of two branches *i.e.*, the support branch and the query branch. During training, the feature set  $S$  of a support image is partitioned into a positive sample set  $S^+$  and a negative sample set  $S^-$  guided by the ground-truth mask.  $S^+$  and  $S^-$  are respectively used to train  $\mu^+$  and  $\mu^-$ , which are used to activate query features in a duplex way (P-Conv and P-Match) for semantic segmentation. “ASPP” refers to the Atrous Spatial Pyramid Pooling (ASPP) [4].

to learn a segmentation model that is trained by numbers of images different from the task query image categories. The training image set is split into many small subsets and within every subset one image serves as the query and the other(s) as the support image(s) with known ground-truth(s). Once the model is trained, the segmentation model is fixed and requires no optimization when tested on a new dataset [19]. The proposed few-shot segmentation model follows a metric learning framework, Fig. 2, which consists of two network branches *i.e.*, the support branch (above) and the query branch (below). Over the support branch, PMMs are estimated for the support image(s). In the support and query branches, two CNNs with shared weights are used as the backbone to extract features. Let  $S \in \mathcal{R}^{W \times H \times C}$  denote the features of the support image where  $W \times H$  denotes the resolution of feature maps and  $C$  the number of feature channels. The features for a query image are denoted as  $Q \in \mathcal{R}^{W \times H \times C}$ .

Without loss of generality, the network architecture and models are illustrated for 1-shot setting, which can be extended to 5-shot setting by feeding five support images to the PMMs to estimate prototypes.

### 3.2 Prototype Mixture Models

During training, features  $S \in \mathcal{R}^{W \times H \times C}$  for the support image are considered as a sample set with  $W \times H$   $C$ -dimensional samples.  $S$  is spatially partitioned into foreground samples  $S^+$  and background samples  $S^-$ , where  $S^+$  corresponds to feature vectors within the mask of the support image and  $S^-$  those outside the mask.  $S^+$  is used to learn foreground PMMs<sup>+</sup> corresponding to object parts, Fig. 3, and  $S^-$  to learn background PMMs<sup>-</sup>. Without loss generality, the models



**Fig. 3.** Foreground sample distribution of support images. The black points on tSNE maps denote positive prototypes correlated to object parts. (Best viewed in color)

and learning procedure are defined for PMMs, which represent either  $\text{PMMs}^+$  or  $\text{PMMs}^-$ .

**Models.** PMMs are defined as a probability mixture model which linearly combine probabilities from base distributions, as

$$p(s_i|\theta) = \sum_{k=1}^K w_k p_k(s_i|\theta) \quad (1)$$

where  $w_k$  denotes the mixing weights satisfying  $0 \leq w_k \leq 1$  and  $\sum_{k=1}^K w_k = 1$ .  $\theta$  denotes the model parameters which are learned when estimating PMMs.  $s_i \in S$  denotes the  $i^{\text{th}}$  feature sample and  $p_k(s_i|\theta)$  denotes the  $k^{\text{th}}$  base model, which is a probability model based on a Kernel distance function, as

$$p_k(s_i|\theta) = \beta(\theta) e^{Kernel(s_i, \mu_k)}, \quad (2)$$

where  $\beta(\theta)$  is the normalization constant.  $\mu_k \in \theta$  is one of the parameter. For the Gaussian mixture models (GMMs) with fixed co-variance, the Kernel function is a radial basis function (RBF),  $Kernel(s_i, \mu_k) = -\|(s_i - \mu_k)\|_2^2$ . For the von Missies-Fisher (VMF) model [32], the kernel function is defined as a cosine distance function, as  $Kernel(s_i, \mu_k) = \frac{\mu_k^T s_i}{\|\mu_k\|_2 \|s_i\|_2}$ , where  $\mu_k$  is the mean vector of the  $k^{\text{th}}$  model. Considering the metric learning framework used, the vector distance function is more appropriate in our approach, as is validated in experiments. Based on the vector distance, PMMs are defined as

$$p_k(s_i|\theta) = \beta_c(\kappa) e^{\kappa \mu_k^T s_i}, \quad (3)$$

where  $\theta = \{\mu, \kappa\}$ .  $\beta_c(\kappa) = \frac{\kappa^{c/2-1}}{(2\pi)^{c/2} I_{c/2-1}(\kappa)}$  is the normalization coefficient, and  $I_\nu(\cdot)$  denotes the Bessel function.  $\kappa$  denotes the concentration parameter, which is empirically set as  $\kappa = 20$  in experiments.

**Model Learning.** PMMs are estimated using the EM algorithm which includes iterative E-steps and M-steps. In each E-step, given model parameters and sample features extracted, we calculate the expectation of the sample  $s_i$  as

$$E_{ik} = \frac{p_k(s_i|\theta)}{\sum_{k=1}^K p_k(s_i|\theta)} = \frac{e^{\kappa \mu_k^T s_i}}{\sum_{k=1}^K e^{\kappa \mu_k^T s_i}}. \quad (4)$$

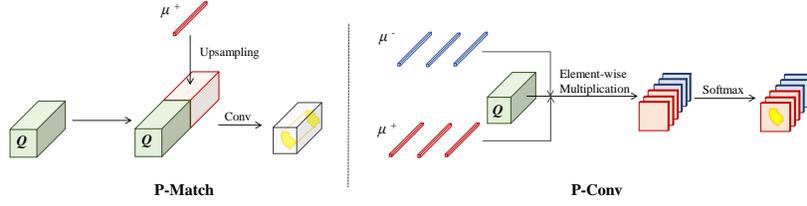


Fig. 4. Illustration of P-Match and P-Conv operations for query feature ( $Q$ ) activation.

---

**Algorithm 1** Learning PMMs for few-shot segmentation

---

**Input:**

Support images, mask  $M_s$  for each support image, query image;

**Output:**

Network parameters  $\alpha$ , prototypes  $\mu^+$  and  $\mu^-$  for each support image;

**for** (each support image) **do**

**Calculate** support and query features  $S$  and  $Q$  by forward propagation with  $\theta$ ;

**Partition**  $S$  into  $S^+$  and  $S^-$  according to  $M_s$ ;

**Learn** PMMs;

        Estimate  $\mu^+$  upon  $S^+$  by iterative EM steps defined by Eqs. 4 and 5;

        Estimate  $\mu^-$  upon  $S^-$  by iterative EM steps defined by Eqs. 4 and 5;

**Activate**  $Q$  using P-Match and P-Conv defined on  $\mu^+$  and  $\mu^-$ , Fig. 4;

**Predict** a segmentation mask and calculate the segmentation loss;

**Update**  $\alpha$  to minimize the cross-entropy loss at the query branch, Fig. 2.

**end for**

---

In each M-step, the expectation is used to update PMMs' mean vectors, as

$$\mu_k = \frac{\sum_{i=1}^N E_{ik} s_i}{\sum_{i=1}^N E_{ik}}, \quad (5)$$

where  $N = W \times H$  denotes the number of samples.

After model learning, the mean vectors  $\mu^+ = \{\mu_k^+, k = 1, \dots, K\}$  and  $\mu^- = \{\mu_k^-, k = 1, \dots, K\}$  are used as prototype vectors to extract convolutions features for the query image. The mixture coefficient  $w_k$  is ignored so that each prototype vectors have same importance for semantic segmentation. Obviously, each prototype vector is the mean of a cluster of samples. Such a prototype vector can represent a region around an object part in the original image for the reception field effect, Fig. 3.

### 3.3 Few-shot Segmentation

During inference, the learned prototype vectors  $\mu^+ = \{\mu_k^+, k = 1, \dots, K\}$  and  $\mu^- = \{\mu_k^-, k = 1, \dots, K\}$  are duplexed to activate query features for semantic segmentation, Fig. 2.

**PMMs as Representation (P-Match).** Each positive prototype vector squeezes representative information about an object part and all prototypes in-

corporate representative information about the complete object extent. Therefore, prototype vectors can be used to match and activate the query features  $Q$ , as

$$Q' = \text{P-Match}(\mu_k^+, Q), k = 1, \dots, K, \quad (6)$$

where P-Match refers to an activation operation consists of prototype upsampling, feature concatenation, and semantic segmentation using convolution, Fig. 4. The convolution operation on concatenated features implements a channel-wise comparison, which activates feature channels related to foreground while suppressing those associated with backgrounds. With P-Match, semantic information about the extent of the complete object is incorporated into the query features for semantic segmentation.

**PMMs as Classifiers (P-Conv).** On the other hand, each prototype vector incorporating discriminative information across feature channels can be seen as a classifier, which produces probability maps  $M_k = \{M_k^+, M_k^-\}$  using the P-Conv operation, as

$$M_k = \text{P-Conv}(\mu_k^+, \mu_k^-, Q), k = 1, \dots, K. \quad (7)$$

As shown in Fig. 4, P-Conv first multiplies each prototype vector with the query feature  $Q$  in an element-wise manner. The output maps are then converted to probability maps  $M_k$  by applying Softmax across channels.

After P-Conv, the produced probability maps  $M_k^+, k = 1, \dots, K$  and  $M_k^-, k = 1, \dots, K$  are respectively summarized to two probability maps, as

$$\begin{aligned} M_p^+ &= \sum_k M_k^+, \\ M_p^- &= \sum_k M_k^-, \end{aligned} \quad (8)$$

which are further concatenated with the query features to activate objects of interest, as

$$Q'' = M_p^+ \oplus M_p^- \oplus Q', \quad (9)$$

where  $\oplus$  denotes the concatenation operation.

After the P-Match and P-Conv operations, the semantic information across channels and discriminative information related to object parts are collected from the support feature  $S$  to activate the query feature  $Q$ . in a dense comparison manner [19]. The activated query features  $Q''$  are further enhanced with Atrous Spatial Pyramid Pooling (ASPP) and fed to a convolutional module to predict the segmentation mask, Fig. 2.

**Segmentation Model Learning.** The segmentation model is implemented as an end-to-end network, Fig. 2. The learning procedure for the segmentation model is described in Algorithm 1. In the feed forward procedure, the support features are partitioned into backgrounds and foreground sample sets  $S^+$  and  $S^-$  according to the ground-truth mask. PMMs are learned on  $S^+$  and  $S^-$  and the learned prototype vectors  $\mu^+$  and  $\mu^-$  are leveraged to activate query features to

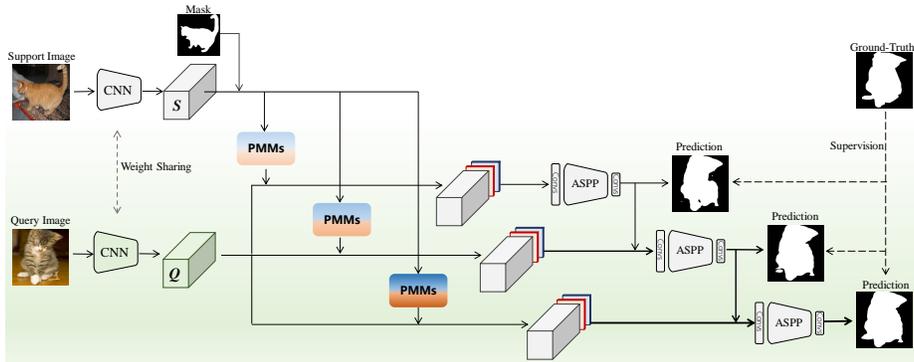


Fig. 5. Network architecture of residual prototype mixture models (RPMMS).

predict segmentation mask of the query image. In the back-propagation procedure, the network parameters  $\theta$  are updated to optimize the segmentation loss at the query branch. With multiple training iterations, rich feature representation about diverse object parts is absorbed into the backbone network. During the inference procedure, the learned feature representation together with PMMs of the support image(s) is used to segment the query image.

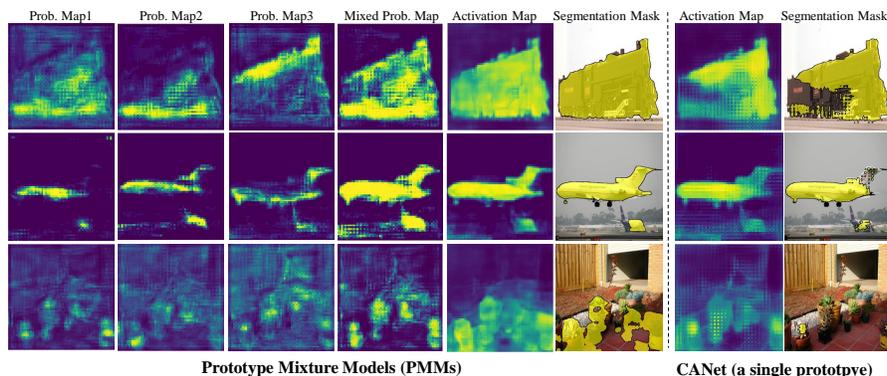
### 3.4 Residual Prototype Mixture Models

To further enhance the model capacity, we implement model ensemble by stacking multiple PMMs, Fig. 5. Stacked PMMs, termed residual PMMs (RPMMS), leverage the residual from the previous query branch to supervise the next query branch for fine-grained segmentation. RPMMS not only further improve the performance but also defines a new model ensemble strategy. This incorporates the advantages of model residual learning, which is inspired by the idea of side-output residual [33] but has the essential difference to handle models instead of features. This is also different from the ensemble of experts [13], which generates an ensemble of the support features guided by the gradient of loss.

## 4 Experiments

### 4.1 Experimental settings

**Implementation Details.** Our approach utilizes CANet [19] without iterative optimization as the baseline, which uses VGG16 or ResNet50 as backbone CNN for feature extraction. During training, four data augmentation strategies including normalization, horizontal flipping, random cropping and random resizing are used [19]. Our approach is implemented upon the PyTorch 1.0 and run on Nvidia 2080Ti GPUs. The EM algorithm iterates 10 rounds to calculate PMMs for each image. The network with a cross-entropy loss is optimized by SGD with the



**Fig. 6.** Activation maps by PMMs and CANet [19]. PMMs produce multiple probability maps and fuse them to a mixed map, which facilitates activating and segmenting complete object extent (first two rows) or multiple objects (last row). CANet that uses a single prototype to segment object tends to miss object parts. (Best viewed in color)

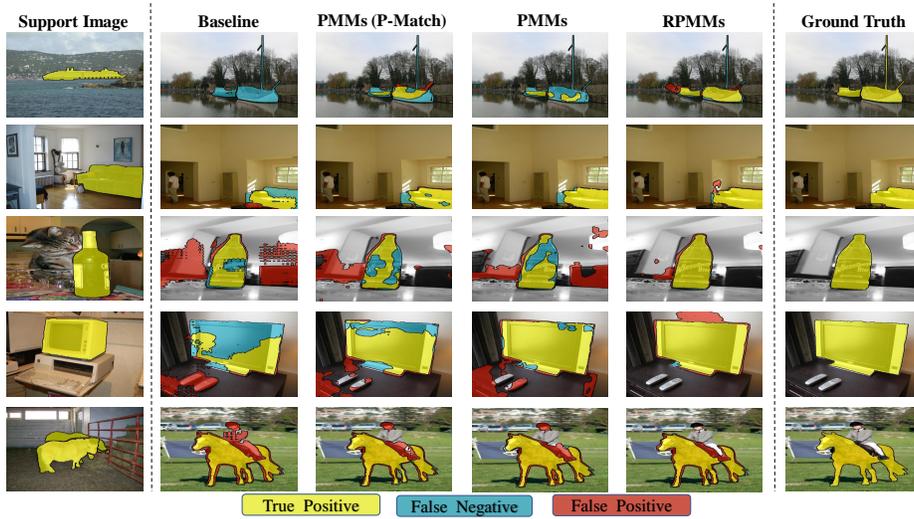
initial learning of 0.0035 and the momentum of 0.9 for 200,000 iterations with 8 pairs of support-query images per batch. The learning rate reduces following the “poly” policy defined in DeepLab [4]. For each training step, the categories in the train split are randomly selected and then the support-query pairs are randomly sampled in the selected categories.

**Datasets.** We evaluate our model on Pascal-5<sup>i</sup> and COCO-20<sup>i</sup>. Pascal-5<sup>i</sup> is a dataset specified for few-shot semantic segmentation in OSLSM [14], which consists of the Pascal VOC 2012 dataset with extra annotations from extended SDS [34]. 20 object categories are partitioned into four splits with three for training and one for testing. At test time, 1000 support-query pairs were randomly sampled in the test split [19]. Following FWB [13], we create COCO-20<sup>i</sup> from MSCOCO 2017 dataset. The 80 classes are divided into 4 splits and each contains 20 classes and the *val* dataset is used for performance evaluation. The other setting is the same as that in Pascal-5<sup>i</sup>.

**Evaluation Metric.** Mean intersection over-union (mIoU) which is defined as the mean IoUs of all image categories was employed as the metric for performance evaluation. For each category, the IoU is calculated by  $\text{IoU} = \frac{\text{TP}}{\text{TP} + \text{FP} + \text{FN}}$ , where TP, FP and FN respectively denote the number of true positive, false positive and false negative pixels of the predicted segmentation masks.

## 4.2 Model Analysis

In Fig. 6, we visualize probability maps produced by positive prototypes of PMMs. We also visualize and compare the activation maps and segmentation masks produced by PMMs and CANet. PMMs produce multiple probability maps and fuse them to a mixed probability map, which facilitates activating complete object extent (first two rows). The advantage in terms of representation



**Fig. 7.** Semantic segmentation results. ‘Baseline’ refers to the CANet method [19] without iterative optimization. (Best viewed in color)

**Table 1.** Ablation study. ‘Mean’ denotes mean mIoU on Pascal-5<sup>t</sup> with PMMs using three prototypes. The first row is the baseline method without using the PMMs or the RPMMs method.

PMMs <sup>+</sup> (P-Match)	PMMs(P-Match & P-Conv)	RPMMs	Mean
			51.93
✓			54.63
✓	✓		55.27
✓	✓	✓	56.34

capacity is that PMMs perform better than CANet when segmenting multiple objects within the same image (last row). By comparison, CANet using a single prototype to activate object tends to miss object parts or whole objects. The probability maps produced by PMMs validate our idea, *i.e.*, prototypes correlated to multiple regions and alleviate semantic ambiguity.

In Fig. 7, we compare the segmentation results by the baseline method and the proposed modules. The segmentation results show that PMMs<sup>+</sup>(P-Match) can improve the recall rate by segmenting more target pixels. By introducing background prototypes, PMMs reduce the false positive pixels, which validates that the background mixture models can improve the discriminative capability of the model. RPMMs further improve the segmentation results by refining object boundaries about hard pixels.

**Table 2.** Performance (mIoU%) on prototype number  $K$ .

$K$	Pascal-5 <sup>0</sup>	Pascal-5 <sup>1</sup>	Pascal-5 <sup>2</sup>	Pascal-5 <sup>3</sup>	Mean
1	49.38	66.42	51.29	47.68	53.69
2	50.85	66.65	<b>51.89</b>	48.25	54.41
3	51.88	66.72	51.14	<b>48.80</b>	<b>54.63</b>
4	<b>51.89</b>	<b>66.96</b>	51.36	47.91	54.53
5	50.76	66.89	50.76	47.94	54.09

**Table 3.** Performance comparison of Kernel functions.

Kernel	Pascal-5 <sup>0</sup>	Pascal-5 <sup>1</sup>	Pascal-5 <sup>2</sup>	Pascal-5 <sup>3</sup>	Mean
Gaussian	50.94	66.70	50.59	47.91	54.04
<b>VMF</b>	<b>51.88</b>	<b>66.72</b>	<b>51.14</b>	<b>48.80</b>	<b>54.63</b>

### 4.3 Ablation Study

**PMMs.** In Table 1, with P-Match modules, PMMs improve segmentation performance by 2.70% (54.63% vs. 51.93%), which validates that the prototypes generated by the PMMs perform better than the prototype generated by global average pooling. By introducing the duplex strategy, PMMs further improve the performance by 0.64% (55.27% vs. 54.63%), which validates that the probability map generated by the combination of foreground and background prototypes can suppress backgrounds and reduce false segmentation. In total, PMMs improve the performance by 3.34% (55.27% vs. 51.93%), which is a significant margin in semantic segmentation. This clearly demonstrates the superiority of the proposed PMMs over previous prototype methods.

**RPMMs.** RPMMs further improve the performance by 1.07% (56.34% vs. 55.27%), which validates the effectiveness of the residual ensemble strategy. Residual from the query prediction output of the previous branch of PMMs can be used to supervise the next branch of PMMs, enforcing the stacked PMMs to reduce errors, step by step.

**Number of Prototypes.** In Table 2, ablation study is carried out to determine the number of prototypes using PMMs<sup>+</sup> with P-Match.  $K = 2$  significantly outperforms  $K = 1$ . The best Pascal-5 <sup>$i$</sup>  performance occurs at  $K = 2, 3, 4$ . When  $K = 3$  the best mean performance is obtained. When  $K = 4, 5$ , the performance slightly decreases. One reason lies in that the PMMs are estimated on a single support image, which includes limited numbers of samples. The increase of  $K$  substantially decreases the samples of each prototype and increases the risk of over-fitting.

**Kernel Functions.** In Table 3, we compare the Gaussian and VMF kernels for sample distance calculation when estimating PMMs. The better results from VMF kernel show that the cosine similarity defined by VMF kernel is preferable.

**Inference Speed.** The size of PMMs model is 19.5M, which is slightly larger than that of the baseline CANet [19] (19M) but much smaller than that of

**Table 4.** Performance of 1-way 1-shot semantic segmentation on Pascal-5<sup>i</sup>. CANet reports multi-scale test performance.

Backbone	Method	Pascal-5 <sup>0</sup>	Pascal-5 <sup>1</sup>	Pascal-5 <sup>2</sup>	Pascal-5 <sup>3</sup>	Mean
VGG16	OSLSM [14]	33.60	55.30	40.90	33.50	40.80
	co-FCN [35]	36.70	50.60	44.90	32.40	41.10
	SG-One [15]	40.20	58.40	48.40	38.40	46.30
	PANet [18]	42.30	58.00	51.10	41.20	48.10
	FWB [13]	47.04	59.64	<b>52.51</b>	48.27	51.90
	<b>RPMMs (ours)</b>	<b>47.14</b>	<b>65.82</b>	50.57	<b>48.54</b>	<b>53.02</b>
Resnet50	CANet [19]	49.56	64.97	49.83	<b>51.49</b>	53.96
	<b>PMMs (ours)</b>	51.98	<b>67.54</b>	51.54	49.81	55.22
	<b>RPMMs (ours)</b>	<b>55.15</b>	66.91	<b>52.61</b>	50.68	<b>56.34</b>

**Table 5.** Performance of 1-way 5-shot semantic segmentation on Pascal-5<sup>i</sup>.

Backbone	Method	Pascal-5 <sup>0</sup>	Pascal-5 <sup>1</sup>	Pascal-5 <sup>2</sup>	Pascal-5 <sup>3</sup>	Mean
VGG16	OSLSM [14]	35.90	58.10	42.70	39.10	43.95
	SG-One [15]	41.90	58.60	48.60	39.40	47.10
	FWB [13]	50.87	62.86	56.48	<b>50.09</b>	55.08
	<b>PANet [18]</b>	<b>51.80</b>	64.60	<b>59.80</b>	46.05	<b>55.70</b>
	RPMMs (ours)	50.00	<b>66.46</b>	51.94	47.64	54.01
Resnet50	CANet [19]	-	-	-	-	55.80
	<b>PMMs (ours)</b>	55.03	<b>68.22</b>	52.89	<b>51.11</b>	56.81
	<b>RPMMs (ours)</b>	<b>56.28</b>	67.34	<b>54.52</b>	51.00	<b>57.30</b>

**Table 6.** Performance of 1-shot and 5-shot semantic segmentation on MS COCO. FWB uses the ResNet101 backbone while other approaches use the ResNet50 backbone.

Settings	Method	COCO-20 <sup>0</sup>	COCO-20 <sup>1</sup>	COCO-20 <sup>2</sup>	COCO-20 <sup>3</sup>	Mean
1-shot	PANet [18]	-	-	-	-	20.90
	FWB [13]	16.98	17.98	20.96	28.85	21.19
	Baseline	25.08	30.25	24.45	24.67	26.11
	<b>PMMs (ours)</b>	29.28	34.81	27.08	27.27	29.61
	<b>RPMMs (ours)</b>	<b>29.53</b>	<b>36.82</b>	<b>28.94</b>	<b>27.02</b>	<b>30.58</b>
5-shot	FWB [13]	19.13	21.46	23.93	30.08	23.65
	PANet [18]	-	-	-	-	29.70
	Baseline	25.95	32.38	26.11	26.98	27.86
	<b>PMMs (ours)</b>	33.00	40.55	30.29	33.27	34.28
	<b>RPMMs (ours)</b>	<b>33.82</b>	<b>41.96</b>	<b>32.99</b>	<b>33.33</b>	<b>35.52</b>

OSLSM [14] (272.6M). Because the prototypes are  $1 \times 1 \times C$  dimensional vectors, they do not significantly increase the model size or computational cost. In one shot setting, with  $K = 3$ , our inference speed on single 2080Ti GPU is 26 FPS,

which is slightly lower than that of CANet (29 FPS). With RPMMs the speed decreases to 20 FPS while the model size (19.6M) does not significantly increase.

#### 4.4 Performance

**PASCAL-5<sup>i</sup>.** In Table 4 and Table 5, PMMs and RPMMs are compared with the state-of-the-art methods. They outperform state-of-the-art methods in both 1-shot and 5-shot settings. With the 1-shot setting and a Resnet50 backbone, RPMMs achieve 2.38% (56.34% vs. 53.96%) performance improvement over the state-of-the-art, which is a significant margin.

With the 5-shot setting and a Resnet50 backbone, RPMMs achieve 1.50% (57.30% vs. 55.80%) performance improvement over the state-of-the-art, which is also significant. With the VGG16 backbone, our approach is comparable with the state-of-the-arts. Note that the PANet and FWB used additional  $k$ -shot fusion strategy while we do not use any post-processing strategy to fuse the predicted results from five shots.

**MS COCO.** Table 6 displays the evaluation results on MS COCO dataset following the evaluation metric on COCO-20<sup>i</sup> [13]. Baseline is achieved by running CANet without iterative optimization. PMMs and RPMMs again outperform state-of-the-art methods in both 1-shot and 5-shot settings. For the 1-shot setting, RPMMs improves the baseline by 4.47%, respectively outperforms the PANet and FWB methods by 9.68% and 9.39%.

For the 5-shot setting, it improves the baseline by 7.66%, and respectively outperforms the PANet and FWB by 5.82% and 11.87%, which are large margins for the challenging few-shot segmentation problem. Compared to PASCAL VOC, MS COCO has more categories and images for training, which facilitates learning richer representation related to various object parts and backgrounds. Thereby, the improvement on MS COCO is larger than that on Pascal VOC.

## 5 Conclusion

We proposed prototype mixture models (PMMs), which correlate diverse image regions with multiple prototypes to solve the semantic ambiguity problem. During training, PMMs incorporate rich channel-wised and spatial semantics from limited support images. During inference, PMMs are matched with query features in a duplex manner to perform accurate semantic segmentation. On the large-scale MS COCO dataset, PMMs improved the performance of few-shot segmentation, in striking contrast with state-of-the-art approaches. As a general method to capture the diverse semantics of object parts given few support examples, PMMs provide a fresh insight for the few-shot learning problem.

**Acknowledgement:** This work was supported in part by the National Natural Science Foundation of China (NSFC) under Grant 61836012, 61671427, and 61771447.

## References

1. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: IEEE CVPR. (2017) 6230–6239
2. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI. Volume 9351. (2015) 234–241
3. Chen, L., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Semantic image segmentation with deep convolutional nets and fully connected crfs. In: ICLR. (2015) 6230–6239
4. Chen, L., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell.* **40**(4) (2018) 834–848
5. Chen, L., Papandreou, G., Schroff, F., Adam, H.: Rethinking atrous convolution for semantic image segmentation. *CoRR* [abs/1706.05587](https://arxiv.org/abs/1706.05587) (2017)
6. Chen, L., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: ECCV. (2018) 833–851
7. Hwang, J., Yu, S.X., Shi, J., Collins, M.D., Yang, T., Zhang, X., Chen, L.: Segsort: Segmentation by discriminative sorting of segments. In: IEEE ICCV. (2019) 7334–7344
8. He, K., Gkioxari, G., Dollár, P., Girshick, R.B.: Mask R-CNN. *IEEE Trans. Pattern Anal. Mach. Intell.* **42**(2) (2020) 386–397
9. Wan, F., Liu, C., Ke, W., Ji, X., Jiao, J., Ye, Q.: C-mil: Continuation multiple instance learning for weakly supervised object detection. In: IEEE CVPR. (2019) 2199–2208
10. Wan, F., Wei, P., Han, Z., Jiao, J., Ye, Q.: Min-entropy latent model for weakly supervised object detection. *IEEE Trans. Pattern Anal. Machine Intell.* **41**(10) (2019) 2395–2409
11. Zhang, X., Wan, F., Liu, C., Ji, R., Ye, Q.: Freeanchor: Learning to match anchors for visual object detection. In: NeurIPS. (2019) 147–155
12. Tokmakov, P., Wang, Y., Hebert, M.: Learning compositional representations for few-shot recognition. In: IEEE ICCV. (2019) 6372–6381
13. Nguyen, K., Todorovic, S.: Feature weighting and boosting for few-shot segmentation. In: IEEE ICCV. (2019) 622–631
14. Shaban, A., Bansal, S., Liu, Z., Essa, I., Boots, B.: One-shot learning for semantic segmentation. In: BMVC. (2017)
15. Zhang, X., Wei, Y., Yang, Y., Huang, T.: Sg-one: Similarity guidance network for one-shot semantic segmentation. *CoRR* [abs/1810.09091](https://arxiv.org/abs/1810.09091) (2018)
16. Dong, N., Xing, E.P.: Few-shot semantic segmentation with prototype learning. In: BMVC. (2018) 79
17. Hao, F., He, F., Cheng, J., Wang, L., Cao, J., Tao, D.: Collect and select: Semantic alignment metric learning for few-shot learning. In: IEEE ICCV. (2019) 8460–8469
18. Wang, K., Liew, J., Zou, Y., Zhou, D., Feng, J.: Panet: Few-shot image semantic segmentation with prototype alignment. (2019) 622–631
19. Zhang, C., Lin, G., Liu, F., Yao, R., Shen, C.: Canet: Class-agnostic segmentation networks with iterative refinement and attentive few-shot learning. In: IEEE CVPR. (2019) 5217–5226
20. Shelhamer, E., Long, J., Darrell, T.: Fully convolutional networks for semantic segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**(4) (2017) 640–651

21. Kirillov, A., Girshick, R., He, K., Dollar, P.: Panoptic feature pyramid networks. In: IEEE CVPR. (2019) 6399–6408
22. Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., Wierstra, D.: Matching networks for one shot learning. In: NeurIPS. (2016) 3630–3638
23. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H.S., Hospedales, T.M.: Learning to compare: Relation network for few-shot learning. In: IEEE CVPR. (2018) 1199–1208
24. Wang, Y., Hebert, M.: Learning to learn: Model regression networks for easy small sample learning. In: ECCV. (2016) 616–634
25. Ravi, S., Larochelle, H.: Optimization as a model for few-shot learning. In: ICLR. (2017)
26. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: ICML. (2017) 1126–1135
27. Jamal, M.A., Qi, G.J.: Task agnostic meta-learning for few-shot learning. In: IEEE ICCV. (2019) 111719–111727
28. Hariharan, B., Girshick, R.B.: Low-shot visual recognition by shrinking and hallucinating features. In: IEEE ICCV. (2017) 3037–3046
29. Wang, Y., Girshick, R.B., Hebert, M., Hariharan, B.: Low-shot learning from imaginary data. In: IEEE CVPR. (2018) 7278–7286
30. Chen, W.Y., Liu, Y.C., Kira, Z., Yu-Chiang Wang: A closer look at few-shot classification. In: IEEE ICLR. (2019)
31. Snell, J., Swersky, K., Zemel, R.S.: Prototypical networks for few-shot learning. In: NeurIPS. (2017) 4077–4087
32. Banerjee, A., Dhillon, I.S., Ghosh, J., Sra, S.: Clustering on the unit hypersphere using von mises-fisher distributions. *J. Mach. Learn. Res.* **6** (2005) 1345–1382
33. Ke, W., Chen, J., Jiao, J., Zhao, G., Ye, Q.: SRN: side-output residual network for object symmetry detection in the wild. In: IEEE CVPR. (2017) 302–310
34. Hariharan, B., Arbelaez, P., Bourdev, L.D., Maji, S., Malik, J.: Semantic contours from inverse detectors. In: IEEE ICCV. (2011) 991–998
35. Rakelly, K., Shelhamer, E., Darrell, T., Efros, A.A., Levine, S.: Conditional networks for few-shot semantic segmentation. In: ICLR Workshop. (2018)