

Differentiable Meta-Learning Model for Few-Shot Semantic Segmentation

Pinzhuo Tian,¹ Zhangkai Wu,¹ Lei Qi,¹ Lei Wang,² Yinghuan Shi,^{1*} Yang Gao¹

¹National Key Laboratory for Novel Software Technology, Nanjing University, China

²School of Computing and Information Technology, University of Wollongong, Australia

{tianpinzhuo, qilei.cs}@gmail.com wuzhangkai@smail.nju.edu.cn leiw@uow.edu.au {syh, gaoy}@nju.edu.cn

Abstract

To address the annotation scarcity issue in some cases of semantic segmentation, there have been a few attempts to develop the segmentation model in the few-shot learning paradigm. However, most existing methods only focus on the traditional 1-way segmentation setting (*i.e.*, one image only contains a single object). This is far away from practical semantic segmentation tasks where the K -way setting ($K > 1$) is usually required by performing the accurate multi-object segmentation. To deal with this issue, we formulate the few-shot semantic segmentation task as a learning-based pixel classification problem, and propose a novel framework called MetaSegNet based on meta-learning. In MetaSegNet, an architecture of embedding module consisting of the global and local feature branches is developed to extract the appropriate meta-knowledge for the few-shot segmentation. Moreover, we incorporate a linear model into MetaSegNet as a base learner to directly predict the label of each pixel for the multi-object segmentation. Furthermore, our MetaSegNet can be trained by the episodic training mechanism in an end-to-end manner from *scratch*. Experiments on two popular semantic segmentation datasets, *i.e.*, PASCAL VOC and COCO, reveal the effectiveness of the proposed MetaSegNet in the K -way few-shot semantic segmentation task.

Introduction

Recently, deep learning has made significant breakthroughs in many applications (*e.g.*, image classification (He et al. 2016), object detection (Girshick 2015) and semantic segmentation (Long, Shelhamer, and Darrell 2015)). However, its main challenge is that a large amount of labeled data is usually required to train deep models, which is impractical in real-world applications. Meanwhile, the accurate labeling is extremely laborious and expensive, particularly for pixel-wise annotation (*e.g.*, semantic segmentation) and 3-D delineation (*e.g.*, medical tissue labeling). Therefore, few-shot learning has recently drawn an increasing interest in the machine learning community, which can aid deep models to effectively learn knowledge from a few samples.

Formally, few-shot learning aims to achieve good performance on the new classes with only a few labeled training data. Recently, most existing methods employ the meta-learning algorithm to deal with the few-shot learning problem (Snell, Swersky, and Zemel 2017; Finn, Abbeel, and

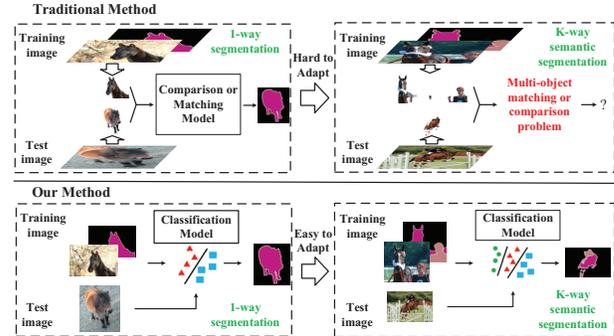


Figure 1: Comparison between our framework and traditional methods. The traditional methods use comparison or matching model in the 1-way segmentation task to measure the similarity between training and test image, which is hard to extend to K -way semantic segmentation task. Differently, our method leverages learning-based model to classify the pixels in the test image for K -way semantic segmentation.

Levine 2017), and obtain incredible success in the classification task. However, there are only a few methods that attempt to develop few-shot learning algorithms for a more difficult task (*e.g.*, semantic segmentation).

Most existing methods for few-shot semantic segmentation usually introduce the extra knowledge to acquire abundant prior knowledge to help the few-shot learning. However, they still suffer from several limitations as:

1. **Distribution divergence.** Most existing methods require their model to be pre-trained on ImageNet (Deng et al. 2009) as a prerequisite. However, the distribution of these images in the few-shot segmentation task might be very different from ImageNet.
2. **Hard to extend.** The existing methods usually solve the few-shot segmentation by utilizing metric-learning methods to measure the one-to-one similarity between training and test images in the 1-way setting. However, these methods cannot be directly adapted to measure multi-object similarity in the K -way setting.
3. **Difficult to train.** These existing methods usually employ complex embedding models to learn feature representation for image comparison, which are hard to optimize. Thus, if we directly train them from scratch, they

*Corresponding author

cannot work effectively in the target task.

To address the aforementioned limitations, we propose a novel framework for few-shot semantic segmentation termed as MetaSegNet. Different from existing works, we formulate the few-shot segmentation as a dense prediction task, which classifies each pixel and is readily extended to the K -way setting ($K > 1$). Specifically, we integrate a linear classification model into our framework as classifier, which could guarantee that our framework can be effectively optimized by the episodic training mechanism with meta-learning paradigm. Furthermore, a novel embedding module is designed in our framework to combine the global context and local information for better segmentation. Compared with two separate embedding modules in existing works, our embedding module is simple but effective. Meanwhile, the linear classification model combined with the new embedding module allows our framework to be effectively trained by an end-to-end manner from scratch for the multi-object semantic segmentation.

We evaluate our framework on two semantic segmentation datasets, *i.e.*, PASCAL VOC 2012 (Hariharan et al. 2015) and COCO 2014 (Lin et al. 2014), in the K -way, N -shot semantic segmentation task. Thanks to the effectiveness and efficiency of linear classifier and embedding module, our MetaSegNet can achieve the state-of-the-art performance in K -way, N -shot semantic segmentation task by using extremely small deep network without any extra knowledge. Our main contributions are summarized as follows.

- We formulate the few-shot semantic segmentation from the classification perspective, and propose a novel framework for dealing with the K -way, N -shot semantic segmentation problem.
- As far as we know, we are the first one to leverage the linear classifier instead of nonlinear layer to effectively train deep models in few-shot semantic segmentation.
- A simple embedding module with novel architecture is proposed, which can extract global and local information for semantic segmentation.
- The proposed novel framework is trained in an end-to-end manner from scratch. Also, it does not need any prior knowledge, which is the merit for the few-shot learning.

Related Work

In this part, we review the related works including meta-learning and few-shot segmentation.

Meta-learning for few-shot Learning. Meta-learning studies what aspects of the learner (commonly referred to bias or prior) effect generalization across a distribution of tasks (Vilalta and Drissi 2002). Nowadays, meta-learning approaches for few-shot learning can be broadly divided into three categories. Basically, (1) the *metric-based methods* learn a sophisticated comparison model conditioned on distance or metric between training and test data. For example, the matching networks (Vinyals et al. 2016) develop an end-to-end differentiable nearest neighbor to perform comparison. Snell *et al.* (Snell, Swersky, and Zemel 2017) propose a prototypical network which represents each category

by the mean embedding of the examples, and utilize cosine distance to measure the similarity between test images and the prototypes. Li *et al.* (Li et al. 2019) design a covariance metric network for distance metric. (2) The *model-based methods* wish to learn a parameterized predictor to estimate model parameters for new tasks. For example, Ravi *et al.* (Ravi and Larochelle 2017) utilize LSTM (Hochreiter and Schmidhuber 1997) to predict parameters analogous to a few steps of gradient descent in parameter space. (3) The *optimization-based methods* employ the optimization-based methods (*e.g.*, gradient descent) to directly adapt the parameters to study knowledge across task. For example, Finn *et al.* (Finn, Abbeel, and Levine 2017) propose MAML algorithm aiming to learn a good initialization in a new task. Franceschi *et al.* (Franceschi et al. 2018) use bi-level optimization method to optimize meta-learning algorithm.

Few-shot segmentation. Previous methods for few-shot segmentation generally use the metric-based meta-learning framework, meanwhile two-branch separate embedding modules pre-trained on the ImageNet are used to compare the similarity between training and test images. These methods can be divided into two categories, according to the task setting. (1) 1-way segmentation. Shaban *et al.* (Shaban et al. 2017) first consider the few-shot learning problem in segmentation, and propose the two-branch comparison structure. Hu *et al.* (Hu et al. 2019) use attention and multi-context to enhance the two-branch comparison model. Also, Zhang *et al.* (Zhang et al. 2019) introduce a two-branch class-agnostic segmentation network which contains a dense comparison module and an iterative optimization module. (2) K -way semantic segmentation. Dong *et al.* (Dong and Xing 2018) and Wang *et al.* (Wang et al. 2019) employ prototypes to represent the different semantic categories in the K -way semantic segmentation. However, they leverage different modules to align and measure the similarity between different objects.

Differently, our approach employs the optimization based meta-learning for few-shot segmentation. The linear classifier is integrated into our framework to deal with pixel classification. Different from 1-way segmentation methods, our framework is easily extended to K -way semantic segmentation setting. Moreover, compared with (Dong and Xing 2018; Wang et al. 2019), our framework is much simpler so that it could be optimized by an end-to-end manner without auxiliary information or data.

Our Method

In this section, we first introduce the K -way, N -shot segmentation task. Then we formulate the meta-learning framework. Finally, we elaborate each component in our method.

K -way, N -shot Few-shot Semantic Segmentation

In this paper, we consider the K -way, N -shot segmentation task which is seldom investigated ($K > 1$). Here, K denotes the number of classes, and N is the number of training examples per class. Especially, in the few-shot learning setting, N is usually small, *e.g.*, $N \in \{1, 5\}$. The training task $\mathcal{T}_i = (D_i^{\text{train}}, D_i^{\text{test}})$ in meta-training set is sampled from a

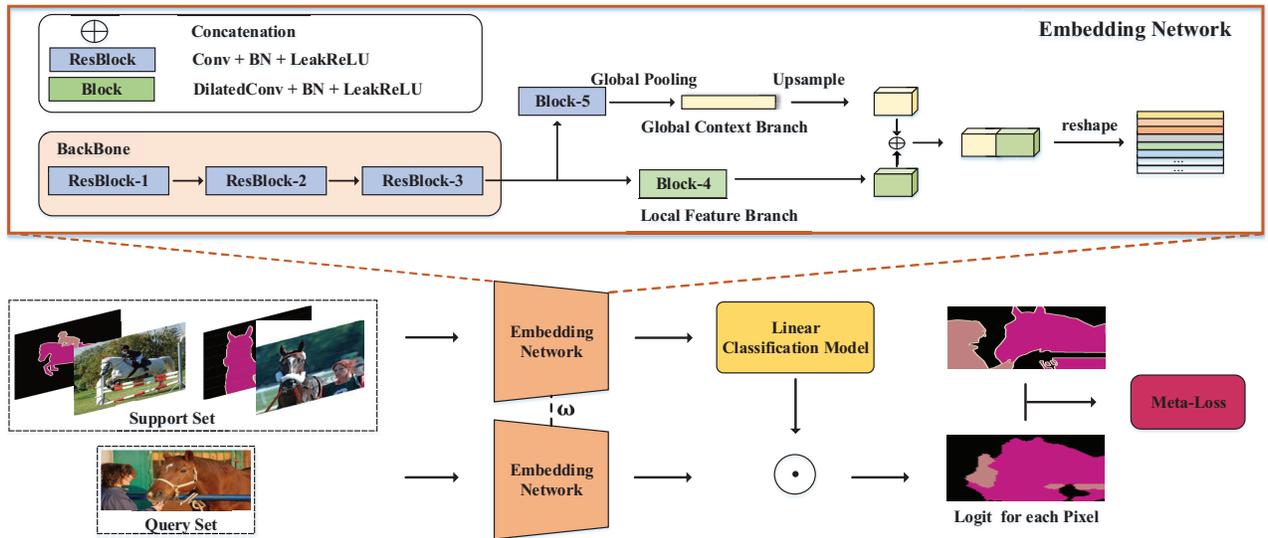


Figure 2: Pipeline of MetaSegNet for 2-way semantic segmentation. In this framework, the support set is used to train a linear classification model, and we make use of meta-loss to train an embedding model, which can generalize well across tasks. Particularly, we put forward a novel embedding model, where two branches are utilized to extract local and global features for the pixel-level classification.

large-scale annotated dataset that contains multiple different training categories (*e.g.*, we use C^{train} to represent the set of training classes). The training categories should be non-overlapping with the new classes (*e.g.*, C^{novel}) in meta-testing set, *i.e.*, $C^{\text{train}} \cap C^{\text{novel}} = \emptyset$. Firstly, the K categories $C_i = \{C_i^k\}_{k=1}^K$ in $\mathcal{D}_i^{\text{train}}$ are randomly selected from C^{train} . Then, N images per category are sampled to constitute the whole support (training) set $\mathcal{D}_i^{\text{train}}$ (*i.e.*, $\mathcal{D}_i^{\text{train}} = \{(x_n, y_n) \mid n = 1, \dots, N \times K, y_n \in C_i\}$). Finally, the query (test) set $\mathcal{D}_i^{\text{test}} = \{(x_n, y_n) \mid n = 1, \dots, Q \times K, y_n \in C_i\}$ consists of Q images per category.

However, most exiting few-shot segmentation methods follow the paradigm and notations in (Shaban et al. 2017), which is a special case of K -way, N -shot semantic segmentation, *i.e.*, $K = 1$. Although the model can forward K times to finish K -way semantic segmentation task, these models will suffer from the multi-object alignment problem. For example, one object may be classified to different semantic categories during K -time running, which obstructs traditional 1-way segmentation model achieving good performance in the K -way setting.

Problem Formulation for Meta-learning

Formally, given M tasks, we denote the training task set (*i.e.*, meta-training set) as $\mathcal{T} = \{(\mathcal{D}_i^{\text{train}}, \mathcal{D}_i^{\text{test}})\}_{i=1}^M$, where the tuple $(\mathcal{D}_i^{\text{train}}, \mathcal{D}_i^{\text{test}})$ represents the training data (*i.e.*, support set) and test data (*i.e.*, query set) for the i -th task \mathcal{T}_i . The goal of meta-learning is to enable a learning algorithm to effectively adapt to new task by generalizing from \mathcal{T} .

In particular, meta-learning approaches for few-shot learning usually introduce a generic embedding model f_ϕ parameterized by ϕ as a meta learner to map the task-specific domain into a common feature space. Also, we denote a base

learner as Λ . It is a kind of learning model, which works in the individual task to estimate parameter \mathbf{w}_i for the task \mathcal{T}_i . Specifically, this process can be formulated as:

$$\Lambda(f_\phi(\mathcal{D}_i^{\text{train}})) = \min_{\mathbf{w}_i} \mathbb{E}_{(x_t, y_t) \in \mathcal{D}_i^{\text{train}}} [\mathcal{L}^{\text{base}}(f_\phi(x_t), y_t; \mathbf{w}_i)], \quad (1)$$

where $\mathcal{L}^{\text{base}}$ represents the loss function of the base learner. x_t, y_t denote the training sample and its label in support set $\mathcal{D}_i^{\text{train}}$ of task \mathcal{T}_i , respectively. The squared loss is used as $\mathcal{L}^{\text{base}}$ in this paper.

However, if we ignore the knowledge that can be transferred between tasks and only re-train a predictor for new task with a few examples, the predictor cannot achieve satisfactory performance. For this issue, the meta learner aims to minimize the generalization (or test) error across tasks, and helps the base learner to improve performance. The objective used to calculate the parameter ϕ of the meta learner f_ϕ can be written as:

$$\begin{aligned} & \min_{\phi} \sum_{i=1}^M \mathcal{L}^{\text{meta}}(\mathcal{D}_i^{\text{test}}, \mathbf{w}_i, \phi) \\ & = \min_{\phi} \sum_{i=1}^M \mathcal{L}^{\text{meta}}(\mathcal{D}_i^{\text{test}}; \Lambda(\mathcal{D}_i^{\text{train}}; \phi), \phi), \end{aligned} \quad (2)$$

where $\mathcal{L}^{\text{meta}}$ denotes meta-learning loss function. In this paper, we adopt the cross-entropy loss for $\mathcal{L}^{\text{meta}}$.

In fact, CNN is usually used as the embedding model f_ϕ . In order to optimize the ϕ in Equation (2) by back-propagation and stochastic gradient descent (SGD), Λ is preferred to be a simple and efficient model (*e.g.*, nearest neighbor methods or linear models) (Bertinetto et al. 2019). Once

the embedding model f_ϕ is learned, the base learner for a new task $\mathcal{T}_{\text{new}} = (\mathcal{D}_{\text{new}}^{\text{train}}, \mathcal{D}_{\text{new}}^{\text{test}})$ can be got by Equation (1).

The Proposed MetaSegNet

In this section, we describe each module of the proposed MetaSegNet in detail.

Embedding network. In order to effectively extract the meta-knowledge for K -way, N -shot semantic segmentation, we propose a novel architecture of embedding network as a meta learner in our framework as illustrated in Figure 2.

The proposed embedding network consists of two sub-modules, *i.e.*, the feature extractor and the feature fusion module. The architecture of our embedding network is different from the existing methods which utilize two separate modules to extract features of support and query set, respectively (Shaban et al. 2017; Hu et al. 2019). In our framework, we use one feature extractor to extract the local and global feature simultaneously for support and query set, and the fusion module is used to fuse local and global information for better predicting the label of each pixel.

Besides, most existing methods directly use VGG-16 or ResNet-50 pre-trained on the ImageNet as the backbone. To tackle the distribution divergence as aforementioned, we train MetaSegNet from scratch without introducing any extra information. So, we just use ResNet-9 which has fewer parameters as backbone in order to prevent the possible over-fitting. With the same goal of these previous semantic segmentation methods, our framework is also required to conduct the dense prediction for a full-resolution output. Specifically, to handle the conflicting demands of multi-scale reasoning and full-resolution dense prediction, we first remove the max-pooling layers behind block-3 and block-4, and use the dilated convolution (Yu and Koltun 2015) instead of the traditional convolution in block-4 to aggregate the multi-scale contextual information.

In particular, we argue that the dilated convolution in block-4 is not enough to extract the local information for semantic segmentation. Inspired by (Liu, Rabinovich, and Berg 2015), we add an additional branch to extract the global context to help semantic segmentation. Typically, two max-pooling layers are inserted ahead and behind the block-5 in the global context branch to expand the receptive fields.

In the feature fusion module, we first unpool (replicate) global feature to the same size of local feature map spatially and then concatenate them together. Since the two kinds of features could have different scales and norms, we apply ℓ_2 -norm in each channel for normalization. After this normalization, we reshape the combined feature map to pixel feature map for pixel-wise classification.

Differentiable linear base learner. According to (Franceschi et al. 2018; Balcan, Khodak, and Talwalkar 2019), the choice of base learner is important for optimization based meta-learning methods. Because the remaining embedding model is optimized to generalize across different tasks in Equation (2), the parameter of base learner Λ needs to be calculated in an efficient way. Moreover, when we apply the stochastic gradient descent (SGD) to estimate the parameters ϕ of the embedding model, the gradients of

Model	Params(M)	Without pre-training	End-to-end	Mean IoU(%)
Fine-tuning	14.2	×	✓	28.6
Prototype	14.2	×	✓	35.0
SG-One	19	×	✓	39.4
PANet	14.2	✓	✓	41.3
PLSEG	28.4	×	×	42.6
MetaSegNet (ours)	13.1	✓	✓	43.3

Table 1: The performance of 2-way, 5-shot semantic segmentation on PASCAL-5². The proposed MetaSegNet needs no extra knowledge. Moreover, the parameter of MetaSegNet is least among all models.

the test error $\mathcal{L}^{\text{meta}}(\mathcal{D}^{\text{test}}; \mathbf{w}, \phi)$ with respect to ϕ should be precisely computed.

The no-parametric metric based model avoids the optimizing parameter in base learner. And the linear model is easy to compute the derivative of the linear model parameter \mathbf{w} with respect to ϕ . Their efficacy has been proven in few-shot classification problem (Lee et al. 2019). However, the classical semantic segmentation frameworks (Long, Shelhamer, and Darrell 2015; Chen et al. 2018) use deconvolution layer as segmentation head to predict label for each pixel. It is difficult to calculate the analytic or optimal solution for deconvolution layer, thus if we directly use it as base learner in Equation (2), the whole framework will be hard to train. Therefore, we choose the differentiable linear model (*i.e.*, ridge regression) as base learner in our framework. Hence, the base learner Λ can be written as:

$$\Lambda(\mathbf{X}) = \arg \min_{\mathbf{w}} \|\mathbf{X}\mathbf{w} - \mathbf{y}\|^2 + \lambda \|\mathbf{w}\|^2, \quad (3)$$

where $\mathbf{X} \in \mathbb{R}^{n \times c}$ is the pixel feature matrix obtained from the embedding network. $\mathbf{y} \in \mathbb{R}^n$ is the label of each pixel.

Furthermore, there is a closed form solution for Equation (3). Support for task $\mathcal{T}_i = (\mathcal{D}_i^{\text{train}}, \mathcal{D}_i^{\text{test}})$, we can directly obtain \mathbf{w} as:

$$\mathbf{w} = (\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I})^{-1} \mathbf{X}^\top \mathbf{y}. \quad (4)$$

Due to the property of Equation (3) (*i.e.*, there exists the closed solution), it is very easy to integrate ridge regression into meta-learning optimization framework. We can directly optimize parameter ϕ of the embedding model in Equation (2) by back-propagation algorithm.

In fact, the ridge regression was originally designed for the regression task, we also adjust the prediction of base linear Λ by Equation (5), as in (Bertinetto et al. 2019).

$$\hat{\mathbf{y}} = \alpha \mathbf{X}' \mathbf{w} + \beta, \quad (5)$$

where $\mathbf{X}' \in \mathbb{R}^{n \times c}$ is the feature matrix of the test image. The $\alpha \in \mathbb{R}$ and $\beta \in \mathbb{R}$ denote the scale and bias, respectively.

Meta-learning objective. In the proposed framework, we adopt the cross-entropy loss to evaluate the pixel-wise prediction. And we regard λ in Equation (4) and α, β in Equa-

tion (5) as learnable parameters. Finally, the meta-learning objective can be written as:

$$\mathcal{L}^{\text{meta}}(\mathcal{T}; \mathbf{w}, \phi, \alpha, \beta, \lambda) = -\frac{1}{|\mathcal{T}| \cdot |\mathcal{D}_i^{\text{test}}| |\mathcal{I}|} \sum_{i \in \mathcal{T}} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}_i^{\text{test}}} \sum_{p \in \mathcal{I}} \log(S_{py_p}). \quad (6)$$

\mathcal{I} is the set of pixels in the image. Let s_{pc} be the score of pixel p and class c , which can be got by Equation (5). $S_{pc} = \exp(s_{pc}) / \sum_{k=1}^N \exp(s_{pk})$ represents the softmax probability of class c at pixel p .

Experiments

Dataset, Metric and Implementation

Dataset. We perform the extensive evaluation on two semantic segmentation benchmark datasets, *i.e.*, PASCAL VOC and COCO: (1) PASCAL VOC 2012 (Hariharan et al. 2015) contains 20 different object categories which consists of 10,582 and 1,449 images as the training and validation sets, respectively. (2) COCO 2014 (Lin et al. 2014) dataset is a challenging large-scale dataset which contains 80 different object categories. In COCO, 82,783 and 40,504 images are used for training and validation, respectively.

Evaluation metric. To compare the performance of the different models, the average of mean intersection-over-union (mIoU) for each task is introduced for evaluation.

Implementation details. Our proposed MetaSegNet involves five residual blocks. Each residual block consists of three modules, and each module includes a 3×3 convolution with k filters, a batch normalization and a LeakyReLU (0.1). The size of five residual blocks are set as 64, 128, 256, 512 and 512, respectively. We set block-1 to block-2 as the backbone which is followed by a 2×2 max-pooling.

In the global feature branch, two 2×2 max-pooling layers are inserted before and after block-5, respectively. A global pooling is employed to extract the global feature. In the local feature branch, the dilated convolution is introduced here to extract the multi-scale local features. Specifically, in the 1-way segmentation setting, the dilation for block-4 is set as 2, 4 and 8, respectively. While, in the K -way setting, the dilated convolution is used in block-3 with dilation as 1, 2 and 4, and the dilation for block-4 is set as 8, 16 and 32.

For optimization, we use Adam (Kingma and Ba 2014) with learning rate as 0.001. For PASCAL, our model is meta-trained for 40 epochs, and each epoch consists of 1000 episodes. For COCO, our model is meta-trained for 80 epoches, and each epoch consists of 500 episodes. During meta-training, we adopt horizontal flip, and randomly rotate the image with 0, 90, 180 or 270 degree for data augmentation. We implement our method by PyTorch on two NVIDIA 2080 Ti GPUs with 12 GB memory.

Results on PASCAL (K -way, N -shot)

Setting. PASCAL-5² (Shaban et al. 2017) is a very popular few-shot segmentation benchmark, thus we also report our results on this dataset. In fact, PASCAL-5² comes from PASCAL, the 20 object categories in PASCAL are divided

Model	Shot=1	Shot=5	Shot=10
Prototype	11.0	12.8	14.3
Fine-tuning	29.3	28.2	29.1
MetaSegConv	26.9	27.3	26.8
MetaSegNet (ours)	33.2	37.9	38.3

Table 2: The performance of 2-way semantic segmentation on COCO. All the models are not pre-trained.

into 4 splits with three splits for training and the rest one for testing. Typically, because of the unbalance of the multi-object image in each split, we choose PASCAL-5² as the test split, and evaluate our method in 2-way semantic segmentation setting, which are the same with (Dong and Xing 2018). The images that only contain one object in training splits are used to complement the meta-training set. The images in PASCAL-5² which contain *person* and another held-out class are sampled as support and query set for new tasks. In fact, there are four situations of the new tasks, *i.e.*, {(table, person), (dog, person), (horse, person), (motobike, person)}. We randomly sample 100 tasks for each situation, and calculate the average performance of the four situations. Query set in each new task has two images.

To evaluate our MetaSegNet in 2-way semantic segmentation task on PASCAL-5², we compare it with two baseline models (*i.e.*, baseline fine-tuning and baseline prototype model (Dong and Xing 2018)), one 1-way segmentation model SG-One (running 2 times) (Zhang et al. 2018) and two state-of-the-art K -way, N -shot semantic segmentation models (*i.e.*, PLSEG (Dong and Xing 2018) and PANet (Wang et al. 2019)). The fine-tuning model is a standard FCN that employs a VGG-16 pretrained on the ImageNet as the backbone. The images from training classes are first used to train this backbone, then the support set in new task is used to fine-tune this model.

Results. We report the results in Table 1. In the 2-way, 5-shot setting, our framework achieves the state-of-the-art performance in spite of simple embedding model and no extra information. Although the information from ImageNet provides extra knowledge for baselines and SG-One model, the multi-object metric or matching problem hinders these approaches from reaching satisfied results in K -way setting. Differently, the proposed MetaSegNet can handle multi-object matching problem from classification perspective, and achieve good performance in this setting. Furthermore, compared with the other state-of-the-art K -way semantic segmentation models (*i.e.*, PLSEG and PANet), MetaSegNet can learn more useful cross-task meta knowledge, which can effectively guide the semantic segmentation in query set.

Results on COCO (K -way, N -shot)

Setting. Similar to (Zhang et al. 2019), we choose a subset of the original COCO dataset to evaluate our model and baselines. We choose 18 categories as training classes, and 7 categories as new classes. The training class IDs are {1, 2, 4, 5, 7, 8, 9, 19, 20, 21, 23, 24, 25, 62, 63, 64, 65, 70}, and the testing class IDs are {3, 6, 16, 17, 18, 22, 67}. Because

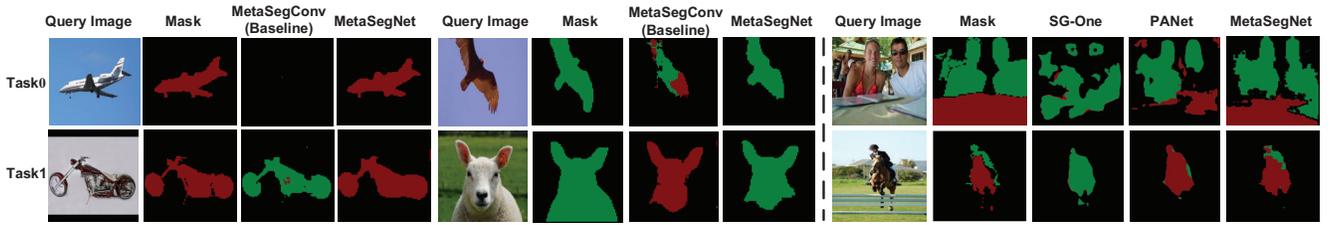


Figure 3: Visual examples for baselines and MetaSegNet. The left part is the results of two tasks in 2-way, 5-shot single object semantic segmentation on PASCAL-5ⁱ_{val}. The right part is the results of two tasks in 2-way, 5-shot multi-object semantic segmentation on PASCAL-5². Red and green regions represent the labels of different semantic objects.

Model	Params (M)	Extra images	Mean IoU(%)
co-FCN (param.)	34.2	1.2×10^6	58.3
OSLSM	276.7	1.2×10^6	61.5
MetaSegNet (ours)	13.1	0	53.4
MetaSegNet (ours)	13.1	6×10^4	59.5

Table 3: The results of 1-way, 5-shot segmentation on PASCAL-5ⁱ. All results are computed by taking the average of 4 splits on PASCAL-5ⁱ. Extra images represent the number of images for per-training the model. Although, MetaSegNet uses 1/20 extra images and 1/21 model size compared with OSLSM, the Mean IoU of MetaSegNet just decreases 2% (61.5 vs. 59.5) accuracy.

of the limitation of images that contain multiple classes, we evaluate MetaSegNet on 2-way with 1-shot, 5-shot and 10-shot few-shot semantic segmentation tasks. Considering that the images in the training classes are more abundant than that in PASCAL, we directly sample images from training classes which contain 2 classes to complement the meta-training set. In addition, these images are cropped according to the bounding box of semantic objects to reduce the size of background. There are 1500 tasks in total sampled from new classes as meta-testing set, where the query set in each task includes 2 images.

According to our best knowledge, our work is the first attempt to deal with K -way, N -shot semantic segmentation on COCO 2014. For comparison, we design three baseline models. As the same as the previous experiment, the fine-tuning model is also a variant of traditional FCN. The embedding module in MetaSegNet without global context branch and 1×1 conv layer are used as backbone and segmentation head for fine-tuning model, respectively. The images from training classes are used to train the model, and the support set for new task is used to fine-tune this model. The MetaSegNet utilizes the same model with fine-tuning model, however, MetaSegConv is trained by the optimization based meta-learning framework. The 1×1 conv layer is regarded as a base learner that is optimized by one step by Adam with learning rate as 0.001. As for prototype baseline, the embedding module is the same as MetaSegNet. However, the ridge regression is replaced by Euclidean distance, and the mean of pixel features is as the prototype for one

semantic category as in (Snell, Swersky, and Zemel 2017).

Results. The results are reported in Table 2. Although the COCO 2014 is more difficult than PASCAL as we know, our MetaSegNet still outperforms the baselines with a large margin in all cases. Also, we could notice that the performance of MetaSegNet increases with the increasing number of shot, while all the other baselines have no obvious variance.

Extension to 1-way, N -shot

To fully investigate the performance of our MetaSegNet, we degenerate our proposed method from the K -way, N -shot task to the traditional 1-way, N -shot task and report the performance on PASCAL-5ⁱ.

Setting. Since most of conventional methods are specifically designed to solve this 1-way, N -shot task, for a fair comparison, we directly follow the same experiment settings as (Shaban et al. 2017). In particular, one class is sampled from training classes as the foreground, and the other categories are regarded as the background. Thus, the mask in this case is a binary map. The four splits on PASCAL-5ⁱ are all used to evaluate MetaSegNet. In each split, 1000 tasks are sampled from new classes as the meta-testing set, and there is one image in query set for each new task.

Two state-of-the-art 1-way segmentation models, *i.e.*, co-FCN (Rakelly et al. 2018) and OSLSM (Shaban et al. 2017), are introduced for comparison. Since there is no multi-object metric case in this setting, the superiority of MetaSegNet is not so obvious when compared with these models. We also use extra images to pre-train our MetaSegNet. As using the entire ImageNet for pre-training is too heavy and time-consuming, we only pre-train our MetaSegNet on the mini-ImageNet (60,000 images) (Vinyals et al. 2016). Besides, we also evaluate MetaSegNet without pre-training.

Results. We report the results in Table 3. It is very convenient for metric-based methods to measure the similarity between foregrounds when each foreground only contains one object. Therefore, the metric-based approaches perform well in this setting. Although MetaSegNet is designed for K -way, N -shot task by using simple model and much less extra knowledge, it can still achieve a competitive result when compared with other 1-way methods in the 1-way setting.

Ablation Study

Considering that there are no validation categories on PASCAL-5ⁱ, we use a new few-shot semantic segmenta-

Model	Episodic Training	Linear model	GC -branch	Mean IoU(%)
Fine-tuning				27.28
MetaSegConv	✓			35.00
MetaSegNet-NG	✓	✓		37.75
MetaSegNet (ours)	✓	✓	✓	38.92

Model	Episodic Training	Linear model	GC -branch	Mean IoU(%)
Fine-tuning				31.95
MetaSegConv	✓			36.04
MetaSegNet-NG	✓	✓		48.34
MetaSegNet (ours)	✓	✓	✓	49.44

Table 4: Ablation study on PASCAL-5ⁱ_{val} in 2-way setting. The top table is the results of 1-shot, and the bottom table is the results of 5-shot. GC-branch represents the global context branch in embedding module.

tion dataset, *i.e.*, PASCAL-5ⁱ_{val}, in ablation study. Similar to PASCAL-5ⁱ, the PASCAL VOC 2012 is also divided into 4 splits to evaluate our method. However, in each split, three classes are chosen as validation set (*i.e.*, 6-8 in split 0; 11-13 in split 1; 16-18 in split 2 and 1-3 in split 3), and the other 12 classes are used as training classes. All images in new categories are sampled to make up support set and query set for new tasks in meta-testing set, just like miniImageNet (Vinyals et al. 2016) utilized for the few-shot classification task. Because the images containing multiple objects are limited on PASCAL-5ⁱ_{val}, all the methods are evaluated in 2-way single object semantic segmentation. Each image in query set contains one semantic object which belongs to the K categories. We randomly sample 500 tasks as meta-testing set. There are four images in query set for new tasks (*i.e.*, two images for each class).

Global context branch. To verify efficacy of the global context branch, we evaluate the performance of MetaSegNet-NG which uses the embedding module without the global context branch. The dilation of block-4 is set as 2, 4, 8, which is the same with MetaSegNet. The experimental results are shown in Table 4. As seen, the global context branch indeed benefits the few-shot semantic segmentation.

Linear classification model. The MetaSegConv is used to evaluate the impact of the ridge regression, which employs the same embedding module with MetaSegNet-NG. Besides, MetaSegConv leverages 1×1 conv layer as base learner, unlike the ridge regression is utilized in MetaSegNet-NG. According to Table 4, the mIoU of MetaSegNet-NG outperforms MetaSegConv by 2.75% (37.75 vs. 35.00) and 12.3% (48.34 vs. 36.04) in 1-shot and 5-shot, respectively. The main reason is that the conv layer is too complex as base learner and it is not conducive to optimization whereas the ridge regression can avoid this issue.

Also, Table 4 shows the performance of fine-tuning in K -way, N -shot semantic segmentation to further confirm the power of meta-learning in few-shot semantic segmentation.

Efficacy of shot in support set. In this part, we research the impact of shot in support shot during meta-testing. As

seen in Figure 4, the performance of our method is in positive correlation with the shot in support set. However, the change of the baselines is not very obvious. This is attributable to our framework can learn how to efficiently use the information from support set by meta-learning framework. Just like human, we cannot learn well in extremely short of information (*e.g.*, 1-shot) in a new task, but our skill to solve a new task is in positive correlation with attempts.

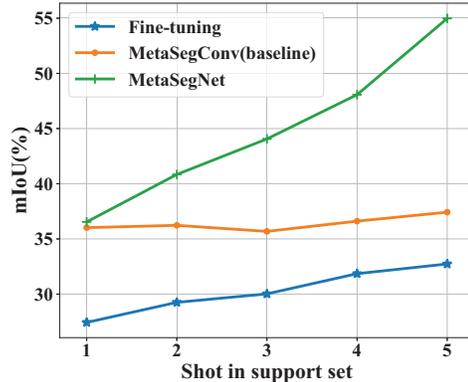


Figure 4: The results of different shots in support set for 2-way semantic segmentation on PASCAL-5¹_{val}.

Conclusion

In this paper, we aim to solve the K -way, N -shot few-shot semantic segmentation task, which is more difficult than the conventional 1-way setting. Different from the existing works, we formulate this task as a pixel classification problem instead of measuring the similarity between images. Under this formulation, a novel few-shot semantic segmentation framework based on optimization meta-learning is proposed. Thanks to the proposed embedding module and linear model, our framework can achieve impressive performance. Furthermore, our method provides a new alternative to deal with K -way, N -shot few-shot segmentation, which is more meaningful than the conventional setting.

Acknowledge

The work of Yang Gao is supported by Science and Technology Innovation 2030-“New Generation Artificial Intelligence” Major Project (No. 2018AAA0100905), NSFC (No. 61432008) and the Collaborative Innovation Center of Novel Software Technology and Industrialization. The work of Yinghuan Shi is supported by the Fundamental Research Funds for the Central Universities (No. 020214380056) and NSFC (No. 61673203).

References

Balcan, M.; Khodak, M.; and Talwalkar, A. 2019. Provable guarantees for gradient-based meta-learning. In *ICML*.

- Bertinetto, L.; Henriques, J. F.; Torr, P. H. S.; and Vedaldi, A. 2019. Meta-learning with differentiable closed-form solvers. In *ICLR*.
- Chen, L.; Papandreou, G.; Kokkinos, I.; Murphy, K.; and Yuille, A. L. 2018. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *TPAMI*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.; Li, K.; and Li, F. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*.
- Dong, N., and Xing, E. P. 2018. Few-shot semantic segmentation with prototype learning. In *BMVC*.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*.
- Franceschi, L.; Frasconi, P.; Salzo, S.; Grazzi, R.; and Pontil, M. 2018. Bilevel programming for hyperparameter optimization and meta-learning. In *ICML*.
- Girshick, R. B. 2015. Fast R-CNN. In *ICCV*.
- Hariharan, B.; Arbeláez, P. A.; Girshick, R. B.; and Malik, J. 2015. Hypercolumns for object segmentation and fine-grained localization. In *CVPR*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural Computation*.
- Hu, T.; Yang, P.; Zhang, C.; Yu, G.; Mu, Y.; and Snoek, C. G. M. 2019. Attention-based multi-context guiding for few-shot semantic segmentation. In *AAAI*.
- Kingma, D. P., and Ba, J. 2014. Adam: A Method for Stochastic Optimization. arXiv.
- Lee, K.; Maji, S.; Ravichandran, A.; and Soatto, S. 2019. Meta-Learning with Differentiable Convex Optimization. arXiv.
- Li, W.; Xu, J.; Huo, J.; Wang, L.; Gao, Y.; and Luo, J. 2019. Distribution consistency based covariance metric networks for few-shot learning. In *AAAI*.
- Lin, T.; Maire, M.; Belongie, S. J.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft COCO: common objects in context. In *ECCV*.
- Liu, W.; Rabinovich, A.; and Berg, A. C. 2015. ParseNet: Looking Wider to See Better. arXiv.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *CVPR*.
- Rakelly, K.; Shelhamer, E.; Darrell, T.; Efros, A. A.; and Levine, S. 2018. Conditional networks for few-shot semantic segmentation. In *ICLR, Workshop*.
- Ravi, S., and Larochelle, H. 2017. Optimization as a model for few-shot learning. In *ICLR*.
- Shaban, A.; Bansal, S.; Liu, Z.; Essa, I.; and Boots, B. 2017. One-shot learning for semantic segmentation. In *BMVC*.
- Snell, J.; Swersky, K.; and Zemel, R. S. 2017. Prototypical networks for few-shot learning. In *NIPS*.
- Vilalta, R., and Drissi, Y. 2002. A perspective view and survey of meta-learning. *AI*.
- Vinyals, O.; Blundell, C.; Lillicrap, T.; Kavukcuoglu, K.; and Wierstra, D. 2016. Matching networks for one shot learning. In *NIPS*.
- Wang, K.; Liew, J. H.; Zou, Y.; Zhou, D.; and Feng, J. 2019. Panet: Few-shot image semantic segmentation with prototype alignment. In *ICCV*.
- Yu, F., and Koltun, V. 2015. Multi-Scale Context Aggregation by Dilated Convolutions. arXiv.
- Zhang, X.; Wei, Y.; Yang, Y.; and Huang, T. 2018. Sg-one: Similarity guidance network for one-shot semantic segmentation. arXiv.
- Zhang, C.; Lin, G.; Liu, F.; Yao, R.; and Shen, C. 2019. CANet: Class-Agnostic Segmentation Networks with Iterative Refinement and Attentive Few-Shot Learning. arXiv.