

Adaptive Prototype Learning and Allocation for Few-Shot Segmentation

Gen Li^{1,3}, Varun Jampani², Laura Sevilla-Lara¹, Deqing Sun², Jonghyun Kim³, Joongkyu Kim^{3*}

¹University of Edinburgh ²Google Research ³Sungkyunkwan University

Abstract

Prototype learning is extensively used for few-shot segmentation. Typically, a single prototype is obtained from the support feature by averaging the global object information. However, using one prototype to represent all the information may lead to ambiguities. In this paper, we propose two novel modules, named superpixel-guided clustering (SGC) and guided prototype allocation (GPA), for multiple prototype extraction and allocation. Specifically, SGC is a parameter-free and training-free approach, which extracts more representative prototypes by aggregating similar feature vectors, while GPA is able to select matched prototypes to provide more accurate guidance. By integrating the SGC and GPA together, we propose the Adaptive Superpixel-guided Network (ASGNet), which is a lightweight model and adapts to object scale and shape variation. In addition, our network can easily generalize to k -shot segmentation with substantial improvement and no additional computational cost. In particular, our evaluations on COCO demonstrate that ASGNet surpasses the state-of-the-art method by 5% in 5-shot segmentation.¹

1. Introduction

Humans have a remarkable ability to learn how to recognize novel objects after seeing only a handful of exemplars. On the other hand, deep learning based computer vision systems have made tremendous progress, but have largely depended on large-scale training sets. Also, deep networks mostly work with predefined classes and are incapable of generalizing to new ones. The field of few-shot learning studies the development of such learning ability in artificial learning systems, where only a few examples of the new category are available.

In this work, we tackle the few-shot segmentation problem, where the target is learning to segment objects in a given query image while only a few support images with ground-truth segmentation masks are available. This is a

*Corresponding author

¹Code is available at <https://git.io/ASGNet>.

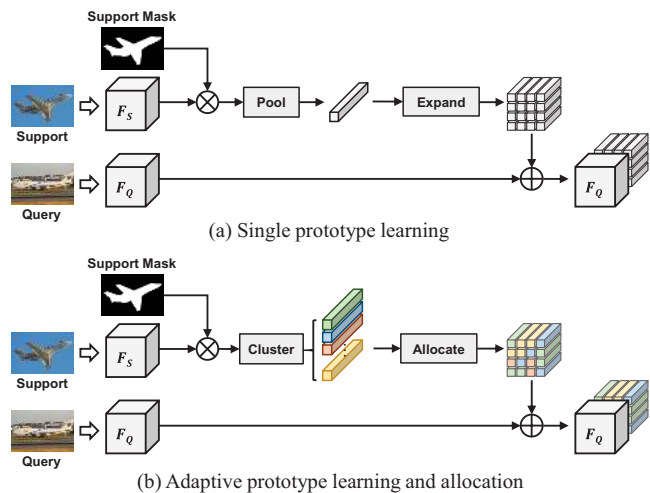


Figure 1. Comparison between (a) single prototype learning and (b) proposed adaptive prototype learning and allocation. We utilize superpixel-guided clustering to generate multiple prototypes and then allocate them pixel-wise to query feature.

challenging problem as the test data are novel categories which do not exist in the training set, and there are usually large variations in appearance and shape between the support and query images.

Current few-shot segmentation networks usually extract features from both query and support images, and then propose different approaches for feature matching and object mask transfer from support to query image. This feature matching and mask transfer are usually performed in one of two ways: prototypical feature learning or affinity learning. Prototypical learning techniques condense the masked object features in a support image into a single or few prototypical feature vectors. Then these techniques find the pixel locations of similar features in the query image to segment the desired object. A key advantage of prototype learning is that the prototypical features are more robust to noise than pixel features. However, prototypical features inevitably drop spatial information, which is important when there is a large variation in the object appearance between the support and query images. In addition, most prototypical learning

networks [40, 42, 32, 29, 4] merely generate a single prototype by masked average pooling as shown in Figure 1(a), thus losing information as well as discriminability.

Affinity learning techniques [39, 36, 31] on the other hand, directly try to match object pixels in a support image to query image pixels thereby transferring the object mask. These techniques predict cross-image pixel affinities (also called connection strengths) using learned features, which perform feature matching while preserving spatial information better than prototypical learning approaches. However, affinity learning techniques are prone to over-fitting on training data as they try to solve an under-constrained pixel-matching problem with dense affinity matrices.

In this work, we propose a novel prototypical learning technique that addresses some of the main shortcomings of existing ones. In particular, we want to adaptively change the number of prototypes and their spatial extent based on the image content, making the prototypes content-adaptive and spatially-aware. This adaptive, multi-prototype strategy is important to deal with large variations in object scales and shapes across different images. Intuitively, when an object occupies a large portion of the image, it carries more information and thus requires more prototypes to represent all the necessary information. On the contrary, if the object is fairly small and the proportion of the background is large, then a single or few prototypes are sufficient. In addition, we want the support region (spatial extent) for each of the prototypes to be adaptive to object information that is present in the support image. Concretely, we aim to divide the support feature into several representative areas according to the feature similarity. We also want to adaptively choose more important prototypes while finding similar features in a query image. As different object parts are visible in different image regions and in different query images, we want to dynamically allocate different prototypes across query image for feature matching. For example, some parts of the object can be occluded in a query image and we want to dynamically choose the prototypes that are corresponding to the visible parts in the query image.

We achieve this adaptive, multi-prototype learning and allocation with our Adaptive Superpixel-guided Network (ASGNet) that leverages superpixels for adapting both the number and support regions of the prototypes. The schematic illustration is presented in Figure 1(b). In particular, we propose two modules, Superpixel-guided Clustering (SGC) and Guided Prototype Allocation (GPA), which form the core of ASGNet. The SGC module does fast feature-based superpixel extraction on the support image and the resulting superpixel centroids are considered as prototypical features. Since superpixel shapes and numbers are adaptive to the image content, the resulting prototypes also become adaptive. The GPA module uses an attention-like mechanism to allocate most relevant support prototype features to

each pixel in a query image. In summary, the SGC module provides adaptive prototype learning both in terms of the number of prototypes and their spatial extents, and the GPA module provides adaptive allocation of the learned prototypes when processing query features. These two modules make ASGNet highly flexible and adaptive to varying object shapes and sizes, allowing it to generalize better to unseen object categories. We make the following contributions:

- We propose the Adaptive Superpixel-guided Network (ASGNet), a flexible prototypical learning approach for few-shot segmentation that is adaptive to different object scales, shapes and occlusions.
- We introduce two novel modules, namely Superpixel-guided Clustering (SGC) and Guided prototype allocation (GPA), for adaptive prototype extraction and allocation respectively. They can serve as effective plug-and-play components on feature matching.
- ASGNet achieves top-performing results with fewer parameters and less computation. Specifically, the proposed method obtains mIoUs of 64.36%/42.48% in the 5-shot setting on Pascal-5ⁱ/COCO-20ⁱ, exceeding the state-of-the-art by 2.40%/5.08%.

2. Related Work

Semantic Segmentation. Most existing semantic segmentation methods are based on fully convolutional networks (FCNs) [19], which replace fully connected layers with fully convolutional ones for pixel-level prediction. Recent breakthroughs in semantic segmentation have mainly come from multi-scale feature aggregation [43, 2, 35, 10] or attention mechanisms [7, 38, 15, 28, 3, 44, 41]. These methods often use dilated convolution kernels [37] and set up an encoder-decoder structure to obtain a large receptive field while maintaining the feature resolution. Although these methods achieve tremendous success, they need long training time and a large amount of pixel-level labeled ground truth to fully supervise the network. Also, in the inference stage, trained models cannot recognize new classes that do not exist in the training set.

Few-shot Learning. Few-shot learning focuses on the generalization ability of models, so that they can learn to predict new classes given only a few annotated examples. Existing methods mainly concentrate on metric-learning [30, 27, 25] and meta-learning [33, 6, 22]. The core idea of metric learning is distance measurement, and it is generally formulated as an optimization of distance/similarity between images or regions. In meta-learning approaches, the main idea is to define specific optimization or loss functions to achieve fast learning capability. Among these methods, the concept of prototypical networks has been extensively adopted

in few-shot segmentation, which largely reduces computational budget while maintaining high performance. Most methods focused on image classification, while recently few-shot segmentation has received growing attention.

Few-shot Segmentation. Few-shot segmentation is an extension of few-shot classification, and it tackles a more challenging task of predicting a label at each pixel instead of predicting a single label for the entire image. This research problem was introduced by Shaban *et al.* [23], who proposed a classical two-branch network. Later on, PL [4] introduced the idea of using prototypes. In that work, each prediction is generated by measuring the similarity between prototypes and pixels in query image. SG-One [42] proposed masked average pooling to obtain the object-related prototype, which has been widely used in subsequent works. PANet [32] introduced a novel prototype alignment regularization to fully exploit the support knowledge, making the final prediction with only the measurement of cosine distance. Based on masked average pooling, CANet [40] expanded the prototype to the same size of the query feature and concatenated them together. This work also used an iterative optimization module to refine the segmentation result. PGNet, BriNet and DAN [39, 36, 31] introduced dense pixel-to-pixel connections between support and query features to maintain spatial information. More recently, PMMs [34] leveraged the expectation-maximization (EM) algorithm to generate multiple prototypes. However, all prototypes have the same relevance in this model, which can potentially be sensitive to poorly matched prototypes. Instead, in our work we utilize the similarity between each prototype and the query feature to select the most relevant prototype at each pixel location.

Superpixel Segmentation. A superpixel is defined as a set of pixels with similar characteristics (color, texture, category). Superpixels have been effective in many computer vision tasks, and have been used as the basic units for the image segmentation task. Superpixels carry more information than pixels, and can provide more compact and convenient image representations for the downstream vision tasks. For more details about traditional and existing methods on superpixels, please refer to [1, 26].

Our research is inspired by the maskSLIC [12] and superpixel sampling network (SSN) [13]. MaskSLIC adapts SLIC [1] to a defined region of interest (RoI), and the main contribution is in the placement of seed points within the RoI. SSN [13] proposed the first end-to-end trainable superpixel algorithm by making the SLIC algorithm differentiable. Inspired by the insights of these two techniques, we propose the masked superpixel clustering in feature space, which can gather similar features together and generate superpixel centroids as prototypes. Instead of representing the information of the entire object, superpixel centroids stand for the parts of the object with similar characteristics.

3. Problem Definition

The key difference between few-shot segmentation and a general semantic segmentation is that the categories in training and testing set do not intersect. It means that, in the inference stage, the testing set has classes totally unseen in the training. Specifically, given a training set $S_{train} = \{(I_{S/Q}, M_{S/Q})\}$ and testing set $S_{test} = \{(I_{S/Q}, M_{S/Q})\}$, the categories of the two sets do not intersect ($S_{train} \cap S_{test} = \emptyset$). Here $I \in \mathbb{R}^{H \times W \times 3}$ indicates the RGB image and $M \in \mathbb{R}^{H \times W}$ denotes the segmentation mask. Subscripts S and Q represent support and query, respectively. Following the first one-shot segmentation work [23], we align training and testing with the episodic paradigm [30]. In each episode, the input to the model is composed by a query image I_Q and K samples $(I_S^i, M_S^i), i \in \{1, \dots, K\}$ from the support set. All support and query images have the same class c . We estimate the query mask \tilde{M}_Q to approximate the ground truth mask M_Q .

4. Proposed Method

In this section, we first introduce the proposed two modules for prototype generation and matching, which are superpixel-guided clustering (SGC) and guided prototype allocation (GPA). Then, we discuss the adaptive ability of these two modules. After that, we introduce the overall network architecture, named Adaptive Superpixel-guided Network (ASGNet), which integrates the SGC and GPA modules in one model. The overall structure is shown in Figure 2. Finally, we elaborate the k -shot setting in ASGNet.

4.1. Superpixel-guided Clustering

The core idea of SGC is inspired by the superpixel sampling network (SSN) [13] and MaskSLIC [12]. SSN was the first end-to-end trainable deep network for superpixel segmentation. The key contribution in SSN is converting the nearest-neighbor operation in SLIC [1] into a differentiable one. The traditional SLIC superpixel algorithm uses k-means clustering iteratively with two steps: *Pixel-superpixel association* and *Superpixel centroid update*. Based on the color similarity and proximity, pixels are assigned to different superpixel centroids. Specifically, the input image $I \in \mathbb{R}^{n \times 5}$ is usually in a five-dimensional space (labxy) with n pixels, where lab represents the pixel vector in CIELAB color space and xy indicates the pixel location. After iterative clustering, the algorithm outputs the association map where each pixel n is assigned to one of the m superpixels.

This simple method inspires us with an insightful idea, which is to aggregate the feature map into multiple superpixel centroids in a clustering way, and here superpixel centroids can serve as prototypes. Therefore, instead of computing the superpixel centroids in image space, we estimate

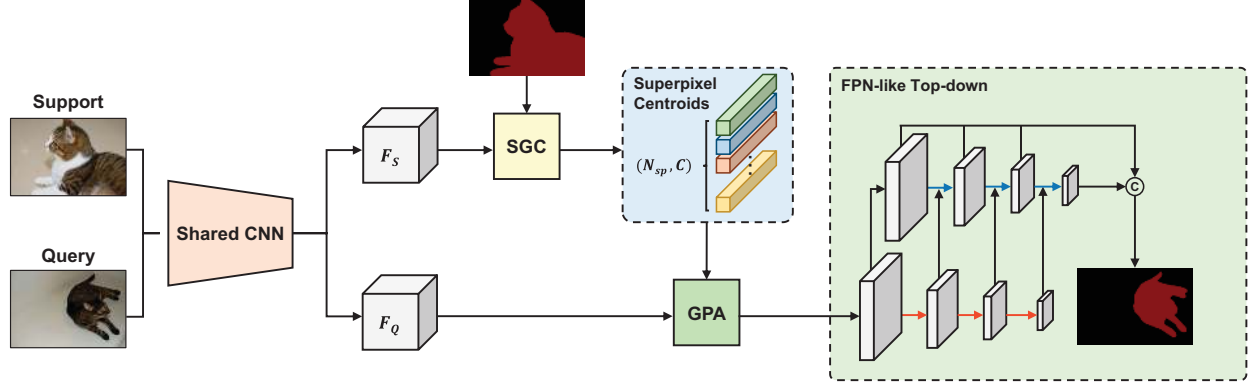


Figure 2. Overall architecture of the proposed Adaptive Superpixel-guided Network.

Algorithm 1 Superpixel-guided Clustering (SGC)

Input: Support feature F_s , support mask M_s , and initial superpixel seeds S^0

concatenate the absolute coordinates to F_s

extract masked features F'_s via support mask M_s

for each iteration t **do**

 Compute association between each pixel p and superpixel i , $Q_{pi}^t = e^{-\|F'_p - S_i^{t-1}\|^2}$

 Update superpixel centroids,
 $S_i^t = \frac{1}{Z_i^t} \sum_{p=1}^{N_m} Q_{pi}^t F'_p$; $Z_i^t = \sum_p Q_{pi}^t$

end for

remove coordinates information

return final superpixel centroids $S \triangleright (N_{sp}, C)$

them in feature space by clustering similar feature vectors. The whole SGC process is delineated in Algorithm 1.

Given support feature $F_s \in \mathbb{R}^{c \times h \times w}$, support mask $M_s \in \mathbb{R}^{h \times w}$ and initial superpixel seeds $S^0 \in \mathbb{R}^{c \times N_{sp}}$ (N_{sp} is the number of superpixels), we aim to obtain the final superpixel centroids, which act as multiple compact prototypes. First, we concatenate the coordinates of each pixel with the support feature map to introduce positional information. Then, we define the distance function D following SLIC [1], which consists of feature and spatial distance:

$$D = \sqrt{(d_f)^2 + (d_s/r)^2}, \quad (1)$$

where d_f , d_s are the Euclidean distance for features and coordinate values, and r is a weighting factor. We filter out the background information with the support mask and only keep the masked features, compressing the feature map from $F_s \in \mathbb{R}^{c \times h \times w}$ to $F'_s \in \mathbb{R}^{c \times N_m}$, where N_m is the number of pixels inside the support mask.

Then we compute superpixel-based prototypes in an iterative fashion. For each iteration t , we first compute the association map Q^t between each pixel p and all superpix-




Figure 3. Illustration of proposed Superpixel-guided Clustering.

els according to the distance function D :

$$Q_{pi}^t = e^{-D(F'_p, S_i^{t-1})} = e^{-\|F'_p - S_i^{t-1}\|^2}. \quad (2)$$

Then, new superpixel centroids are updated as the weighted sum of masked features:

$$S_i^t = \frac{1}{Z_i^t} \sum_{p=1}^{N_m} Q_{pi}^t F'_p, \quad (3)$$

where $Z_i^t = \sum_p Q_{pi}^t$ is a constant for normalization. The above process is visualized in Figure 3.

Here, we elaborate the selection of initial seeds. Generally, in superpixel algorithm, a $H \times W$ image is evenly partitioned into regular grid cells of size $h \times w$, and each grid cell is considered as an initial seed (i.e., superpixel). However, this initialization is not suitable for our purposes where a foreground mask is given for the support image and we only need to initialize the seeds inside this foreground region. To uniformly initialize seeds in the masked region, we refer to MaskSLIC [12] for iteratively placing each initial seed, and the pipeline is depicted in Figure 5. This seed initialization results in faster convergence of superpixel-guided clustering with only a few iterations.

4.2. Guided Prototype Allocation

After extracting prototypes, previous methods mostly follow the design of CANet [40], expanding a single prototype to the same size of the query feature and concatenating them together. However, this operation results in equal

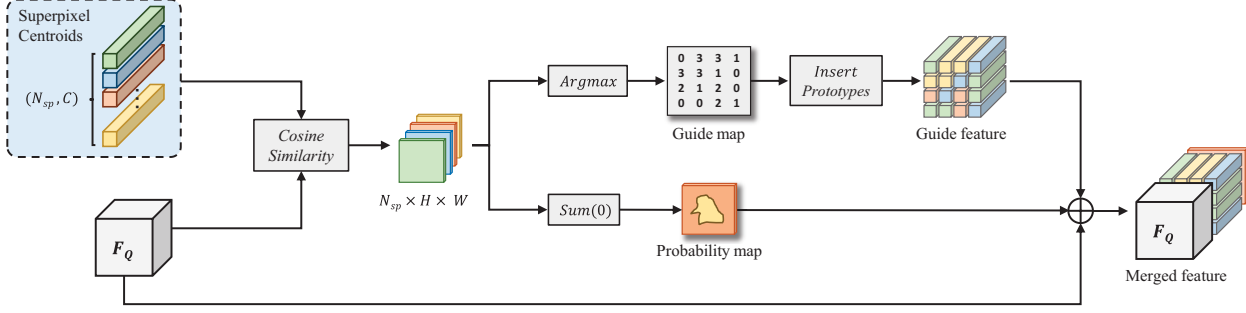


Figure 4. Illustration of proposed Guided Prototype Allocation.




Figure 5. The pipeline of placing initial seeds. (1) Step 1: We conduct Distance Transform and choose the maximum point. (2) Step 2: We set that point to False and then repeat Step 1. After that, we iterate Step 2 until we get coordinates of all initial seeds.

guidance to every location in the query feature. To make prototype matching more adaptive to the query image content, we propose the Guided Prototype Allocation (GPA), illustrated in Figure 4. We first compute the cosine distance to measure the similarity between each prototype and each query feature element:

$$C_i^{x,y} = \frac{S_i \cdot F_q^{x,y}}{\|S_i\| \cdot \|F_q^{x,y}\|} \quad i \in \{1, 2, \dots, N_{sp}\}, \quad (4)$$

where $S_i \in \mathbb{R}^{c \times 1}$ is the i_{th} superpixel centroid (prototype), and $F_q^{x,y} \in \mathbb{R}^{c \times 1}$ is the feature vector at location (x, y) of the query feature. We use this similarity information as input to a two-branch structure. The first branch computes which prototype is the most similar at each pixel location as follows:

$$G^{x,y} = \arg \max_{i \in \{0, \dots, N_{sp}\}} C_i^{x,y}, \quad (5)$$

where the argmax operator is adopted to obtain $G^{x,y}$, which is a single index value representing a particular prototype. Putting all index values together, we get a guide map $G \in \mathbb{R}^{h \times w}$. Then, by placing corresponding prototype in each position of the guide map, we obtain the guide feature $F_G \in \mathbb{R}^{c \times h \times w}$ to achieve pixel-wise guidance. While in the other branch, all the similarity information C is added up across all the superpixels to get the probability map P .

Finally, we concatenate the probability map and the guide feature with the original query feature F_Q to provide the guiding information, and thus obtain the refined query feature F'_Q :

$$F'_Q = f(F_Q \oplus F_G \oplus P), \quad (6)$$

where \oplus indicates the concatenation operation along channel dimension, and $f(\cdot)$ is a 1×1 convolution.

4.3. Adaptability

As mentioned before, we argue that one of the key attributes in the proposed network is its adaptive ability for few-shot semantic segmentation. In Figure 6, we provide some examples illustrating the adaptive ability of SGC and GPA. In SGC, to make it adaptive to object scale, we define a criterion to regulate the number of superpixel centroids as

$$N_{sp} = \min\left(\left\lfloor \frac{N_m}{S_{sp}} \right\rfloor, N_{max}\right), \quad (7)$$

where N_m is the number of pixels in the support mask; S_{sp} is the average area assigned to each initial superpixel seed, and we set to 100 empirically. When the foreground is fairly small, $N_{sp} = 0$ or 1, this method degrades to a general masked average pooling as shown in Figure 6(a). In addition, to reduce the computational burden, we set a hyper-parameter N_{max} to constrain the maximum number of prototypes. In GPA, we can observe its adaptability to object shape. In other words, it is resilient to occlusion. When severe occlusion exists in query image, e.g., Figure 6(b), GPA can choose the best matched prototype for each query feature location.

4.4. Adaptive Superpixel-guided Network

Based on the above SGC and GPA modules, we propose the Adaptive Superpixel-guided Network (ASGNet) for few-shot semantic segmentation as illustrated in Figure 2. First, the support and query images are fed into a shared CNN (pretrained on ImageNet [14]) to extract features. Then, by passing the support features through SGC with support mask, we obtain the superpixel centroids, which are considered as prototypes. After that, for more accurate pixel-wise guidance, we adopt GPA module to match the prototypes with the query feature. Finally, we use the feature enrichment module [29] and set up a FPN-like [16] top-down structure to introduce multi-scale information. As demonstrated in [29], transferring features from




Figure 6. Illustration of the model adaptability in prototype learning and allocation. SGC and GPA indicate the proposed superpixel-guided clustering and guided prototype allocation respectively, and MAP denotes masked average pooling proposed in [42].

fine to coarse promotes the feature interaction, so we follow their design for fast multi-scale aggregation. Finally, all different scales are concatenated, and each scale yields a segmentation result for computing the loss.

4.5. K-shot Setting

In previous work, the k -shot setting is usually tackled via feature averaging or attention-based fusion. However, it turns out that the improvements from such strategies is minor, while requiring heavy computation. In contrast, based on the proposed SGC and GPA, our ASGNet can easily adopt an efficient k -shot strategy without collapsing support features. Specifically, in each support image and mask pair, we implement SGC to obtain the superpixel centroids. By collecting all superpixel centroids together, we get the overall superpixel centroids S from k shots:

$$S = (S^1, S^2, \dots, S^k), \quad (8)$$

where the number of superpixel centroids is $N_{sp} = \sum_{i=1}^k N_{sp}^i$. By doing so, the GPA can receive a larger range of selections from multiple shots, and thus yield more accurate guidance to segment the object in the query image.

5. Experiments

5.1. Datasets and Evaluation Metric

We choose Pascal-5ⁱ and COCO-20ⁱ, two widely used datasets for few-shot semantic segmentation, to analyze our model performance. Pascal-5ⁱ [23] includes images from the PASCAL VOC 2012 [5] and extra annotations from SBD [9]. The total 20 categories are evenly partitioned into four splits, and the model training is conducted in a cross-validation manner. Specifically, three splits are selected in the training process and the remaining one is used for testing. During inference, 1000 support-query pairs are randomly sampled for evaluation [23]. Different from Pascal-5ⁱ, MSCOCO [17] is a large-scale dataset with 82,081 images in the training set. Following FWB [20], the overall 80 classes from MSCOCO [17] are evenly divided into four splits with the same cross-validation strategy. For more stable results, we randomly sample 20,000 pairs during evaluation [29].

We use mean intersection-over-union (mIoU) as a primary evaluation metric for the ablation study, as it is commonly recognized in image segmentation. In addition, for more consistent comparisons, results of foreground-background IoU (FB-IoU) are also reported.

5.2. Implementation Details

We adopt ResNet [11] as the backbone network [40], and we concatenate block2 and block3 to generate feature maps. We train the model with a SGD optimizer on Pascal-5ⁱ for 200 epochs and COCO-20ⁱ for 50 epochs. We set the initial learning rate to 0.0025 with batch size 4 on Pascal-5ⁱ, and 0.005 with batch size 8 on COCO-20ⁱ. The number of iterations in the SGC module is set to 10 during training and 5 during inference. We use data augmentation during training - input images are transformed with random scale, horizontal flip and rotation from [-10, 10], and then all images are cropped to 473 × 473 (Pascal) or 641 × 641 (COCO) as training samples. We implement our model using Pytorch and run our experiments on a workstation with Nvidia Tesla V100 GPU. To increase the variance of the cosine measurement, we remove the ReLU layer before both support and query features, so the similarity metric is bounded in [-1, 1] rather than [0, 1].

5.3. Ablation Study

To verify the effectiveness of the proposed modules, we implement extensive ablation studies with a ResNet-50 backbone on Pascal-5ⁱ. We use floating point operations (FLOPs) to represent the amount of computation, and take both addition and multiplication into account.

Number of Superpixel Centroids. To explore the effect of the number of superpixel centroids, we conduct experiments with different N_{max} in 1-shot segmentation. As stated in Section 4.3, N_{max} is a hyper-parameter set to regulate the maximum number of prototypes. Particularly, when $N_{max} = 1$, the prototype generation process degrades to the masked average pooling. As shown in Table 1, when N_{max} equals 5, ASGNet achieves the best results in split 0 and 2, as well as the best mean performance, which demonstrates the validity of superpixel-guided clustering. The results get improved gradually when N_{max} increases from 1 to 5,

and then decrease slightly after that, denoting that excessive prototypes bring negative effect and are prone to overfitting. Finally, we choose N_{max} as 5 for both 1-shot and 5-shot segmentation. In addition, we analyzed the impact of the proposed adaptive scheme on the number of superpixel centroids N_{sp} and show the results in Table 2. Compared with using the fixed number (5) of superpixels, the results show that the adaptive setting (Eqn. 7) can reduce redundant computation while obtaining better performance.

N_{max}	s-0	s-1	s-2	s-3	mean
1	57.40	67.52	55.69	54.36	58.74
3	57.92	67.86	55.86	53.74	58.85
5	58.84	67.86	56.79	53.66	59.29
7	58.84	68.11	55.85	54.00	59.20
9	57.61	67.97	56.68	53.27	58.88

Table 1. Ablation study on the maximum number of superpixel centroids N_{max} . $s-x$ denotes different cross-validation splits.

Centroids	s-0	s-1	s-2	s-3	mean
Fixed	57.87	67.62	57.10	53.61	59.05
Adaptive	58.84	67.86	56.79	53.66	59.29

Table 2. Performance comparison between fixed and adaptive superpixel centroids.

SGC and GPA. To demonstrate the effectiveness of proposed SGC and GPA modules, we conduct diagnostic experiments on prototype generation and matching. We first implement a baseline model with the single prototype learning proposed in PFENet [29]. Then, we introduce SGC to generate multiple prototypes, and fuse them in a dense manner with reference to PMMs [34]. Finally, we replace the prototype expansion with the proposed allocation scheme. For ablations on the SGC module, as shown in Table 3, in the 1-shot setting we observe that replacing mean average pooling (MAP) with SGC can worsen the representations and lead to degraded performance, but the model benefits in the 5-shot scenario. We deem the main reason is that excessive prototypes become highly similar on a single support sample, and cosine distance can not distinguish them apart. Finally, when GPA module is adopted for prototype matching, the performance improves by 2.70% compared to the prototype expansion, and also the computational overhead is much lower.

K-shot Fusion Setting. As mentioned in Section 4.5, ASGNet is able to tackle the k -shot learning problem without collapsing the support features. To explore the effect of different fusion methods, we compare our k -shot setting with two other commonly used solutions: 1) feature average fusion [21] and 2) attention weighted summation [40]. For simplicity, experiments of this ablation study are only

Generation		Matching		mIoU		FLOPs Δ
MAP	SGC	Expand	GPA	1-shot	5-shot	
✓		✓		58.96	60.19	0.9G
	✓	✓		58.31	61.24	K*8.5G+0.5G
	✓		✓	59.29	63.94	K*5.5M+0.9G

Table 3. Ablation study on prototype generation (Masked average pooling (MAP) vs. SGC) and matching (Expand vs. GPA). FLOPs Δ denotes the computational cost from prototype matching process, and K is the adaptive number of prototype ($K \leq 5$).

conducted in Pascal-5⁰. As reported in Table 4, our simple strategy achieves the best performance and the largest increment over the 1-shot baseline (4.82%) without additional computation. On the contrary, attention-based fusion requires a large amount of computation, and has limited performance improvement. This demonstrates that the GPA module is highly effective when given a large number of selections.

5-shot setting	mIoU	IoU Δ	FLOPs Δ
1-shot baseline	58.84	-	-
Attention	60.69	1.85	42.6G
Feature-avg	62.10	3.26	3.7M
Ours	63.66	4.82	None

Table 4. Ablation study of different k -shot fusion strategy.

5.4. Comparison to State-of-the-art

Pascal-5ⁱ. Comparisons to state-of-the-art methods are shown in Table 5 and 6, where two different metrics are adopted. In Table 5, with ResNet-101 as the backbone, ASGNet outperforms recent methods with a considerable margin of 2.40% in 5-shot segmentation, while being on par with state-of-the-arts under the 1-shot setting. In Table 6, we further make comparisons on FB-IoU and the number of trainable parameters. Once again, the proposed ASGNet achieves significant improvement over state-of-the-arts in 5-shot setting (75.2% vs.73.9%), and it also has the largest performance increment of 5.0% over 1-shot result. In addition, we have far fewer trainable parameters than other methods. In Figure 7, we show some representative segmentation examples. We observe that the proposed ASGNet can generate accurate segmentation results even when there are large variations in appearance and pose between support and query images.

COCO-20ⁱ. In Table 7, we present the performance comparison of mean IoU and FB-IoU on COCO-20ⁱ. As can be seen, ASGNet achieves state-of-the-art results in both 1-shot and 5-shot setting in terms of the mean IoU, and it substantially outperforms previous methods. Particularly, ASGNet achieves a margin of 3.98% and 6.96% higher

Backbone	Methods	1-shot					5-shot					Δ
		s-0	s-1	s-2	s-3	mean	s-0	s-1	s-2	s-3	mean	
VGG-16	OSLSM [23]	33.60	55.30	40.90	33.50	40.80	35.90	58.10	42.70	39.10	43.95	3.15
	co-FCN [21]	36.70	50.60	44.90	32.40	41.10	37.50	50.00	44.10	33.90	41.40	0.30
	AMP [24]	41.90	50.20	46.70	34.40	43.40	40.30	55.30	49.90	40.10	46.40	3.00
	SG-One [42]	40.20	58.40	48.40	38.40	46.30	41.90	58.60	48.60	39.40	47.10	0.80
	PANet [32]	42.30	58.00	51.10	41.20	48.10	51.80	64.60	59.80	46.50	55.70	7.60
	FWB [20]	47.04	59.64	52.61	48.27	51.90	50.87	62.86	56.48	50.09	55.08	3.18
ResNet50	CANet \dagger [40]	52.50	65.90	51.30	51.90	55.40	55.50	67.80	51.90	53.20	57.10	1.70
	PGNet \dagger [39]	56.00	66.90	50.60	50.40	56.00	57.70	68.70	52.90	54.60	58.50	2.50
	RPMMs [34]	55.15	66.91	52.61	50.68	56.34	56.28	67.34	54.52	51.00	57.30	0.96
	SimPropNet [8]	54.82	67.33	54.52	52.02	57.19	57.20	68.50	58.40	56.05	60.04	2.85
	PPNet [18]	47.83	58.75	53.80	45.63	51.50	58.39	67.83	64.88	56.73	61.96	10.46
	PFENet [29]	61.70	69.50	55.40	56.30	60.80	63.10	70.70	55.80	57.90	61.90	1.10
	ASGNet (ours)	58.84	67.86	56.79	53.66	59.29	63.66	70.55	64.17	57.38	63.94	4.65
ResNet101	FWB [20]	51.30	64.49	56.71	52.24	56.19	54.84	67.38	62.16	55.30	59.92	3.73
	DAN \dagger [31]	54.70	68.60	57.80	51.60	58.20	57.90	69.00	60.10	54.90	60.50	2.30
	PFENet [29]	60.50	69.40	54.40	55.90	60.10	62.80	70.40	54.90	57.60	61.40	1.30
	ASGNet (ours)	59.84	67.43	55.59	54.39	59.31	64.55	71.32	64.24	57.33	64.36	5.05

Table 5. Comparison with state-of-the-art on Pascal-5ⁱ. \dagger indicates multi-scale inference is adopted. Δ means increment over 1-shot segmentation result.

Methods	FB-IoU		Δ	#Params
	1-shot	5-shot		
OSLSM [23]	61.3	61.5	0.2	272.6M
co-FCN [21]	60.1	60.2	0.1	34.2M
AMP [24]	62.2	63.8	1.6	14.9M
SG-One [42]	63.1	65.9	2.8	19.0M
CANet \dagger [40]	66.2	69.6	3.4	10.5M
PGNet \dagger [39]	69.9	70.5	0.6	-
PANet [32]	66.5	70.7	4.2	14.7M
DAN \dagger [31]	71.9	72.3	0.4	-
SimPropNet [8]	73.0	72.9	-0.1	-
PFENet [29]	73.3	73.9	0.6	10.8M
ASGNet (RN-50)	69.2	74.2	5.0	10.4M
ASGNet (RN-101)	71.7	75.2	3.5	10.4M

Table 6. Comparison of FB-IoU and the number of trainable parameters on Pascal-5ⁱ. \dagger indicates multi-scale inference is adopted. Δ means increment over 1-shot segmentation result.

Backbone	Methods	mIoU		FB-IoU	
		1-shot	5-shot	1-shot	5-shot
ResNet101	FWB [20]	21.19	23.65	-	-
	DAN [31]	24.20	29.60	62.30	63.90
	PFENet [29]	32.40	37.40	58.60	61.90
ResNet50	RPMMs [34]	30.58	35.52	-	-
	ASGNet	34.56	42.48	60.39	66.96

Table 7. Comparison with state-of-the-arts on COCO-20ⁱ.

mean IoU over RPMMs [34] in 1-shot and 5-shot segmentation respectively. Also, our ASGNet obtains competitive 1-shot results and top-performing 5-shot results with respect to FB-IoU. These results demonstrate that the proposed method is capable of handling more complex cases, as MSCOCO is a much more challenging dataset with diverse samples and categories.




Figure 7. Qualitative examples of 1-shot segmentation on the Pascal-5ⁱ. Best viewed in color and zoom in.

6. Conclusions

In this paper we propose ASGNet for few-shot image segmentation. Targeting the limitations of existing single prototype based models, we introduce two new modules, named Superpixel-guided Clustering (SGC) and Guided Prototype Allocation (GPA), for adaptive prototype learning and allocation. Concretely, SGC aggregates similar feature vectors with feature-based superpixel clustering, and GPA aims to allocate the most relevant prototype to each query feature element by measuring the similarity with cosine distance. Extensive experiments and ablation studies have demonstrated the superiority of ASGNet, and we achieve state-of-the-art performance on both Pascal-5ⁱ and COCO-20ⁱ without any additional post-processing steps.

Acknowledgements This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No.2020R1A2C101215911).

References

- [1] Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *PAMI*, 34(11):2274–2282, 2012.
- [2] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, pages 801–818, 2018.
- [3] Sungha Choi, Joanne T Kim, and Jaegul Choo. Cars can't fly up in the sky: Improving urban-scene segmentation via height-driven attention networks. In *CVPR*, pages 9373–9383, 2020.
- [4] Nanqing Dong and Eric P Xing. Few-shot semantic segmentation with prototype learning. In *BMVC*, 2018.
- [5] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 88(2):303–338, 2010.
- [6] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017.
- [7] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *CVPR*, pages 3146–3154, 2019.
- [8] Siddhartha Gairola, Mayur Hemani, Ayush Chopra, and Balaji Krishnamurthy. Simpropnet: Improved similarity propagation for few-shot image segmentation. In *IJCAI*, 2020.
- [9] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *ICCV*, pages 991–998. IEEE, 2011.
- [10] Junjun He, Zhongying Deng, Lei Zhou, Yali Wang, and Yu Qiao. Adaptive pyramid context network for semantic segmentation. In *CVPR*, pages 7519–7528, 2019.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [12] Benjamin Irving. maskslc: regional superpixel generation with application to local pathology characterisation in medical images. *arXiv preprint arXiv:1606.09518*, 2016.
- [13] Varun Jampani, Deqing Sun, Ming-Yu Liu, Ming-Hsuan Yang, and Jan Kautz. Superpixel sampling networks. In *ECCV*, pages 352–368, 2018.
- [14] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.
- [15] Xia Li, Zhisheng Zhong, Jianlong Wu, Yibo Yang, Zhouchen Lin, and Hong Liu. Expectation-maximization attention networks for semantic segmentation. In *ICCV*, pages 9167–9176, 2019.
- [16] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 2117–2125, 2017.
- [17] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755. Springer, 2014.
- [18] Yongfei Liu, Xiangyi Zhang, Songyang Zhang, and Xuming He. Part-aware prototype network for few-shot semantic segmentation. In *ECCV*, 2020.
- [19] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, pages 3431–3440, 2015.
- [20] Khoi Nguyen and Sinisa Todorovic. Feature weighting and boosting for few-shot segmentation. In *ICCV*, pages 622–631, 2019.
- [21] Kate Rakelly, Evan Shelhamer, Trevor Darrell, Alyosha Efros, and Sergey Levine. Conditional networks for few-shot semantic segmentation. In *ICLR Workshop*, 2018.
- [22] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017.
- [23] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots. One-shot learning for semantic segmentation. In *BMVC*, 2017.
- [24] Mennatullah Siam, Boris N Oreshkin, and Martin Jagersand. Amp: Adaptive masked proxies for few-shot segmentation. In *ICCV*, pages 5249–5258, 2019.
- [25] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, pages 4077–4087, 2017.
- [26] David Stutz, Alexander Hermans, and Bastian Leibe. Superpixels: An evaluation of the state-of-the-art. *Computer Vision and Image Understanding*, 166:1–27, 2018.
- [27] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, pages 1199–1208, 2018.
- [28] Andrew Tao, Karan Sapra, and Bryan Catanzaro. Hierarchical multi-scale attention for semantic segmentation. *arXiv preprint arXiv:2005.10821*, 2020.
- [29] Z Tian, H Zhao, M Shu, Z Yang, R Li, and J Jia. Prior guided feature enrichment network for few-shot segmentation. *PAMI*, 2020.
- [30] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *NeurIPS*, pages 3630–3638, 2016.
- [31] Haochen Wang, Xudong Zhang, Yutao Hu, Yandan Yang, Xianbin Cao, and Xiantong Zhen. Few-shot semantic segmentation with democratic attention networks. In *ECCV*, 2020.
- [32] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *ICCV*, pages 9197–9206, 2019.
- [33] Yu-Xiong Wang and Martial Hebert. Learning to learn: Model regression networks for easy small sample learning. In *ECCV*, pages 616–634. Springer, 2016.
- [34] Boyu Yang, Chang Liu, Bohao Li, Jianbin Jiao, and Qixiang Ye. Prototype mixture models for few-shot semantic segmentation. In *ECCV*, 2020.
- [35] Maoke Yang, Kun Yu, Chi Zhang, Zhiwei Li, and Kuiyuan Yang. Denseaspp for semantic segmentation in street scenes. In *CVPR*, pages 3684–3692, 2018.

- [36] Xianghui Yang, Bairun Wang, Kaige Chen, Xinchu Zhou, Shuai Yi, Wanli Ouyang, and Luping Zhou. Brinet: Towards bridging the intra-class and inter-class gaps in one-shot segmentation. In *BMVC*, 2020.
- [37] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015.
- [38] Yuhui Yuan and Jingdong Wang. Ocnet: Object context network for scene parsing. *arXiv preprint arXiv:1809.00916*, 2018.
- [39] Chi Zhang, Guosheng Lin, Fayao Liu, Jiushuang Guo, Qingyao Wu, and Rui Yao. Pyramid graph networks with connection attentions for region-based one-shot semantic segmentation. In *ICCV*, pages 9587–9595, 2019.
- [40] Chi Zhang, Guosheng Lin, Fayao Liu, Rui Yao, and Chunhua Shen. Canet: Class-agnostic segmentation networks with iterative refinement and attentive few-shot learning. In *CVPR*, pages 5217–5226, 2019.
- [41] Fan Zhang, Yanqin Chen, Zhihang Li, Zhibin Hong, Jingtuo Liu, Feifei Ma, Junyu Han, and Errui Ding. Acfnets: Attentional class feature network for semantic segmentation. In *ICCV*, pages 6798–6807, 2019.
- [42] Xiaolin Zhang, Yunchao Wei, Yi Yang, and Thomas S Huang. Sg-one: Similarity guidance network for one-shot semantic segmentation. *IEEE Transactions on Cybernetics*, 2020.
- [43] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, pages 2881–2890, 2017.
- [44] Zhen Zhu, Mengde Xu, Song Bai, Tengpeng Huang, and Xiang Bai. Asymmetric non-local neural networks for semantic segmentation. In *ICCV*, pages 593–602, 2019.