

Beyond the Prototype: Divide-and-conquer Proxies for Few-shot Segmentation

Chunbo Lang*, Binfei Tu*, Gong Cheng†, Junwei Han

School of Automation, Northwestern Polytechnical University, Xi’an, China
 {langchunbo, binfeitu}@mail.nwpu.edu.cn, {gcheng, jhan}@nwpu.edu.cn

Abstract

Few-shot segmentation, which aims to segment unseen-class objects given only a handful of densely labeled samples, has received widespread attention from the community. Existing approaches typically follow the prototype learning paradigm to perform meta-inference, which fails to fully exploit the underlying information from support image-mask pairs, resulting in various segmentation failures, *e.g.*, incomplete objects, ambiguous boundaries, and distractor activation. To this end, we propose a simple yet versatile framework in the spirit of divide-and-conquer. Specifically, a novel self-reasoning scheme is first implemented on the annotated support image, and then the coarse segmentation mask is *divided* into multiple regions with different properties. Leveraging effective masked average pooling operations, a series of support-induced proxies are thus derived, each playing a specific role in *conquering* the above challenges. Moreover, we devise a unique parallel decoder structure that integrates proxies with similar attributes to boost the discrimination power. Our proposed approach, named divide-and-conquer proxies (DCP), allows for the development of appropriate and reliable information as a guide at the “episode” level, not just about the object cues themselves. Extensive experiments on PASCAL-5ⁱ and COCO-20ⁱ demonstrate the superiority of DCP over conventional prototype-based approaches (up to 5~10% on average), which also establishes a new state-of-the-art. Code is available at github.com/chunbolang/DCP.

1 Introduction

Benefiting from the superiority of deep convolutional neural networks, various computer vision tasks have made tremendous progress over the past few years, including image classification [He *et al.*, 2016], object detection [Ren *et al.*, 2015], and semantic segmentation [Long *et al.*, 2015], to name a

*Equal contribution.

†Corresponding author.

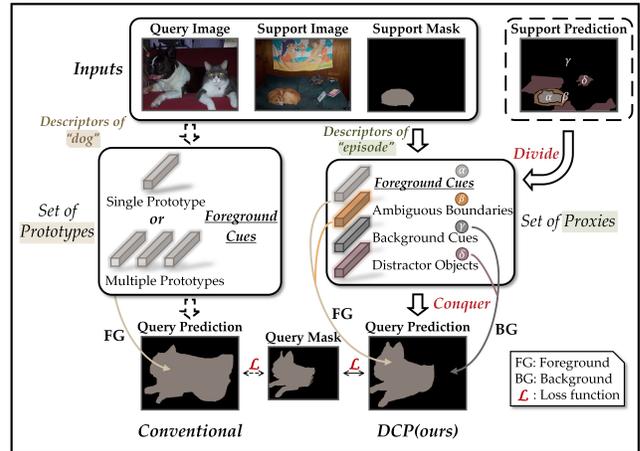


Figure 1: Comparison of conventional FSS approaches (left part) and our DCP approach (right part). **Conventional approaches** typically leverage support features and corresponding masks to generate single or multiple foreground *prototypes* for guiding query image segmentation, while **DCP** utilizes the visual reasoning results of support images to derive many different “*proxies*” from a broader episodic perspective (*divide*), each of which plays a specific role (*conquer*). Overall, such operations could provide pertinent and reliable guidance, including foreground/background cues, ambiguous boundaries, and distractor objects, rather than merely object-related information.

few. However, such an effective technique also has its inherent limitation: the strong demand for a considerable number of annotated samples from well-established datasets to achieve satisfactory performance [Deng *et al.*, 2009]. For the tasks such as semantic segmentation that performs dense predictions on given images, the acquisition of training sets with sufficient data costs even more, substantially hindering the development of deep learning systems. Few-shot learning (FSL) has emerged as a promising research direction to address this issue, which intends to learn a generic model that can recognize new concepts with very limited information available [Vinyals *et al.*, 2016] [Lang *et al.*, 2022].

In this paper, we undertake the few-shot segmentation (FSS) task that can be regarded as a natural application of FSL technology in the field of image segmentation [Shaban *et al.*, 2017]. Specifically, FSS models aim to retrieve the

foreground regions of a specific category in the query image while only providing a handful of support images with corresponding annotated masks. Prevalent FSS approaches [Zhang *et al.*, 2019; Tian *et al.*, 2022] follow the prototype learning paradigm to perform meta-reasoning, where single or multiple prototypes are generated by the average pooling operation [Zhang *et al.*, 2020], furnishing query object inference with essential foreground cues. However, their segmentation capabilities could be fragile when confronted with daunting few-shot tasks with factors such as blurred object boundaries and irrelevant object interference. We argue that one possible reason for these failures lies in the underexplored information from support image-mask pairs, and it is far from enough to rely solely on squeezed object-related features to guide segmentation. Taking the 1-shot segmentation task in Figure 1 as an example, conventional approaches tend to confuse several distractor objects (*e.g.*, *cat* and *sofa*) and yield inaccurate segmentation boundaries (*e.g.*, *dog*'s paw and head), revealing the limitations of the scheme based on such descriptors.

To alleviate the abovementioned problems, we propose a simple yet versatile framework in the spirit of divide-and-conquer (refer to the right part of Figure 1), where the descriptors for guidance are derived from a broader *episodic* perspective, rather than just the objects themselves. Concretely, we first implement a novel self-reasoning scheme on the annotated support image, and then *divide* the coarse segmentation mask into multiple regions with different properties. Using effective masked average pooling operations, a series of support-induced proxies are thus generated, each playing a specific role in *conquering* the typical challenges of daunting FSS tasks. Furthermore, we devise a unique parallel decoder structure in which proxies with similar attributes are integrated to boost the discrimination power. Compared with conventional prototype-based approaches, our proposed framework, named divide-and-conquer proxies (DCP), allows for the development of appropriate and reliable information as a guide at the “episode” level, not just about the object cues.

Our primary contributions can be summarized as follows:

- We present a simple yet versatile FSS framework in the spirit of divide-and-conquer, where the descriptors for guidance are derived from a broader episodic perspective to tackle the typical challenges of daunting tasks.
- Compared with conventional prototype-based approaches, the proposed framework allows for the development of appropriate and reliable information, rather than merely the object cues themselves, providing a fresh insight for future works.
- Extensive experiments on two standard benchmarks demonstrate the superiority of our framework over the baseline approach (up to 5~10% on average), establishing new state-of-the-arts in few-shot literature.

2 Related Works

Semantic Segmentation. Semantic segmentation is a fundamental and essential computer vision task, the goal of which is to classify each pixel in the image according to a predefined set of semantic categories. Most advanced approaches follow

fully convolutional networks (FCNs) [Long *et al.*, 2015] to perform dense predictions, vastly boosting the segmentation performance [Zhao *et al.*, 2017]. On top of that, the rapid development of this field has also brought some impressive techniques, such as dilated convolutions [Chen *et al.*, 2017] and attention mechanism [Huang *et al.*, 2019]. In this work, we adopt the Atrous Spatial Pyramid Pooling (ASPP) module [Chen *et al.*, 2017] based on dilated convolution to build the decoder, enjoying a large receptive field and multi-scale feature aggregation.

Few-shot Learning. Few-shot learning (FSL) is proposed to tackle the generalization problem of conventional deep networks, where meta-knowledge is shared across tasks to facilitate the recognition of new concepts with very little support information. Representative FSL approaches can be grouped into three categories: optimization-based method [Finn *et al.*, 2017], hallucination-based method [Chen *et al.*, 2019], and metric-based method [Vinyals *et al.*, 2016]. Our work closely relates to the prototypical network (PN) [Snell *et al.*, 2017] assigned to the third category. PN directly leverages the feature representations (*i.e.*, prototypes) computed through global average pooling operations to perform *non-parametric* nearest neighbor classification. Our DCP, on the other hand, utilizes foreground/background prototypes for dense feature comparison [Zhang *et al.*, 2019], providing essential information for subsequent *parametric* query image segmentation.

Few-shot Segmentation. Few-shot segmentation (FSS) has emerged as a promising research direction to tackle the dense prediction problem in a low-data regime. Existing approaches typically employ a two-branch structure, in which the annotation information travels from the support branch to the query branch [Zhang *et al.*, 2019] [Siam *et al.*, 2019]. Generally speaking, how to achieve more effective information interaction between the two branches is the primary concern in this field. OSLSM [Shaban *et al.*, 2017], the pioneering work of FSS, proposed to generate classifier weights for query image segmentation in the conditioning (support) branch. Whereafter, the prototype learning paradigm [Snell *et al.*, 2017] was introduced into state-of-the-art approaches for better information exchange. SG-One [Zhang *et al.*, 2020] computed the affinity scores between the prototype and query features, yielding the spatial similarity map as a guide. CANet [Zhang *et al.*, 2019] proposed to conduct a novel dense comparison operation based on the prototype, which exhibited decent segmentation performance under few-shot settings. However, these methods are limited to the utilization of features of foreground objects (*i.e.*, prototype) to facilitate segmentation and fail to provide pertinent and reliable guidance from a broader *episodic* perspective. Our method draws on the idea of prototype learning while proposing a set of support-induced proxies that go beyond it.

3 Problem Setup

Few-shot segmentation (FSS) aims to segment the objects of a specific category in the query image using only a few labeled samples. To improve the generalization ability for unseen classes, existing approaches widely adopt the meta-learning strategy for training, also known as episodic learn-

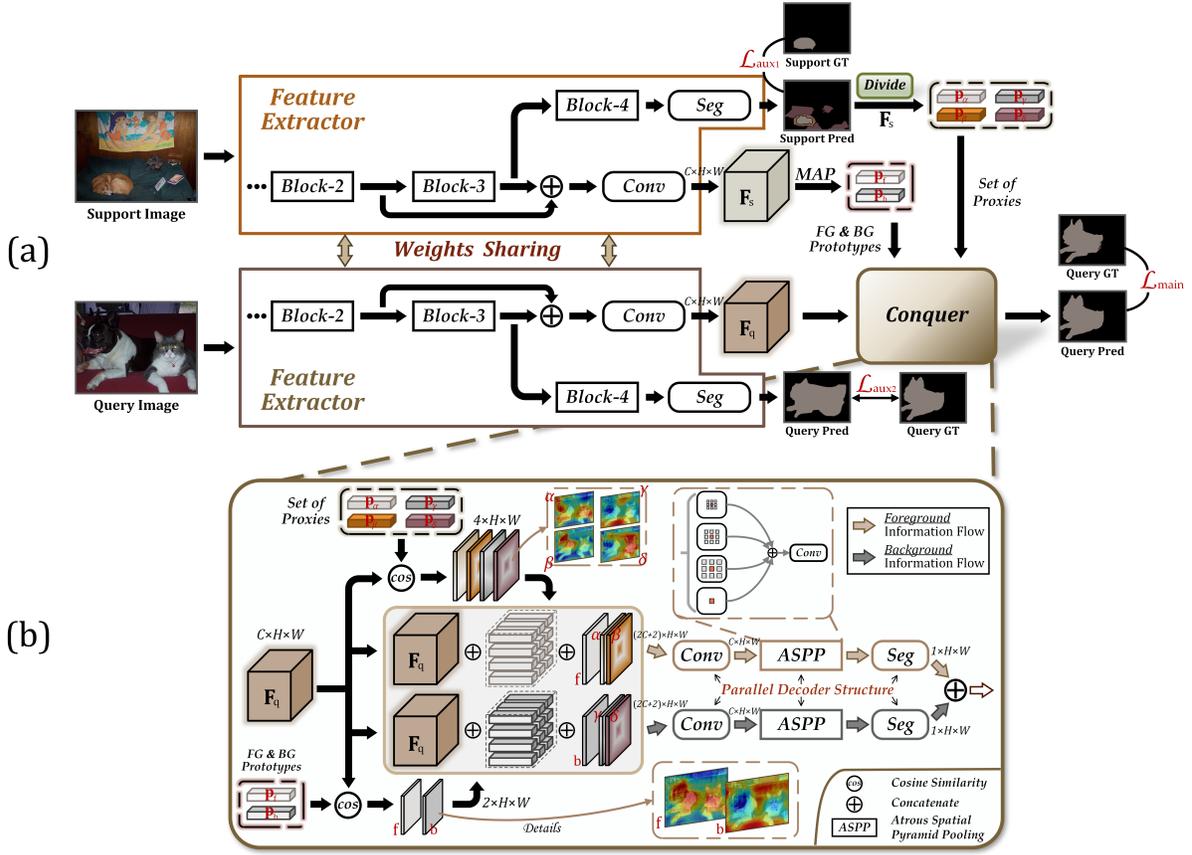


Figure 2: DCP for 1-shot segmentation. (a) Overall pipeline of the proposed approach. (b) Implementation details of the conquer module.

ing [Vinyals *et al.*, 2016]. Specifically, given a training set $\mathcal{D}_{\text{train}}$ and a testing set $\mathcal{D}_{\text{test}}$ that are disjoint with regard to object classes, a series of *episodes* are sampled from them to simulate the few-shot scenario. For a K -shot segmentation task, each episode is composed of **1**) a support set $\mathcal{S} = \{(\mathbf{I}_s^k, \mathbf{M}_{s,c}^k)\}_{c \in \mathcal{C}_{\text{episode}}, k=1, \dots, K}$, where \mathbf{I}_s^k is the k -th support image, $\mathbf{M}_{s,c}^k$ is the corresponding mask for class c , and $\mathcal{C}_{\text{episode}}$ represents the class set of the given episode; and **2**) a query set $\mathcal{Q} = \{\mathbf{I}_q, \mathbf{M}_{q,c}\}$ where \mathbf{I}_q is the query image and $\mathbf{M}_{q,c}$ is the ground-truth mask *available* during training while *unknown* during testing.

4 Method

In this section, we propose a simple yet effective few-shot segmentation framework, termed divide-and-conquer proxies (DCP). The unique feature of our approach, as also the source of improvements, lies in the spirit of divide-and-conquer. The *divide* module generates a set of proxies according to the prediction and ground-truth mask of the support image; the *conquer* module utilizes these distinguishing proxies to provide essential segmentation cues for the query branch, as depicted in Figure 2.

4.1 Divide

Self-reasoning Scheme. Following the paradigm of advanced FSS approaches [Zhang *et al.*, 2019], we freeze the backbone network to boost generalization capabilities. Taking ResNet [He *et al.*, 2016] for example, the intermediate feature maps after each convolutional block of the support branch are represented as $\{\mathbf{F}_s^b\}_{b=1}^B$. We use the high-level features generated by the last block, *i.e.*, *block4*, to perform the self-reasoning scheme:

$$\hat{\mathbf{M}}_s^{\text{aux}} = f_{\text{seg}}(\mathbf{F}_s^4), \quad (1)$$

where $f_{\text{seg}}(\cdot)$ is a lightweight segmentation decoder with residual connections, including three convolutions with 256 filters, one convolution with 2 filters, and an $\arg \max(\cdot)$ operation for producing coarse prediction masks $\hat{\mathbf{M}}_s^{\text{aux}} \in \{0, 1\}^{H \times W}$. H and W denote the height and weight respectively, and their product $H \times W$ is the minimum resolution of all feature maps. Unlike the widely adopted prototype-guided segmentation framework, such a *mask-agnostic* self-reasoning scheme is more efficient and can also provide reliable auxiliary information.

Proxy Acquisition. Given the coarse prediction $\hat{\mathbf{M}}_s^{\text{aux}}$ and down-sampled ground-truth mask \mathbf{M}_s of the support image,

* B denotes the total number of convolutional blocks contained in the backbone network, which is equal to 4 for ResNet.

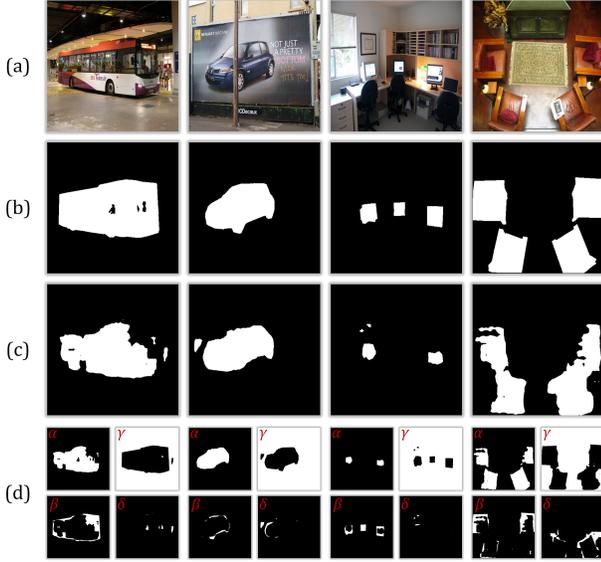


Figure 3: Visualization of the masks for generating proxies. From top to bottom, each row represents (a) raw (support) images, (b) ground-truth masks, (c) prediction masks, and (d) masks for generating various proxies, respectively.

we first evaluate their differences to derive the corresponding mask (*i.e.*, valid region) $\mathbf{M}_{m \in \mathcal{P}_1}$ of each proxy where $\mathcal{P}_1 = \{\alpha, \beta, \gamma, \delta\}$ denotes the index set of proxies:

$$\begin{cases} \mathbf{M}_\alpha^{(x,y)} = \mathbb{1}[\hat{\mathbf{M}}_s^{\text{aux};(x,y)} = \mathbf{M}_s^{(x,y)} = 1] \\ \mathbf{M}_\beta = \mathbf{M}_s - \mathbf{M}_\alpha \\ \mathbf{M}_\gamma^{(x,y)} = \mathbb{1}[\hat{\mathbf{M}}_s^{\text{aux};(x,y)} = \mathbf{M}_s^{(x,y)} = 0] \\ \mathbf{M}_\delta = \mathbf{G} - \mathbf{M}_s - \mathbf{M}_\gamma \end{cases}, \quad (2)$$

where (x, y) indexes the spatial location of the mask, $\mathbb{1}(\cdot)$ is an indicator function that outputs 1 if the condition is satisfied or 0 otherwise, and $\mathbf{G} \in \mathbb{R}^{H \times W}$ represents an all 1's matrix with the same shape as \mathbf{M}_m . \mathbf{M}_α and \mathbf{M}_β aggregate the main and auxiliary foreground information respectively, generally corresponding to the body and boundary regions of objects; while \mathbf{M}_γ and \mathbf{M}_δ gather the primary and secondary background information respectively, which typically correspond to the regions of the generalized background and distractor objects. As described above, these masks are generated according to the segmentation results of the model, therefore the main (primary) and auxiliary (secondary) status may change when encountering complex tasks. The extreme case is $\mathbf{M}_\beta = \mathbf{M}_s$, but even so, the guidance information is still complete. To better understand the effect of each mask, we present several examples in Figure 3.

Then, we conduct masked average pooling operations [Zhang *et al.*, 2020] on the support features $\mathbf{F}_s \in \mathbb{R}^{C \times H \times W}$ and masks \mathbf{M}_m computed above to encode the set of proxies $\{\mathbf{p}_m\}$:

$$\mathbf{p}_m = \frac{\sum_{x,y} \mathbf{F}_s^{(x,y)} \cdot \mathbf{M}_m^{(x,y)}}{\sum_{x,y} \mathbf{M}_m^{(x,y)}}, \quad (3)$$

where $\mathbf{p}_m \in \mathbb{R}^C$ represents a specific proxy in the set, playing some roles in conquering the typical challenges of FSS tasks.

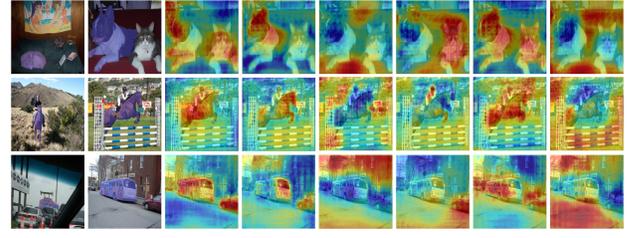


Figure 4: Visualization of activation maps generated by proxies and prototypes. From left to right, each column represents support images, query images, \mathbf{A}_α , \mathbf{A}_β, \dots , \mathbf{A}_f , \mathbf{A}_b , respectively.

Prototype Acquisition. Similar to the acquisition method of proxies (Eq. (3)), the set of prototypes $\{\mathbf{p}_n\}_{n \in \mathcal{P}_2}$ are also generated using the masked average pooling operation where $\mathcal{P}_2 = \{f, b\}$ represents the index set of prototypes:

$$\mathbf{p}_n = \frac{\sum_{x,y} \mathbf{F}_s^{(x,y)} \cdot \mathbb{1}[\mathbf{M}_s^{(x,y)} = j]}{\sum_{x,y} \mathbb{1}[\mathbf{M}_s^{(x,y)} = j]}, \quad (4)$$

where $j=1$ if the foreground prototype \mathbf{p}_f is evaluated and $j=0$ if the background one \mathbf{p}_b is calculated.

4.2 Conquer

Feature Matching. Dense feature comparison is a common approach in the FSS literature that compares the query feature maps \mathbf{F}_q with the global feature vector (*i.e.*, prototype) at all spatial locations. In view of its effectiveness, we follow the same technical route and extend it to a dual-prototype setting (see Figure 3(b)), which can be defined as:

$$\mathbf{F}_{q;n}^{\text{tmp}} = \bigoplus (\mathbf{F}_q, \zeta_{h \times w}(\mathbf{p}_n)), \quad (5)$$

where $\zeta_{h \times w}(\cdot)$ expands the given vector to a spatial size $h \times w$ and \bigoplus represents the concatenation operation along channel dimensions. The superscript “tmp” denotes *temporary*, and the subscript $n \in \mathcal{P}_2$ indicates the index of the prototype.

Feature Activation. Given the proxies $\{\mathbf{p}_m\}_{m \in \mathcal{P}_1}$ and prototypes $\{\mathbf{p}_n\}_{n \in \mathcal{P}_2}$ computed by Eqs. (3-4), we evaluate the cosine similarity between each of them and query features at each spatial location respectively:

$$\mathbf{A}_l^{(x,y)} = \frac{\mathbf{p}_l^\top \cdot \mathbf{F}_q^{(x,y)}}{\|\mathbf{p}_l\| \cdot \|\mathbf{F}_q^{(x,y)}\|}, \quad l \in \mathcal{P}_1 \cup \mathcal{P}_2, \quad (6)$$

where $\mathbf{A}_l \in \mathbb{R}^{H \times W}$ denotes the activation map of the l -th item. Figure 4 presents some examples to illustrate the region of interest for each feature vector \mathbf{p}_l . Actually, there is an alternative approach to propagate information about feature vectors $\{\mathbf{p}_l\}_{l \in \mathcal{P}_1 \cup \mathcal{P}_2}$ to the query branch, that is, to perform the same operation as Eq. (5). Such a scheme, however, introduces additional computational overhead and performance degradation, as discussed in § 5.3.

Parallel Decoder Structure. Most previous works abandon the prototype \mathbf{p}_b due to the complexity of background regions in support images. We argue that specifically designing the network to capture foreground and background features separately could have a positive effect on performance,

whereas a simple fusion strategy would be counterproductive. To this end, a unique parallel decoder structure (PDS) is devised in which proxies/prototypes with similar attributes are integrated to boost the discrimination power:

$$\mathbf{F}_{q;n}^{\text{refined}} = \bigoplus \left(\mathbf{F}_{q;n}^{\text{tmp}}, \mathbf{A}_n, \{\mathbf{A}_m\}_{m \in \mathcal{P}_{\text{tmp}}}, \right), \quad (7)$$

where \mathcal{P}_{tmp} is a temporary set of vector indexes, representing $\{\alpha, \beta\}$ when foreground features $\mathbf{F}_{q:f}^{\text{refined}}$ are evaluated and $\{\gamma, \delta\}$ when background ones $\mathbf{F}_{q:b}^{\text{refined}}$ are computed. These two refined query features are then fed into the corresponding decoder network respectively, and the output single-channel probability maps are concatenated to generate the final segmentation result.

4.3 Training Loss

To guarantee the richness of the information provided by proxies while preventing extreme situations (§ 4.1), we introduce additional constraints on the prediction results of self-reasoning schemes for both support and query branches. The total loss for each episode can be written as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{main}} + \lambda_1 \mathcal{L}_{\text{aux1}} + \lambda_2 \mathcal{L}_{\text{aux2}}, \quad (8)$$

where \mathcal{L} represents the binary cross-entropy (BCE) loss function. λ_1 and λ_2 , the balancing coefficients, are set to 1.0.

5 Experiments

5.1 Setup

Datasets. We evaluate the proposed approach on two standard FSS benchmarks: PASCAL-5ⁱ [Shaban *et al.*, 2017] and COCO-20ⁱ [Nguyen and Todorovic, 2019]. The former is created from PASCAL VOC 2012 [Everingham *et al.*, 2010] with additional mask annotations from SDS [Hariharan *et al.*, 2011], consisting of 20 semantic categories evenly divided into 4 folds: $\{5^i\}_{i=0}^3$; while the latter, built from MS COCO [Lin *et al.*, 2014], is composed of 80 semantic categories divided into 4 folds: $\{20^i\}_{i=0}^3$. Models are trained on 3 folds and tested on the remaining one in a cross-validation manner. We randomly sample 1,000 episodes for evaluation.

Evaluation Metrics. Following previous FSS approaches [Li *et al.*, 2021; Tian *et al.*, 2022], we adopt mean intersection-over-union (mIoU) and foreground-background IoU (FB-IoU) as our evaluation metrics, in which the mIoU metric is primarily used since it better reflects the overall performance of different categories.

Implementation Details. Two different backbone networks (*i.e.*, VGG16 [Simonyan and Zisserman, 2014] and ResNet50 [He *et al.*, 2016]) are adopted for a fair comparison with existing FSS methods. Following [Tian *et al.*, 2022], we pre-train these backbone networks on ILSVRC [Russakovsky *et al.*, 2015] and freeze them during training to boost the generalization capability. The setup of pre-processing technology is the same as [Tian *et al.*, 2022]. The SGD optimizer is utilized with a learning rate of 0.005 for 200 epochs on PASCAL-5ⁱ and 50 epochs on COCO-20ⁱ. All experiments are performed on NVIDIA RTX 2080Ti GPUs with the PyTorch framework. We report the average results of 5-runs with different random seeds to reduce the influence of selected support-query image pairs on performance.

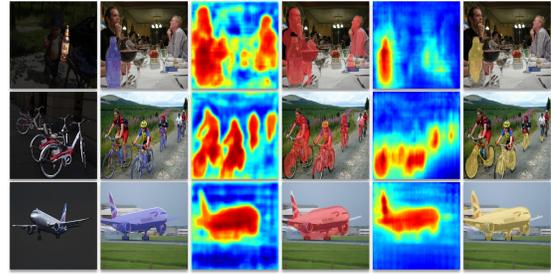


Figure 5: Semantic segmentation results on unseen categories under 1-shot setting. From left to right, each column represents support images, query images (blue), baseline activation maps (AMs), baseline results (red), our AMs, and our results (yellow), respectively.

5.2 Comparison with State-of-the-arts

We evaluate the proposed DCP framework with state-of-the-art methods on two standard FSS datasets. Table 1 presents the 1-shot and 5-shot results assessed under the mIoU metric on PASCAL-5ⁱ. It can be observed that our models outperform existing approaches by a sizeable margin, especially under the 5-shot setting, indicating that more appropriate and reliable information is supplied for query image segmentation. Table 2 summarizes the performance of different methods with ResNet50 backbone on COCO-20ⁱ, in which the proposed approach sets new state-of-the-arts with a small number of learnable parameters. Specifically, our DCP achieves 2.89% (1-shot) and 3.78% (5-shot) of mIoU improvements over the previous best approach (*i.e.*, PFENet), respectively.

For qualitative analysis, we first visualize the segmentation results of DCP and baseline approach, as illustrated in Figure 5. It can be found that the false positives caused by irrelevant objects are significantly reduced (see the first two rows). Meanwhile, more accurate segmentation boundaries are obtained, benefiting from the superiority of \mathbf{p}_β (see the last row). Note that for more details of the baseline approach, please refer to Table 3 in the ablation studies (§ 5.3).

5.3 Ablation Studies

The following ablation studies are performed with VGG16 backbone on PASCAL-5ⁱ unless otherwise stated.

Proxy & Prototype. As discussed in § 4.2, the derived proxies, along with foreground/background prototypes, tend to propagate pertinent information for the query branch in the form of activation maps. In Table 3, we evaluate the impact of each of them on segmentation performance under 1-shot settings. Overall, the following observations could be made: **(i)** By introducing the foreground activation map \mathbf{M}_f , the performance of baseline approach is greatly improved from 58.52% to 59.64%_{+1.12}, demonstrating the importance of foreground cues for guidance. Furthermore, we argue that such a simple yet powerful variant could serve as a new baseline approach for future FSS works. **(ii)** The background information provided by mask annotations needs to be used with *caution* since the performance might even deteriorate without PDS (see the second row). **(iii)** Our proposed proxies show a superior capability compared to the prototypes in facilitating query image segmentation (60.12% vs. 60.63% mIoU), which

Backbone	Method	1-shot					5-shot				
		P.-5 ⁰	P.-5 ¹	P.-5 ²	P.-5 ³	Mean	P.-5 ⁰	P.-5 ¹	P.-5 ²	P.-5 ³	Mean
VGG16	OSLSM [Shaban <i>et al.</i> , 2017]	33.60	55.30	40.90	33.50	40.80	35.90	58.10	42.70	39.10	43.90
	FWB [Nguyen and Todorovic, 2019]	47.04	59.64	52.61	48.27	51.90	50.87	62.86	56.48	50.09	55.08
	PANet [Wang <i>et al.</i> , 2019]	42.30	58.00	51.10	41.20	48.10	51.80	64.60	59.80	46.50	55.70
	PFENet [Tian <i>et al.</i> , 2022]	56.90	68.20	54.40	52.40	58.00	59.00	69.10	54.80	52.90	59.00
	MMNet [Wu <i>et al.</i> , 2021]	57.10	67.20	56.60	52.30	58.30	56.60	66.70	53.60	56.50	58.30
	DCP (ours)	59.67	68.67	63.78	53.11	61.31	64.25	70.66	67.38	61.08	65.84
ResNet50	CANet [Zhang <i>et al.</i> , 2019]	52.50	65.90	51.30	51.90	55.40	55.50	67.80	51.90	53.20	57.10
	SCL [†] [Zhang <i>et al.</i> , 2021]	63.00	70.00	56.50	57.70	61.80	64.50	70.90	57.30	58.70	62.90
	SAGNN [Xie <i>et al.</i> , 2021]	64.70	69.60	57.00	57.20	62.10	64.90	70.00	57.00	59.30	62.80
	MMNet [Wu <i>et al.</i> , 2021]	62.70	70.20	57.30	57.00	61.80	62.20	71.50	57.50	62.40	63.40
	CWT [Lu <i>et al.</i> , 2021]	56.30	62.00	59.90	47.20	56.40	61.30	68.50	68.50	56.60	63.70
	DCP (ours)	63.81	70.54	61.16	55.69	62.80	67.19	73.15	66.39	64.48	67.80

Table 1: Comparison with state-of-the-arts on PASCAL-5ⁱ in mIoU under 1-shot and 5-shot settings. ‘‘P.’’ means PASCAL. **RED/BLUE** represents the 1st/2nd best performance. Superscript ‘‘†’’ indicates that PFENet is served as the baseline.

Method	1-shot					5-shot					#Params.
	C.-20 ⁰	C.-20 ¹	C.-20 ²	C.-20 ³	Mean	C.-20 ⁰	C.-20 ¹	C.-20 ²	C.-20 ³	Mean	
FWB [†] [Nguyen and Todorovic, 2019]	16.98	17.98	20.96	28.85	21.19	19.13	21.46	23.93	30.08	23.65	43.0M
PFENet [†] [Tian <i>et al.</i> , 2022]	36.80	41.80	38.70	36.70	38.50	40.40	46.80	43.20	40.50	42.70	10.8M
MMNet [Wu <i>et al.</i> , 2021]	34.90	41.00	37.20	37.00	37.50	37.00	40.30	39.30	36.00	38.20	10.5M
CWT [Lu <i>et al.</i> , 2021]	32.20	36.00	31.60	31.60	32.90	40.10	43.80	39.00	42.40	41.30	47.3M
DCP (ours)	40.89	43.77	42.60	38.29	41.39	45.82	49.66	43.69	46.62	46.48	11.3M

Table 2: Comparison with state-of-the-arts on COCO-20ⁱ in mIoU. ‘‘C.’’ means COCO. The methods with superscript ‘‘†’’ use the ResNet101 backbone while other approaches use the ResNet50. ‘‘#Params.’’ indicates the number of *learnable* parameters.

Prototype		Proxy				PDS	mIoU	FB-IoU
M _F	M _B	M _α	M _β	M _γ	M _δ			
✓						58.52	72.46	
✓						59.64	73.04	
✓	✓					59.05	72.86	
✓	✓				✓	60.12	73.20	
✓		✓			✓	60.63	73.85	
✓	✓	✓	✓		✓	60.42	73.56	
✓	✓			✓	✓	60.76	74.11	
✓	✓	✓	✓	✓	✓	61.31	74.87	

Table 3: Ablation studies on each component.

Fusion Strategy	mIoU	FB-IoU	#Params.	FLOPs	Speed
Activation Map	61.31	74.87	0.26M	0.24G	16.1FPS
Tile & Concat.	60.50	74.56	0.52M	0.47G	14.9FPS

Table 4: Ablation studies on fusion strategies. ‘‘FLOPs’’ indicates the computational overhead. ‘‘Speed’’ denotes the evaluation of average frame-per-second (FPS) under the 1-shot setting.

is the so-called ‘‘beyond the prototype’’. (iv) There is a complementary relationship between the prototype and proxy, the combination of which could further boost performance. (v) Surprisingly, background-related proxies could be better integrated with prototypes (see rows 6-7). One possible reason for this phenomenon is that the interference of irrelevant objects may have a greater influence on the FSS model among many factors, while M_δ can effectively reduce false positives. **Parallel Decoder Structure.** After obtaining the activation maps (Eq. (6)), we first attempt to leverage a simple fusion scheme to guide the query branch, *i.e.*, concatenating $\mathbf{F}_{q,f}^{\text{tmp}}$ and $\{\mathbf{A}_l\}_{l \in P_1 \cup P_2}$ and feeding the refined features to a single decoder network. Unfortunately, such a fusion scheme does not offer performance gains, which we attribute to the *information confusion* between the expanded foreground pro-

otype $\zeta_{H \times W}(\mathbf{p}_f)$ and background-related activation maps (M_γ and M_δ). To this end, we devise a unique parallel decoder structure (PDS) where proxies/prototypes with similar attributes are integrated, yielding a superior result compared to the previous scheme (see rows 3-4 of Table 3).

Efficiency. We also evaluate the efficiency of different fusion strategies between derived proxies $\{\mathbf{p}_m\}_{m \in P_1}$ and query features \mathbf{F}_q , as presented in Table 4. Please note that the first row represents the previously discussed scheme (see Eq. (7)), while the second row denotes the scheme that concatenates query features with the expanded proxies, similar to Eq. (6). The experiment results indicate that the first fusion strategy exhibits remarkable segmentation performance with fewer parameters (0.26M) and faster speed (16+ FPS), demonstrating the effectiveness of activation maps for guidance.

6 Conclusion

We have proposed an efficient framework with the spirit of divide-and-conquer, which performs a self-reasoning scheme to derive support-induced proxies for tackling the typical challenges of FSS tasks. Qualitative and quantitative results show that the proposed proxies could provide appropriate and reliable information to facilitate query image segmentation, surpassing conventional prototype-based approaches by a considerable margin, *i.e.*, the so-called ‘‘beyond prototype’’.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grants 62136007 and U20B2065, in part by the Shaanxi Science Foundation for Distinguished Young Scholars under Grant 2021JC-16, and in part by the Fundamental Research Funds for the Central Universities.

References

- [Chen *et al.*, 2017] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(4):834–848, 2017.
- [Chen *et al.*, 2019] Zitian Chen, Yanwei Fu, Yu-Xiong Wang, Lin Ma, Wei Liu, and Martial Hebert. Image deformation meta-networks for one-shot learning. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, pages 8680–8689, 2019.
- [Deng *et al.*, 2009] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, pages 248–255, 2009.
- [Everingham *et al.*, 2010] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *Int. J. Comput. Vis.*, 88(2):303–338, 2010.
- [Finn *et al.*, 2017] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proc. IEEE Int. Conf. Mach. Learn.*, pages 1126–1135, 2017.
- [Hariharan *et al.*, 2011] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 991–998, 2011.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, pages 770–778, 2016.
- [Huang *et al.*, 2019] Zilong Huang, Xinggang Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. Cnet: Criss-cross attention for semantic segmentation. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 603–612, 2019.
- [Lang *et al.*, 2022] Chunbo Lang, Gong Cheng, Binfei Tu, and Junwei Han. Learning what not to segment: A new perspective on few-shot segmentation. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, 2022.
- [Li *et al.*, 2021] Gen Li, Varun Jampani, Laura Sevilla-Lara, Deqing Sun, Jonghyun Kim, and Joongkyu Kim. Adaptive prototype learning and allocation for few-shot segmentation. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, pages 8334–8343, 2021.
- [Lin *et al.*, 2014] 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 *Proc. Eur. Conf. Comput. Vis.*, pages 740–755, 2014.
- [Long *et al.*, 2015] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, pages 3431–3440, 2015.
- [Lu *et al.*, 2021] Zhihe Lu, Sen He, Xiatian Zhu, Li Zhang, Yi-Zhe Song, and Tao Xiang. Simpler is better: Few-shot semantic segmentation with classifier weight transformer. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 8741–8750, 2021.
- [Nguyen and Todorovic, 2019] Khoi Nguyen and Sinisa Todorovic. Feature weighting and boosting for few-shot segmentation. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 622–631, 2019.
- [Ren *et al.*, 2015] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Proc. Conf. Adv. Neural Inform. Process. Syst.*, volume 28, 2015.
- [Russakovsky *et al.*, 2015] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.*, 115:211–252, 2015.
- [Shaban *et al.*, 2017] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots. One-shot learning for semantic segmentation. In *Proc. Br. Mach. Vis. Conf.*, 2017.
- [Siam *et al.*, 2019] Mennatullah Siam, Boris N Oreshkin, and Martin Jagersand. Amp: Adaptive masked proxies for few-shot segmentation. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 5249–5258, 2019.
- [Simonyan and Zisserman, 2014] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [Snell *et al.*, 2017] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Proc. Conf. Adv. Neural Inform. Process. Syst.*, volume 30, 2017.
- [Tian *et al.*, 2022] Zhuotao Tian, Hengshuang Zhao, Michelle Shu, Zhicheng Yang, Ruiyu Li, and Jiaya Jia. Prior guided feature enrichment network for few-shot segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(2):1050–1065, 2022.
- [Vinyals *et al.*, 2016] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, koray kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In *Proc. Conf. Adv. Neural Inform. Process. Syst.*, volume 29, 2016.
- [Wang *et al.*, 2019] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 9197–9206, 2019.
- [Wu *et al.*, 2021] Zhonghua Wu, Xiangxi Shi, Guosheng Lin, and Jianfei Cai. Learning meta-class memory for few-shot semantic segmentation. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 517–526, 2021.
- [Xie *et al.*, 2021] Guo-Sen Xie, Jie Liu, Huan Xiong, and Ling Shao. Scale-aware graph neural network for few-shot semantic segmentation. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, pages 5475–5484, 2021.
- [Zhang *et al.*, 2019] 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 *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, pages 5217–5226, 2019.
- [Zhang *et al.*, 2020] Xiaolin Zhang, Yunchao Wei, Yi Yang, and Thomas S Huang. Sg-one: Similarity guidance network for one-shot semantic segmentation. *IEEE Trans. Cybern.*, 50(9):3855–3865, 2020.
- [Zhang *et al.*, 2021] Bingfeng Zhang, Jimin Xiao, and Terry Qin. Self-guided and cross-guided learning for few-shot segmentation. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, pages 8312–8321, 2021.
- [Zhao *et al.*, 2017] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proc. IEEE Int. Conf. Comput. Vis. Pattern Recognit.*, pages 2881–2890, 2017.