

Hierarchical Dense Correlation Distillation for Few-Shot Segmentation

Bohao Peng¹, Zhuotao Tian^{4*}, Xiaoyang Wu², Chengyao Wang¹, Shu Liu⁴, Jingyong Su³, Jiaya Jia^{1,4}

¹The Chinese University of Hong Kong ²The University of Hong Kong

³Harbin Institute of Technology, Shenzhen ⁴SmartMore

Abstract

Few-shot semantic segmentation (FSS) aims to form class-agnostic models segmenting unseen classes with only a handful of annotations. Previous methods limited to the semantic feature and prototype representation suffer from coarse segmentation granularity and train-set overfitting. In this work, we design Hierarchically Decoupled Matching Network (HDMNet) mining pixel-level support correlation based on the transformer architecture. The self-attention modules are used to assist in establishing hierarchical dense features, as a means to accomplish the cascade matching between query and support features. Moreover, we propose a matching module to reduce train-set overfitting and introduce correlation distillation leveraging semantic correspondence from coarse resolution to boost fine-grained segmentation. Our method performs decently in experiments. We achieve 50.0% mIoU on COCO-20ⁱ dataset one-shot setting and 56.0% on five-shot segmentation, respectively. The code is available on the project website¹.

1. Introduction

Semantic segmentation tasks [2, 3, 22, 52] have made tremendous progress in recent years, benefiting from the rapid development of deep learning [13, 32]. However, most existing deep networks are not scalable to previously unseen classes and rely on annotated datasets to achieve satisfying performance. Data collection and annotation cost much time and resources, especially for dense prediction tasks.

Few-shot learning [34, 39, 43] has been introduced into semantic segmentation [5, 38] to build class-agnostic models quickly adapting to novel classes. Typically, few-shot segmentation (FSS) divides the input into the query and support sets [5, 46, 48, 51] following the episode paradigm [41]. It segments the query targets conditioned on

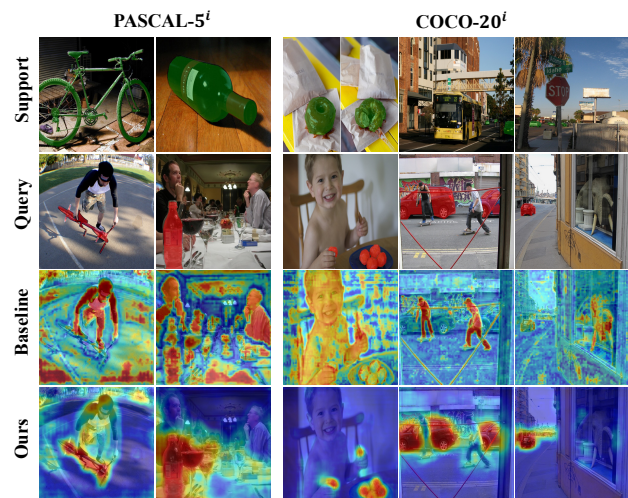


Figure 1. Activation maps of the correlation values on both PASCAL-5ⁱ [29] and COCO-20ⁱ [26]. The baseline is prone to give high activation values to the categories sufficiently witnessed during training, such as the “People” class, even with other support annotations. Then we convert it to the hierarchically decoupled matching structure and adopt correlation map distillation to mine inner-class correlation.

the semantic clues from the support annotations with meta-learning [34, 39] or feature matching [25, 41, 49].

Previous few-shot learning methods may still suffer from coarse segmentation granularity and train-set overfitting [38] issues. As shown in Fig. 1, “people” is the base class that has been sufficiently witnessed during training. But the model is still prone to yield high activation to “people” instead of more related novel classes with the support samples, producing inferior results. This issue stems from framework design, as illustrated in Fig. 2. Concretely, prototype-based [38, 42] and adaptive-classifier methods [1, 23] aim at distinguishing different categories with global class-wise characteristics. It is challenging to compute the correspondence of different components between query and support objects for the dense prediction tasks. In contrast, matching-based methods [49] mine pixel-level correlation but may heavily rely on class-specific fea-

*Corresponding Author

¹<https://github.com/Pbihao/HDMNet>

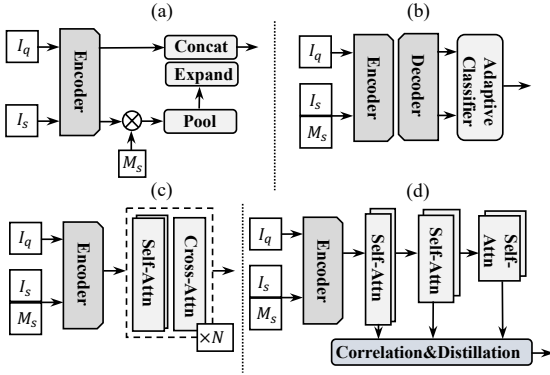


Figure 2. Illustration of different few-shot segmentation frameworks. (a) Prototype-based method. (b) Adaptive-classifier method. (c) Feature matching with transformer architecture. (d) Our Hierarchically Decoupled Matching Network (HDMNet) with correlation map distillation.

tures and cause overfitting and weak generalization.

To address these issues, we propose Hierarchically Decoupled Matching Network (HDMNet) with correlation map distillation for better mining pixel-level support correspondences. HDMNet extends transformer architecture [6, 40, 44] to construct the feature pyramid and performs dense matching. Previous transformer-based methods [35, 49] adopt the self-attention layer to parse features and then feed query and support features to the cross-attention layer for pattern matching, as illustrated in Fig. 2(c). This process stacks the self- and cross-attention layers multiple times, mixes separated embedding features, and accidentally causes unnecessary information interference.

In this paper, we decouple the feature parsing and matching process in a hierarchical paradigm and design a new matching module based on correlation and distillation. This correlation mechanism calculates pixel-level correspondence without directly relying on the semantic-specific features, alleviating the train-set overfitting problem. Further, we introduce correlation map distillation [14, 50] that encourages the shallow layers to approximate the semantic correlation of deeper layers to make the former more aware of the context for high-quality prediction.

Our contribution is the following. 1) We extend the transformer to hierarchical parsing and feature matching for few-shot semantic segmentation, with a new matching module reducing overfitting. 2) We propose correlation map distillation leveraging soft correspondence under multi-level and multi-scale structure. 3) We achieve new state-of-the-art results on standard benchmark of COCO-20ⁱ and PASCAL-5ⁱ without compromising efficiency.

2. Related Work

Few-shot semantic segmentation. Few-Shot Semantic Segmentation (FSS) [5, 9, 10, 38] predicts dense masks for

novel classes with only a few annotations. Previous approaches following metric learning [5, 36, 38, 42] can be divided into prototype- and matching-based methods. Motivated by PrototypicalNet [34] for few-shot learning, the prevalent FSS models [5, 8, 24, 37, 42] utilize prototypes for specific-class representation. Recent work [18, 47, 49] points out that a single prototype has a limitation to cover all regions of an object, especially for pixel-wise dense segmentation tasks. To remedy this problem, methods of [18, 47] use EM and cluster algorithms to generate multiple prototypes for different parts of the objects. Compared with prototype-based methods, matching-based ones [23, 25, 41, 49] are based on intuition to mine dense correspondence between the query images and support annotations. They utilize pixel-level features and supplement more detailed support context.

Transformer. Transformer was first introduced in Natural Language Processing (NLP) [4, 40]. In computer vision, ViT [6] treats an image as a patch sequence and demonstrates that pure transformer architecture can achieve state-of-the-art image classification. Recent work explores combining few-shot semantic segmentation and transformer architecture [19, 35]. In [23], the classifier weight transformer adapts the classifier’s weights to address the intra-class variation issue. CyCTR [49] is a cycle-consistent transformer by generating query and key sequences from the query and support set, respectively. Transformer architecture helps FSS transcend the limitation of semantic-level prototypes and leverage pixel-wise alignment. Previous transformer-based methods are still difficult to handle noise interference and over-fitting. We, in this paper, propose a new transformer structure decoupling the downsampling and matching processes and design the matching module constructed on correlation mechanism and distillation.

Knowledge distillation. Knowledge distillation (KD) [14] was widely used in model compression. Large models typically have higher knowledge capacity. In contrast, small models have fewer parameters, better efficiency, and lower cost. KD attempts to transfer learned knowledge from the large model (a.k.a. the teacher) to another light model (a.k.a. the student) with tolerable loss in performance. Method of [50] processes a self-distillation framework to distill knowledge within itself to improve model accuracy. Self-distillation divides the model into multiple sections and transfers knowledge from deeper portions to shallow ones. Knowledge distillation is also used for extracting semantic information and mining inner correlation. STEGO [11] shows that even unsupervised deep network features have correlation patterns consistent with true labels. STEGO framework applies feature correlation map distillation to excavate the intrinsic semantic correlation at pixel level. Motivated by unsupervised semantic segmentation [11], we design the correlation

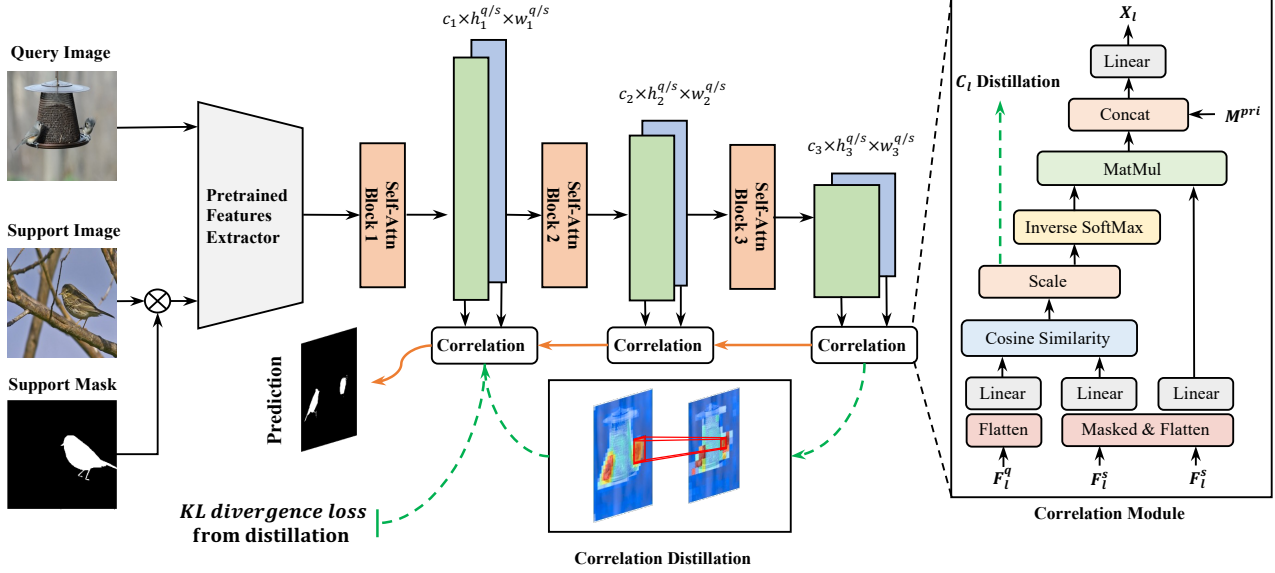


Figure 3. Schematic overview of the proposed few-shot semantic segmentation model. $h_l^{q/s} = \frac{H^{q/s}}{2^{l+2}}$, $w_l^{q/s} = \frac{W^{q/s}}{2^{l+2}}$ indicate the height and width of the l -th stage features, and $H^{q/s}$ and $W^{q/s}$ are the height and width of the input query and support images. c_l represents the channels and $c_{l+1} > c_l$.

mechanism for class-agnostic feature matching and propose multi-level multi-scale correlation map distillation to transfer relation between the query and support set from deep sections to the shallow ones.

3. Task Definition and Our Method

Few-shot segmentation is to train segmentation for novel objects with only a few annotated support images. In definition, the model is trained on D_{train} and is evaluated on D_{test} . Suppose the category sets in D_{train} and D_{test} are C_{train} and C_{test} respectively. There is no intersection between the training and testing sets, i.e., $C_{train} \cap C_{test} = \emptyset$. Following previous work [30, 38, 42, 48], episodes are applied to both train set D_{train} and test set D_{test} .

Each episode is composed of a query set $Q = \{(I^q, M^q)\}$ and a support set $S = \{(I_i^s, M_i^s)\}_{i=1}^K$ with the same class c , where $I^q, I^s \in \mathbb{R}^{H \times W \times 3}$ represent the RGB images and $M^q, M^s \in \mathbb{R}^{H \times W}$ denote their binary masks. Both the query masks M^q and the support masks M^s are used during the training process, while only the support masks M^s are accessible in testing. Since the model parameters are fixed and require no optimization for novel categories during testing, the model is trained to leverage the semantic clues provided by the support set to locate the regions of interest on the query images.

3.1. Hierarchically Decoupled Matching Network

Given the query set $Q = \{(I^q, M^q)\}$ and the support set $S = \{(I_i^s, M_i^s)\}_{i=1}^K$, HDMNet adopts a parameter-fixed encoder to extract rich features of the query and support im-

ages, following [38, 48]. The difference is on the design of new decoder to yield predictions on the query images by decently leveraging pixel-level feature matching between the query and support sets.

Overview of the pipeline is shown in Fig. 3. In the following, we start by putting forward the basic structure of HDMNet in Sec. 3.2, followed by the introduction regarding the proposed correlation calculation strategy in Sec. 3.3. Then, the inter-stage correlation map distillation is presented in Sec. 3.4. Finally, in Sec. 3.5, we instantiate the way to extend to the K-shot setting.

3.2. Overview of the Architecture

Motivation. In previous matching-based methods with transformer architecture, the self-attention and cross-attention layers are interleaved for multiple times for feature parsing and pattern matching respectively [35, 49] as shown in Fig. 4. We note that the cross-attention layers accomplish mutual message exchange between the query and support features. Objects in the background of the query sample may also correlate with the target in the support sample. Thus, they can be enriched with support information. With this finding, the necessary support information may be accumulated to the distracters via multiple stacked cross-attention and self-attention layers, making the decoder harder to distinguish among them.

To ensure the purity of the sequential features and consistency of pattern matching, we propose a new hierarchically matching structure decoupling the down-sampling and matching processes, where only independent self-layers are




Figure 4. Feature Matching with directly stacking cross-attention layers and self-attention layers. It’s intuitive to notice that the cross-attention layer mixes the query and support features, destroying the purity of parsing and matching consistency.

adopted to build hierarchical features.

Decoupled downsampling and matching. First, the query and support features extracted from the backbone are independently sent to sequential transformer blocks with only the self-attention layers to fully exploit self-correlation within the support and query features. We note that the down-sampling layers are inserted between blocks to establish a hierarchical structure that may assist in mining the inter-scale correlations.

Then, the intermediate feature maps of L stages are collected, i.e., $\{\mathbf{F}_l^q\}_{l=1}^L$ and $\{\mathbf{F}_l^s\}_{l=1}^L$. Assume $\{\mathbf{F}_l^q\}$ and $\{\mathbf{F}_l^s\}$ have the same spatial size $[c_l \times h_l^{q/s} \times w_l^{q/s}]$ for simplicity’s sake.

$$h_l^{q/s} = \frac{H^{q/s}}{2^{l+2}}, \quad w_l^{q/s} = \frac{W^{q/s}}{2^{l+2}},$$

l is the stage index, and c_l denotes the feature channel number. Finally, $\{\mathbf{F}_l^q\}_{l=1}^L$ and $\{\mathbf{F}_l^s\}_{l=1}^L$ are used to yield correlations $\{\mathbf{C}_l \in \mathbb{R}^{h_l^q w_l^q \times h_l^s w_l^s}\}_{l=1}^L$ and enriched query features $\{\mathbf{X}_l \in \mathbb{R}^{c_l \times h_l^q \times w_l^q}\}_{l=1}^L$. Detailed formulations are elaborated later in Eqs. (4) and (6) in Sec. 3.3.

Coarse-grained to fine-grained decoder. HDMNet incorporates a simple decoder to predict the final mask for the query image with the hierarchically enriched features $\{\mathbf{X}_l \in \mathbb{R}^{c_l \times h_l^q \times w_l^q}\}_{l=1}^L$ in a coarse-to-fine manner. Specifically, the coarse-grained features \mathbf{X}'_{l+1} are scaled up to have the same spatial size as the fine-grained one \mathbf{X}'_l . Then an MLP layer is adopted to fuse them with a residual connection, written as

$$\mathbf{X}'_l = \text{ReLU}(\text{MLP}(\mathbf{X}_l + \zeta_l(\mathbf{X}'_{l+1}))) + \zeta_l(\mathbf{X}'_{l+1}), \quad (1)$$

where l indicates the hierarchical stage, and $\zeta_l : \mathbb{R}^{H \times W} \mapsto \mathbb{R}^{h_l \times w_l}$ denotes the bilinear-interpolation resize function fitting the input size to that of the output. Finally, we apply a convolution layer with 1×1 kernel size to \mathbf{X}'_l followed by a bilinear up-sampling layer to predict the query mask $\mathbf{M}^{out} \in \mathbb{R}^{H \times W}$.

3.3. Matching Module

Motivation. Previous matching-based methods [23, 49] directly adopt the cross-attention as the matching module

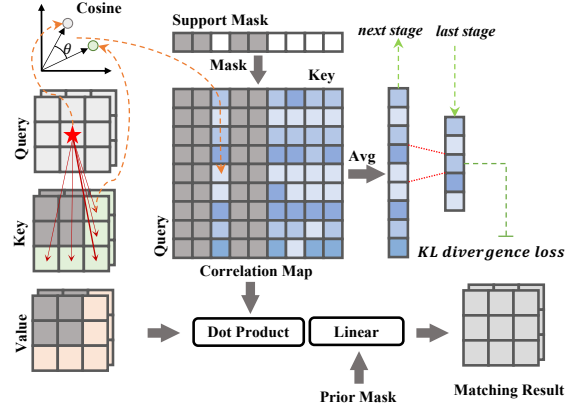


Figure 5. Illustration of our proposed matching module based on correlation mechanism and distillation.

by generating the query and key features from the query and support sets respectively. However, we observe that it leads to overfitting and weak generalization. This issue could be attributed to the fact that the models are more likely to rely on class-specific features to optimize the training objectives. To alleviate this issue, we propose a matching module illustrated in Fig. 5.

Attention. Following the general form [40, 44], the critical element of the transformer block is the dot-product attention layer, formulated as

$$\text{Attn}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}}\right)\mathbf{V}, \quad (2)$$

where $[\mathbf{Q}; \mathbf{K}; \mathbf{V}] = [\mathbf{W}^q \mathbf{F}^q; \mathbf{W}^k \mathbf{F}^s; \mathbf{W}^v \mathbf{F}^s]$, in which \mathbf{F}^q and \mathbf{F}^s denote the query and support features respectively, $\mathbf{W}^q, \mathbf{W}^k, \mathbf{W}^v \in \mathbb{R}^{d \times d}$ are the learnable parameters, d is the hidden dimension.

The cross-attention layer takes essential support information from \mathbf{V} , conditioned on the query-support correlation between \mathbf{Q} and \mathbf{K} . When $\mathbf{F}^q = \mathbf{F}^s$, it functions as a self-attention layer for relating different positions within either the support or query input features.

Our correlation mechanism. Our designed matching module based on the correlation mechanism retrieves the most relevant regions with high cosine similarity and fuses the high-level prior mask generated as that in [38]. Given the query features \mathbf{F}^q and the support features \mathbf{F}^s , we first transform the input features by

$$\begin{aligned} \hat{\mathbf{F}}^q &= \varphi(\mathbf{F}^q), \\ \hat{\mathbf{F}}^s &= \varphi(\mathbf{F}^s \odot \mathbf{M}^s), \end{aligned} \quad (3)$$

where \odot is Hadamard product, $\varphi : \mathbb{R}^{c \times h \times w} \mapsto \mathbb{R}^{hw \times c}$ refers to the reshape function, and \mathbf{M}^s denotes the support mask. To mitigate the risk of overfitting the category-

specific information brought by the feature norms, we measure the cosine similarities of the inner product angle, instead of performing dot product, to calculate the correlation map as $C \in \mathbb{R}^{h^q w^q \times h^s w^s}$ as

$$C = \frac{\langle \mathbf{W}^q \hat{\mathbf{F}}^q, \mathbf{W}^k \hat{\mathbf{F}}^s \rangle}{\left\| \mathbf{W}^q \hat{\mathbf{F}}^q \right\| \left\| \mathbf{W}^k \hat{\mathbf{F}}^s \right\| t}, \quad (4)$$

where $\mathbf{W}^q, \mathbf{W}^k \in \mathbb{R}^{c \times c}$ denote the learnable parameters, $\|\cdot\|$ indicates L^2 norm, and t is a hyperparameter to control the distribution range, empirically set to 0.1 in all experiments. Inspired by [28, 33], we propose the inverse softmax layer that normalizes the correlation matrix along the query axis since we only retrieve the interested region of the query set as

$$\hat{C}(i, j) = \frac{\exp(C(i, j))}{\sum_{k=1}^{h_l^q w_l^q} \exp(C(k, j))}. \quad (5)$$

Finally, we introduce the prior mask $M^{pri} \in \mathbb{R}^{h^q \times w^q}$ calculated the same as [38] by concatenating it with correspondence scores along the channel dimension to generate the matching results of

$$\mathbf{X} = \mathbf{W}^o([\psi(\hat{C}(\mathbf{W}^v \hat{\mathbf{F}}^s)), M^{pri}], \quad (6)$$

where $\mathbf{W}^v \in \mathbb{R}^{c \times c}$, $\mathbf{W}^o \in \mathbb{R}^{c \times (c+1)}$ denote the learnable parameters, $\hat{\mathbf{F}}^s \in \mathbb{R}^{h^s w^s \times c}$, $\mathbf{X} \in \mathbb{R}^{c \times h^q \times w^q}$ are flattened support features and matching output, and $\psi: \mathbb{R}^{h^q w^q \times c} \mapsto \mathbb{R}^{c \times h^q \times w^q}$ is the reshape function.

3.4. Correlation Map Distillation

Motivation. The query and support features are constructed in a hierarchical structure, used to generate multi-level multi-scale dense correlation map. High-level correspondence typically has more informative semantic cues. It only gives rough location limited by the coarse resolution, while the low stages contain more locally detailed features. We propose correlation map distillation to encourage the correlation maps of earlier stages to retain the fine-grained segmentation quality without deprecating the contextual hints.

In addition, compared to the one-hot labels in ground-truth annotation, the soft targets adopted during distillation can reveal the distributions among all categories, providing extra semantic cues regarding the inter-class relation. It facilitates the shallower ones to capture broader context information. On the contrary, the ground-truth annotation may not provide such information and may even let the model over-fit noises contained in the one-hot labels.

Distillation formulation. Eq.(4) calculates the correlation maps $\{C_l \in \mathbb{R}^{h_l^q w_l^q \times h_l^s w_l^s}\}_{l=1}^L$ for the query and support features. We reorganize C_l with mean average and filter the irrelevant information by the support mask M^s as

$$C'_l(i) = \frac{\sum_{j=1}^{h_l^s w_l^s} C_l(i, j) \cdot [\varphi \circ \zeta_l(M^s)(j) > 0]}{\sum_{j=1}^{h_l^s w_l^s} [\varphi \circ \zeta_l(M^s)(j) > 0]}, \quad (7)$$

where l indicates the stage, $C'_l \in \mathbb{R}^{h_l^q w_l^q}$, and ζ_l is the resize function. Given flattened correlation maps, we apply a softmax layer to perform spatial normalization among all positions as

$$\hat{C}'_l(i) = \frac{\exp(C'_l(i)/T)}{\sum_{j=1}^{h_l^q w_l^q} \exp(C'_l(j)/T)}, \quad (8)$$

where l indicates the stage and T denotes the temperature of distillation [14] set to 1. Moreover, the results regarding the temperature T are shown in the supplementary file.

Then the KL (Kullback-Leibler) divergence loss is used as supervision from the teacher to student with their softmax output. The correlation maps of adjacent stages act as the teacher and student respectively, formulated as

$$\begin{aligned} \mathcal{L}_{KL} &= \sum_{x \in \mathcal{X}} \phi_t(x) \log\left(\frac{\phi_t(x)}{\phi_s(x)}\right) \\ &= \sum_{i=1}^{h_l^q w_l^q} \zeta_l(\hat{C}_{l+1})(i) \cdot \log\left(\frac{\zeta_l(\hat{C}_{l+1})(i)}{\hat{C}'_l(i)}\right), \end{aligned} \quad (9)$$

where l indicates the stage, ϕ_t is the teacher model while ϕ_s is the student model, and $\zeta_l: \mathbb{R}^{h_{l+1}^q w_{l+1}^q} \mapsto \mathbb{R}^{h_l^q w_l^q}$ represents resizing. In particular, for the last correlation map without successor, we directly utilize the ground-truth as its teacher.

3.5. Extension to K -shot Setting

In extension to K -shot ($K > 1$) setting, K support images with their annotated masks $S = \{(I_k^s, M_k^s)\}_{k=1}^K$ and the query set $\{(I^q, M^q)\}$ are given. HDMNet can be quickly and easily extended to the new setting based on the matching mechanism.

As indicated in Secs. 3.3 and 3.4, the feature matching and distillation processes are independent of the specific size of the support features, benefiting from the correlation mechanism and preprocessing of the correlation map. To prevent information loss and ensure consistency under different settings, we concatenate the support features $\mathbf{F}_l^s = \text{CONCAT}([\mathbf{F}_{l,1}^s, \mathbf{F}_{l,2}^s, \dots, \mathbf{F}_{l,K}^s])$ along the channel dimension as well as their corresponding masks directly. The decoder processing remains the same as that in the one-shot setting.

4. Experiments

Datasets. Following the setting of [38], we use two benchmark few-shot segmentation datasets, i.e.,

Backbone	Methods	1-shot					5-shot				
		Fold-0	Fold-1	Fold-2	Fold-3	Mean	Fold-0	Fold-1	Fold-2	Fold-3	Mean
VGG-16	PANet [42]	-	-	-	-	20.9	-	-	-	-	29.7
	FWB [26]	18.4	16.7	19.6	25.4	20.0	20.9	19.2	21.9	28.4	22.6
	PRNet [21]	27.5	33.0	26.7	29.0	29.1	31.2	36.5	31.5	32.0	32.8
	PFENet [38]	35.4	38.1	36.8	34.7	36.3	38.2	42.5	41.8	38.9	40.4
	BAM [†] [16]	36.4	47.1	43.3	41.7	42.1	42.9	51.4	48.3	46.6	47.3
	HDMNet (Ours)	40.7	50.6	48.2	44.0	45.9	47.0	56.5	54.1	51.9	52.4
ResNet-50	DCP [17]	40.9	43.8	42.6	38.3	41.4	45.8	49.7	43.7	46.6	46.5
	PFENet [38]	36.5	38.6	34.5	33.8	35.8	36.5	43.3	37.8	38.4	39.0
	RPMs [45]	29.5	36.8	29.0	27.0	30.6	33.8	42.0	33.0	33.3	35.5
	RePRI [1]	32.0	38.7	32.7	33.1	34.1	39.3	45.4	39.7	41.8	41.6
	HSNet [25]	36.3	43.1	38.7	38.7	39.2	43.3	51.3	48.2	45.0	46.9
	CAPL [37]	37.3	43.3	40.2	38.4	39.8	43.1	53.8	48.7	47.4	48.3
	CWT [23]	30.3	36.6	30.5	32.2	32.4	38.5	46.7	39.4	43.2	42.0
	DGPNet [15]	43.6	47.8	44.5	44.2	45.0	54.7	59.1	56.8	54.4	56.2
	CyCTR [49]	38.9	43.0	39.6	39.8	40.3	41.1	48.9	45.2	47.0	45.6
	BAM [†] [16]	39.4	49.9	46.2	45.2	45.2	43.2	53.4	49.4	48.1	48.5
HDMNet (Ours)	43.8	55.3	51.6	49.4	50.0	50.6	61.6	55.7	56.0	56.0	

Table 1. Few-shot semantic segmentation performance comparison on COCO-20ⁱ [26] using mIoU (%) evaluation metric. Numbers in bold indicate the best performance. †: Reproduced following the official configuration with 10,000 test episodes.

PASCAL-5ⁱ [29] and COCO-20ⁱ [26], to evaluate HDMNet. PASCAL-5ⁱ is built from PASCAL VOC 2012 [7] with additional annotations from SDS [12]. It consists of 20 classes. COCO-20ⁱ is generated from MSCOCO [20] and contains 80 categories.

For each selected dataset, cross-validation is conducted by dividing all classes into 4 folds evenly. We use the same split class list as that of [26, 29] on PASCAL-5ⁱ and COCO-20ⁱ respectively. Specifically, three folds serve as training data, while the remaining one is used for testing. For ensuring the performance stability and fairness for comparison, we follow [38] randomly sample 1,000 and 10,000 query/support pairs for PASCAL-5ⁱ and COCO-20ⁱ testing. More analysis and discussion about datasets and test episodes are shown in supplementary materials.

Metrics. We adopt the mean intersection over union (mIoU) as the main evaluation metric and foreground-background IoU (FB-IoU) as the supplement. We denote $mIoU = 1/C \sum_{i=1}^C IoU_i$, where C is the number of classes in each fold, and IoU_i indicates intersection-over-union for class i . $FB-IoU = 1/2(IoU_F + IoU_B)$, where IoU_F and IoU_B represent the foreground and background IoU values, ignoring the class difference and computing the average.

4.1. Implementation Details

HDMNet is built upon the Pytorch [27] framework. All models are trained on 4 NVIDIA GeForce RTX 3090 GPUs and tested on a single GPU. The training augmentations of PASCAL-5ⁱ dataset and COCO-20ⁱ dataset follow that of [38] for fair comparisons, including random crop, scale, rotate, blur and flip.

HDMNet is trained in an episode fashion for 200 and 50 epochs on COCO-20ⁱ and PASCAL-5ⁱ, and the batch sizes are set to 6 and 4 respectively. During training, AdmW optimizer is adopted the same as [25, 49], and the learning rate is set to 0.0001. In addition, the weight decay is 0.01, and the “poly” strategy is used to adjust the learning rate.

We use ResNet-50 [13] and VGG-16 [32] as the encoder to extract features with freezing parameters to verify the effectiveness of the proposed method on different backbones. PSPNet [52] serves as the base learner in all experiments. We apply the PPM module [52], which was widely used in previous semantic segmentation methods, to provide multi-resolution context for feature enrichment after the 4th block of ResNet-50 or VGG-16 to generate the prior mask [38]. Similar to that of [38], we concatenate the prior mask and match features with a 1×1 kernel size convolution layer leveraging the high-level semantic information to boost performance. During testing, predictions are resized back to the original sizes of the input images, keeping the ground-truth labels intact [38].

4.2. Comparison with State-of-the-Art Methods

In Tables 1 and 2, we report comparison of our proposed HDMNet with other state-of-the-art few-shot semantic segmentation approaches in recent years on COCO-20ⁱ [26] and PASCAL-5ⁱ [29] datasets. The mIoU (%) is used as the evaluation metric. To verify the generality, we build our baseline with VGG-16 [32] and Resnet-50 [13]. Both of them gain significant improvement from our method, and our method achieves new state-of-the-art performance on both COCO-20ⁱ and PASCAL-5ⁱ datasets. Especially on COCO-20ⁱ dataset, our model outperforms

Backbone	Methods	1-shot					5-shot				
		Fold-0	Fold-1	Fold-2	Fold-3	Mean	Fold-0	Fold-1	Fold-2	Fold-3	Mean
VGG-16	PANet [42]	42.3	58.0	51.1	41.2	48.1	51.8	64.6	59.8	46.5	55.7
	PFENet [38]	56.9	68.2	54.4	52.4	58.0	59.0	69.1	54.8	52.9	59.0
	HSNet [25]	59.6	65.7	59.6	54.0	59.7	64.9	69.0	64.1	58.6	64.1
	HDMNet (Ours)	64.8	71.4	67.7	56.4	65.1	68.1	73.1	71.8	64.0	69.3
ResNet-50	HSNet [25]	64.3	70.7	60.3	60.5	64.0	70.3	73.2	67.4	67.1	69.5
	PFENet [38]	61.7	69.5	55.4	56.3	60.8	63.1	70.7	55.8	57.9	61.9
	CyCTR [49]	65.7	71.0	59.5	59.7	64.0	69.3	73.5	63.8	63.5	67.5
	SSP [8]	60.5	67.8	66.4	51.0	61.4	67.5	72.3	75.2	62.1	69.3
	DCAMA [31]	67.5	72.3	59.6	59.0	64.6	70.5	73.9	63.7	65.8	68.5
	SD-AANet [53]	60.9	70.8	58.4	57.3	61.9	65.5	71.6	62.5	62.3	65.5
	BAM [16]	69.0	73.6	67.6	61.1	67.8	70.6	75.1	70.8	67.2	70.9
	HDMNet (Ours)	71.0	75.4	68.9	62.1	69.4	71.3	76.2	71.3	68.5	71.8

Table 2. Performance on PASCAL-5ⁱ [29] using the mIoU (%) evaluation metric. Results in bold denote the best performance.

Backbone	Methods	FB-IoU (%)		#learnable params
		1-shot	5-shot	
ResNet-50	ASGNet [18]	60.4	67.0	10.4M
	HSNet [25]	68.2	70.7	2.6M
	BAM [16]	71.1	73.3	4.9M
	HDMNet (Ours)	72.2	77.7	4.2M

Table 3. Comparison of results on COCO-20ⁱ in terms of FB-IoU and the number of learnable parameters.

Ens.	HDM	Corr.	Distill	mIoU (%)	Δ
				44.7	0.0
✓				45.8	+1.1
✓	✓			47.9	+3.2
✓	✓	✓		48.3	+3.6
✓	✓	✓	✓	50.0	+5.3

Table 4. Ablation studies for different components and architecture design in HDMNet.

the prior arts by a significant margin, achieving 4.8% (1-shot) and 7.5% (5-shot) of mIoU improvements over the SOTA with ResNet-50 backbone.

COCO-20ⁱ dataset contains 80 categories compared to 20 classes in PASCAL-5ⁱ dataset and has a particularly larger image capacity. The superiority on COCO-20ⁱ dataset proves that our method is with higher generality and can better adapt to novel categories in more complex scenes. Table 2 shows that our HDMNet also achieves 69.4% (1-shot) and 71.8% (5-shot) of mIoU on Pascal with ResNet-50 backbone, surpassing previous state-of-the-art. Table 3 gives comparison in terms of FB-IoU and model parameter number on COCO-20ⁱ for 1-shot and 5-shot segmentation. HDMNet achieves the best performance without compromising efficiency.

4.3. Ablation Study

We report the ablation study results in this section to investigate the effectiveness of each component and our design choice. All ablation experiments are conducted un-

Matching	mIoU(%)	Loss	mIoU(%)
CA	48.0	w/o	47.9
Cos	49.2	CE	48.4
Cos+SM	49.5	CE+KD	49.1
Cos+Inv-SM	50.0	KL+KD	50.0

Table 5. Ablation study on correlation mechanism. CA: Cross-Attention. Cos: Cosine similarity. Inv-SM: Inversed Softmax.

Table 6. Ablation study on the different loss functions. KD indicates using the adjacent layers' soft predictions.

der COCO-20ⁱ 1-shot setting with ResNet-50 backbone if not otherwise specified.

Component-wise ablation. Table 4 shows ablation results regarding the effectiveness of different components and architecture design, where the mIoU results are averaged over four splits.

The first line is the baseline result. We build our baseline following [38,49] and utilize ResNet-50 [13] as the encoder to extract image features and generate the prior mask [38]. Extracted features from Block-2 and Block-3 of the backbone are fused and fed to the next step with the prior mask. The baseline stacks the self-attention and cross-attention modules three times and applies the cross-attention module for feature matching [49].

For a fair comparison, we use the ensemble module (Ens.) following BAM to filter the categories appearing in the training process. The test result is incrementally improved. Then we convert the framework to our proposed structure (HDM) described in Sec. 3.2, resulting in even better results, with further mIoU increase of 2.1% compared to the one with Ens. only. It indicates that decoupling feature down-sampling and matching reduce noise interference and boost performance. Replacing the attention function with the proposed correlation mechanism (Corr.) as mentioned in Sec. 3.3 within the matching module continuously improves the performance by 0.4% mIoU.

Decoder	mIoU(%)	params(M)	time(ms)	FLOPs(G)
CyCTR [49]	40.3	5.6	54.3	96.7
HSNet [25]	39.2	2.6	25.5	20.6
BAM [16]	45.2	4.1	7.4	26.0
Ours- S_1	47.1	1.3	15.1	8.8
Ours- S_2	48.8	2.1	21.1	10.2
Ours- S_3	50.0	2.8	27.4	10.6
Ours- S_4	48.6	3.6	38.0	10.6

Table 7. Comparison of decoders from different methods in terms of accuracy, efficiency, and model size. S_i indicates constructing our decoder with i matching stages.




Figure 6. Ablation study on correlation map distillation for training and validation on COCO-20ⁱ dataset in terms of mIoU.

Finally, we use the correlation map distillation (Distill) described in Sec. 3.4 to supervise the matching process and leverage the correlation information from different stages, bringing another 1.7% mIoU improvement, which validates the effectiveness of our proposed distillation strategy.

Correlation mechanism ablation. Table 5 compares the alternative choices of the correlation mechanism. We first build HDMNet with the original cross-attention layer (CA) [40] as the matching module and apply the distillation function to the attention map, *i.e.*, dot product between the query and key sequences. Then we replace the matching module with the cosine similarity (Cos) and inverse softmax (InvSM) as described in Sec. 3.3. The results prove that the proposed correlation mechanism is conducive to the final performance.

Distillation ablation. To verify the effectiveness of correlation map distillation, we set up two controlled experiments under COCO-20ⁱ 1-shot setting. Fig. 1 shows the activation maps of correlation.

Fig. 6 plots the training and validation results with or without correlation map distillation on every COCO-20ⁱ split. The correlation map distillation has limited influence on the training process. But it largely improves the validation results.

Finally, we compare different loss functions in Table 6. We directly adopt the cross entropy (CE) loss between the correlation maps and the one-hot labels in ground-truth annotations and then use knowledge distillation (KD) by se-

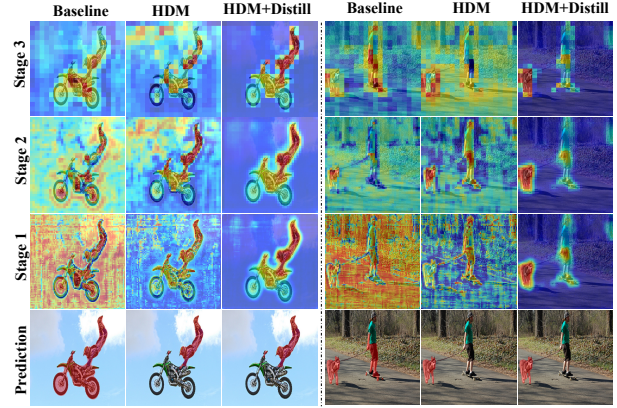


Figure 7. Qualitative correlation maps in 1-3 matching stages. The target classes of the left and right panels are “people” and “dog” respectively.

lecting adjacent layers as the teacher and student. KL denotes the Kullback-Leibler divergence loss.

Effect of the matching pyramid. Table 7 compares the decoders of previous methods and our proposed matching pyramid with different stage numbers in terms of accuracy, efficiency, and model size. Fig. 7 visualizes qualitative results of correlation maps in 1-3 matching stages under distinct designs. We build our baseline by directly interleaving the self- and cross-attention layers and then convert it to the hierarchically decoupled matching (HDM) structure. HDMNet better fights against the interference of other classes in the same image but suffers from mining the correlation information from the shallow stages, which contain more detailed features but fewer semantic cues. We finally adopt correlation map distillation to facilitate the earlier layers to be more aware of the contextual information.

5. Concluding Remarks

In this paper, we have proposed hierarchically decoupled matching network (HDMNet) to tackle the challenging few-shot segmentation problem. HDMNet decouples the downsampling and matching process to prevent information interference. Further, we designed a novel matching module constructed on the correlation mechanism and distillation and performed extensive experiments to demonstrate that this design alleviates the training-class overfitting problem and improves generality. One limitation is that our model sometimes may fail to clearly distinguish different categories that have rather similar appearances and close semantic relations, such as truck and car, chair and sofa. It may be a promising direction for future research.

Acknowledgements. This work is partially supported by ITF Partnership Research Programme (No.PRP/65/20FX) and Shenzhen Science and Technology Program KQTD20210811090149095.

References

- [1] Malik Boudiaf, Hoel Kervadec, Ziko Imtiaz Masud, Pablo Piantanida, Ismail Ben Ayed, and Jose Dolz. Few-shot segmentation without meta-learning: A good transductive inference is all you need? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13979–13988, 2021. 1, 6
- [2] 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 transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. 1
- [3] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. 1
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [5] Nanqing Dong and Eric P Xing. Few-shot semantic segmentation with prototype learning. In *BMVC*, 2018. 1, 2
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [7] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010. 6
- [8] Qi Fan, Wenjie Pei, Yu-Wing Tai, and Chi-Keung Tang. Self-support few-shot semantic segmentation. *arXiv preprint arXiv:2207.11549*, 2022. 2, 7
- [9] Zhibo Fan, Jin-Gang Yu, Zhihao Liang, Jiarong Ou, Changxin Gao, Gui-Song Xia, and Yuanqing Li. Fgn: Fully guided network for few-shot instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9172–9181, 2020. 2
- [10] Siddhartha Gairola, Mayur Hemani, Ayush Chopra, and Balaji Krishnamurthy. Simpropnet: Improved similarity propagation for few-shot image segmentation. *arXiv preprint arXiv:2004.15014*, 2020. 2
- [11] Mark Hamilton, Zhoutong Zhang, Bharath Hariharan, Noah Snavely, and William T Freeman. Unsupervised semantic segmentation by distilling feature correspondences. *arXiv preprint arXiv:2203.08414*, 2022. 2
- [12] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *2011 international conference on computer vision*, pages 991–998. IEEE, 2011. 6
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 1, 6, 7
- [14] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015. 2, 5
- [15] Joakim Johnander, Johan Edstedt, Michael Felsberg, Fahad Shahbaz Khan, and Martin Danelljan. Dense gaussian processes for few-shot segmentation. In *European Conference on Computer Vision*, pages 217–234. Springer, 2022. 6
- [16] Chunbo Lang, Gong Cheng, Binfei Tu, and Junwei Han. Learning what not to segment: A new perspective on few-shot segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8057–8067, 2022. 6, 7, 8
- [17] Chunbo Lang, Binfei Tu, Gong Cheng, and Junwei Han. Beyond the prototype: Divide-and-conquer proxies for few-shot segmentation. *arXiv preprint arXiv:2204.09903*, 2022. 6
- [18] Gen Li, Varun Jampani, Laura Sevilla-Lara, Deqing Sun, Jonghyun Kim, and Joongkyu Kim. Adaptive prototype learning and allocation for few-shot segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8334–8343, 2021. 2, 7
- [19] Zhaoshuo Li, Xingtong Liu, Nathan Drenkow, Andy Ding, Francis X Creighton, Russell H Taylor, and Mathias Unberath. Revisiting stereo depth estimation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6197–6206, 2021. 2
- [20] 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 *European conference on computer vision*, pages 740–755. Springer, 2014. 6
- [21] Jinlu Liu and Yongqiang Qin. Prototype refinement network for few-shot segmentation. *arXiv preprint arXiv:2002.03579*, 2020. 6
- [22] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 1
- [23] Zhihe Lu, Sen He, Xi Tian Zhu, Li Zhang, Yi-Zhe Song, and Tao Xiang. Simpler is better: Few-shot semantic segmentation with classifier weight transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8741–8750, 2021. 1, 2, 4, 6
- [24] Xiaoliu Luo, Zhuotao Tian, Taiping Zhang, Bei Yu, Yuan Yan Tang, and Jiaya Jia. Pfenet++: Boosting few-shot semantic segmentation with the noise-filtered context-aware prior mask. *arXiv preprint arXiv:2109.13788*, 2021. 2
- [25] Juhong Min, Dahyun Kang, and Minsu Cho. Hypercorrelation squeeze for few-shot segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6941–6952, 2021. 1, 2, 6, 7, 8
- [26] Khoi Nguyen and Sinisa Todorovic. Feature weighting and boosting for few-shot segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 622–631, 2019. 1, 6

- [27] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. [6](#)
- [28] Ignacio Rocco, Mircea Cimpoi, Relja Arandjelović, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Neighbourhood consensus networks. *Advances in neural information processing systems*, 31, 2018. [5](#)
- [29] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots. One-shot learning for semantic segmentation. *arXiv preprint arXiv:1709.03410*, 2017. [1](#), [6](#), [7](#)
- [30] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots. One-shot learning for semantic segmentation. *arXiv preprint arXiv:1709.03410*, 2017. [3](#)
- [31] Xinyu Shi, Dong Wei, Yu Zhang, Donghuan Lu, Munan Ning, Jiashun Chen, Kai Ma, and Yefeng Zheng. Dense cross-query-and-support attention weighted mask aggregation for few-shot segmentation. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XX*, pages 151–168. Springer, 2022. [7](#)
- [32] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [1](#), [6](#)
- [33] Samuel L. Smith, David H. P. Turban, Steven Hamblin, and Nils Y. Hammerla. Offline bilingual word vectors, orthogonal transformations and the inverted softmax. In *International Conference on Learning Representations*, 2017. [5](#)
- [34] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30, 2017. [1](#), [2](#)
- [35] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. Loftr: Detector-free local feature matching with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8922–8931, 2021. [2](#), [3](#)
- [36] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1199–1208, 2018. [2](#)
- [37] Zhuotao Tian, Xin Lai, Li Jiang, Shu Liu, Michelle Shu, Hengshuang Zhao, and Jiaya Jia. Generalized few-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11563–11572, 2022. [2](#), [6](#)
- [38] Zhuotao Tian, Hengshuang Zhao, Michelle Shu, Zhicheng Yang, Ruiyu Li, and Jiaya Jia. Prior guided feature enrichment network for few-shot segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 2020. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [39] Joaquin Vanschoren. Meta-learning: A survey. *arXiv preprint arXiv:1810.03548*, 2018. [1](#)
- [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [2](#), [4](#), [8](#)
- [41] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *Advances in neural information processing systems*, 29, 2016. [1](#), [2](#)
- [42] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9197–9206, 2019. [1](#), [2](#), [3](#), [6](#), [7](#)
- [43] Yaqing Wang, Quanming Yao, James T Kwok, and Lionel M Ni. Generalizing from a few examples: A survey on few-shot learning. *ACM computing surveys (csur)*, 53(3):1–34, 2020. [1](#)
- [44] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34:12077–12090, 2021. [2](#), [4](#)
- [45] Boyu Yang, Chang Liu, Bohao Li, Jianbin Jiao, and Qixiang Ye. Prototype mixture models for few-shot semantic segmentation. In *European Conference on Computer Vision*, pages 763–778. Springer, 2020. [6](#)
- [46] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9587–9595, 2019. [1](#)
- [47] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9587–9595, 2019. [2](#)
- [48] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5217–5226, 2019. [1](#), [3](#)
- [49] Gengwei Zhang, Guoliang Kang, Yi Yang, and Yunchao Wei. Few-shot segmentation via cycle-consistent transformer. *Advances in Neural Information Processing Systems*, 34:21984–21996, 2021. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [50] Linfeng Zhang, Jiebo Song, Anni Gao, Jingwei Chen, Chenglong Bao, and Kaisheng Ma. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3713–3722, 2019. [2](#)
- [51] Xiaolin Zhang, Yunchao Wei, Yi Yang, and Thomas S Huang. Sg-one: Similarity guidance network for one-shot semantic segmentation. *IEEE transactions on cybernetics*, 50(9):3855–3865, 2020. [1](#)
- [52] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017. [1](#), [6](#)

- [53] Qi Zhao, Binghao Liu, Shuchang Lyu, and Huojin Chen. A self-distillation embedded supervised affinity attention model for few-shot segmentation. *IEEE Transactions on Cognitive and Developmental Systems*, 2023. 7