

Few-shot Semantic Segmentation with Support-induced Graph Convolutional Network

Jie Liu¹

j.liu5@uva.nl

Yanqi Bao²

yanqibao1997@gmail.com

Wenzhe Yin¹

w.yin@uva.nl

Haochen Wang¹

h.wang3@uva.nl

Yang Gao²

gaoy@nju.edu.cn

Jan-Jakob Sonke³

j.sonke@nki.nl

Efstratios Gavves¹

egavves@uva.nl

¹ University of Amsterdam
Amsterdam, Netherlands

² Nanjing University
Nanjing, China

³ The Netherlands Cancer Institute
Amsterdam, Netherlands

Abstract

Few-shot semantic segmentation (FSS) aims to achieve novel objects segmentation with only a few annotated samples and has made great progress recently. Most of the existing FSS models focus on the feature matching between support and query to tackle FSS. However, the appearance variations between objects from the same category could be extremely large, leading to unreliable feature matching and query mask prediction. To this end, we propose a Support-induced Graph Convolutional Network (SiGCN) to explicitly excavate latent context structure in query images. Specifically, we propose a *Support-induced Graph Reasoning (SiGR)* module to capture salient query object parts at different semantic levels with a Support-induced GCN. Furthermore, an *instance association (IA)* module is designed to capture high-order instance context from both support and query instances. By integrating the proposed two modules, SiGCN can learn rich query context representation, and thus being more robust to appearance variations. Extensive experiments on PASCAL-5ⁱ and COCO-20ⁱ demonstrate that our SiGCN achieves state-of-the-art performance.

1 Introduction

The development of deep convolutional neural networks has achieved great success in many computer vision tasks, such as image classification [8, 26], semantic segmentation [9, 65],

[43, 45], and object detection [6, 8, 9, 20]. Despite the effectiveness of deep networks, their over-reliance on large-scale annotated dataset still remains a fundamental limitation as data annotation requires large amount of human efforts, especially for dense prediction tasks, *e.g.*, semantic segmentation. To cope with these challenges, some weakly-supervised and semi-supervised semantic segmentation techniques [8, 11, 24] try to introduce weak annotations (*e.g.*, points, scribbles, and bounding box) and unlabelled samples. However, these techniques usually fail to work when novel classes that never appear in the training set emerge. In such case, few-shot semantic segmentation (FSS) methods [8, 12, 13, 17, 19, 20, 25, 27, 29, 30, 31, 32, 37, 39, 40] struggle to segment novel objects with only a few pixel-wise annotated samples.




Figure 1: Comparison between (a) previous typical few-shot segmentation (FSS) methods and (b) the proposed support-induced graph convolutional network (SiGCN). (a) Previous FSS methods focus on designing effective feature matching mechanism between support prototypes/features and query feature. However, such matching techniques usually fail to work when the support and query objects exhibit large appearance variations. (b) Our SiGCN aims to address the appearance variation problem by explicitly mining query context information. Based on the activation map, SiGCN constructs a query graph to associate salient query object parts. Then the support prototypes are introduced to guide the graph updating process for effective query context excavation.

Meta-learning [23] is widely adopted in FSS to ensure better model generalization to novel classes. Specifically, both training and testing sets are composed of many sub-tasks (episodes), and each specific sub-task is divided into a support set and a query set, which share the same object classes. The FSS models aim to segment images from the query set given a few pixel-wise annotated images from the support set. Thus, the key challenge of the FSS task is how to leverage the information from the support set and the query set effectively, especially when the support and query objects exhibit large variations (*e.g.*, appearance, scale, and location).

Although existing FSS methods [12, 27, 40] have achieved extensive progress, we observe a significant drop in performance when the support and query images share large appearance variance. For instance, as shown in Figure 1, the birds in the support and query images have completely different appearance, which makes it difficult to segment the query image. Existing methods (Figure 1 (a)) usually generate support foreground features or prototypes and then focus on designing various feature matching mechanisms between the support foreground features/prototypes and the query feature to achieve the query mask prediction. However, such matching with information coming only from the support set is suboptimal when it comes to predicting an accurate mask for the query image. By contrast, different query object parts (*e.g.*, the wings and the main body of the bird) contain rich

contextual information, which is beneficial to the query object segmentation.

In this work, we propose a support-induced graph convolutional network (SiGCN) to tackle the aforementioned challenge. Our main idea is to explicitly enrich the query context representations with the guidance of support information in two ways. Firstly, we propose a novel support-induced graph reasoning (SiGR) module to extract query context from different query object parts. Specifically, as shown in Figure 1(b), we utilize the activation map, which is generated by the feature matching between the support and the query features, to capture local object parts that are discriminative and easily localized in the query image. These salient parts are associated together in the query graph by modeling the interaction between them, thus learning rich context in the query image. The query graph is then updated by the graph convolutional network, in which the state updating matrix is generated by the support prototypes to effectively extract query context. Secondly, an instance association (IA) module is introduced to aggregate query and support context from instance level. Specifically, given activation maps produced from different semantic levels, we get corresponding updated query graphs (i.e., query features) in the SiGR module. These updated query features, which can be seen as different object instances, contain complementary context information. Therefore, we associate the query instances and the support instance together to learn high-order instance context. By proposing the SiGR and IA module, our SiGCN aggregates rich complementary context from both support and query images, thus being more robust to appearance variations.

In summary, our contributions are summarized as follows:

- We propose a *support-induced graph convolution network (SiGCN)*, which utilizes the proposed IA and SiGR modules to capture complementary context from the query and the support set, for addressing the appearance variation problem in the FSS task.
- We propose a novel graph updating mechanism, in which support prototypes guide the node updating in the query graphs. This mechanism brings a significant performance gain to the proposed SiGCN.
- Our method achieves state-of-the-art results on two FSS benchmarks, i.e., PASCAL-5ⁱ and COCO-20ⁱ. Our SiGCN tackles the appearance variation problem from the perspective of extracting query contextual information, which sheds light for future research in this field.

2 Related Work

2.1 Semantic Segmentation

Semantic Segmentation is a fundamental computer vision task which aims to assign each pixel with a class label. Fully convolutional network (FCN) [14] adopts fully convolutional layers for dense prediction and achieves superior performance. Inspired by FCN, many networks have been invented to boost the performance of semantic segmentation. Unet [22] presents an encoder-decoder structure to capture multi-scale context information. PSPNet [45] utilizes the pyramid pooling to aggregate more object details. Dilated convolution [45] and deformable convolution [6] are introduced to FCN to further improve the semantic segmentation performance. Recently, K-Net [43] proposes to segment both instance and semantic categories by a group of learnable kernels. However, these methods still require tons of

samples with pixel-wise annotations in the training stage. In our work, we focus on few-shot semantic segmentation.

2.2 Few-shot Semantic Segmentation

Few-shot semantic segmentation is proposed to segment novel objects given only few annotated samples. Most existing methods adopt episode-base meta-learning [61] strategy to ensure the efficient generalization of the model to the novel classes with few examples. Shaban et al. [23] proposes the first few-shot semantic segmentation work, which generates parameters from the conditioning branch and then segments the query image in the segmentation branch. Later on, prototype-based methods are proposed including CANet [39], PANet [29], SG-One [44], PPNet [13], PMMs [33]. The intuition behind these methods lies in extracting representative foreground or background prototypes for support samples and then interacting these prototypes with query features using different strategies. Besides, matching-based methods [17, 27, 31, 38, 40] also achieve great success in the few shot-segmentation task. PFENet [27] introduces a prior mask by implementing pixel-to-pixel matching between high-level support and query features, while PGNet [38] and SAGNN [6] propose to implement element-wise matching between support and query by graph neural networks. Most recently, HSNet [17] adopts multi-scale dense matching to construct hypercorrelations and uses 4D convolutions to capture contextual information. CyCTR [40] designs a cycle consistency transformer to aggregate pixel-wise support feature to query ones and achieves great success. Different from previous work, we not only consider the matching between support and query, but also highlight context excavation in the query feature, our model is thus more robust to the object appearance variations.

2.3 Graph Convolutional Network

Graph convolutional networks (GCNs) [10], first proposed to tackle the semi-supervised classification, have seen massive application in computer vision tasks, e.g., action localization [18, 56], object detection [16], and semantic segmentation [42]. Recently, GCNs are introduced to few-shot learning and achieve superior performance. Zhang et al [41] propose an adaptive aggregation GCN (AdarGCN) to model the information passing between noisy samples and underlying training samples. Yan et al [52] design a GCN based data selection policy to compute a context-sensitive representation for each unlabeled data by graph message passing. Furthermore, han et al [9] introduce a heterogeneous GCN with three different types of edges to promote the pairwise matching for the few-shot object detection task. Inspired by the powerful relation modelling capacity of GCNs, we design a support-induced graph convolutional network (SiGCN) to enrich query contextual representation for addressing the object appearance variation problem in FSS.

3 Method

3.1 Problem Setting

To improve the model generalizability over previously unseen classes, most methods adopt the episode-based meta-learning paradigm [23, 28]. Models are trained on the base classes \mathcal{C}_{base} and tested on novel classes \mathcal{C}_{novel} . \mathcal{C}_{base} and \mathcal{C}_{novel} are disjoint, i.e., $\mathcal{C}_{base} \cap \mathcal{C}_{novel} = \emptyset$.

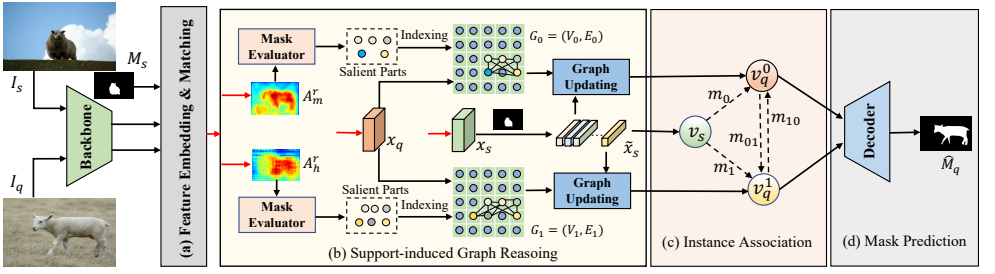


Figure 2: **Overall architecture of our proposed network.** Firstly, the feature embedding and matching network are introduced to generate support feature x_s , query feature x_q , middle-level activation map A_m^r , and high-level activation map A_h^r . Then, the support-induced graph reasoning module captures salient query object parts and associates them by a graph convolutional network. The support prototypes \tilde{x}_s is introduced to the graph updating for effective query context excavation. In addition, the instance association module aims to mine instance-level context by the message passing between query instances (i.e., v_q^0 and v_q^1) and a support instance v_s . Finally, the updated features are concatenated and fed into a decoder for the final query mask prediction.

Both training and inference are based on the episode data $(\mathcal{S}, \mathcal{Q})$, which is composed of a support set \mathcal{S} , and a query set \mathcal{Q} for a specific class c . For the 1-way K -shot setting, the support set $\mathcal{S} = \{(I_s^i, M_s^i)\}_{i=1}^K$ contains K image-mask pairs, where I_s^i represents i th support image and M_s^i indicates corresponding binary mask. Similarly, we define the query set as $\mathcal{Q} = \{(I_q, M_q)\}$, where I_q is query image and its binary mask M_q is only available in the model training phase. In the meta-training stage, the FSS model takes input \mathcal{S} and I_q from a specific class c and generates a predicted mask \hat{M}_q for the query image. Then, the model can be trained with the supervision of a binary cross-entropy loss between M_q and \hat{M}_q , i.e., $\mathcal{L}_{BCE}(M_q, \hat{M}_q)$. Finally, multiple episodes $(\mathcal{S}_i^s, \mathcal{Q}_i^s)_{i=1}^{N_{Ts}}$ are randomly sampled from the testing set \mathcal{D}_{Ts} for evaluation.

3.2 Overview

Figure 2 illustrates the overall architecture of our proposed network, which consists of three main components: feature embedding and matching module, support-induced graph reasoning (SiGR) module and instance association (IA) module. Given the support set $\mathcal{S} = \{(I_s^i, M_s^i)\}_{i=1}^K$ and query image I_q , we first encode them into feature maps by a weight-sharing backbone. Then, the feature embedding and matching network is introduced to generate support/query features $x_s \in \mathbb{R}^{C \times H \times W}$ and $x_q \in \mathbb{R}^{C \times H \times W}$, middle-level/high-level activation map $A_m^r \in \mathbb{R}^{H \times W}$ and $A_h^r \in \mathbb{R}^{H \times W}$, where C and HW are the channel and spatial dimension, respectively. We obtain activation maps by region-to-region matching between support and query features [15]. Afterwards, the support-induced graph reasoning module is designed to enrich the query contextual representation by a novel support-induced graph convolutional network. Finally, we aggregate complementary context from both support and query by instance-level message passing in the instance association module.

3.3 Support-induced Graph Reasoning

The proposed SiGR module consists of two parallel branches. Taking the upper branch as shown in Figure 2 as an example, we explain the support-induced graph reasoning process in details.

Mask Evaluation. The coarse activation map for the query feature is usually not accurate enough and contains noise and non-salient parts. Therefore, we design a mask evaluator to select salient objects parts. Given the activation map A_m^r , we first process it with a min-max normalization to normalize the value to between 0 and 1. Then, we select salient query object parts by introducing a salience matrix $S_m^r \in \mathbb{R}^{H \times W}$:

$$S_m^r(i, j) = \begin{cases} 1, & A_m^r(i, j) \geq t \\ 0, & A_m^r(i, j) < t \end{cases}, \quad (1)$$

where S_m^r is a sparse matrix, in which the elements with value 1 represent target location, otherwise noise or undiscriminating parts location. t is a threshold which decides whether a pixel belongs to salient parts or not.

Graph Construction. Based on the defined salience matrix S_m^r , we select $N_m = \sum_{i,j}^{HW} S_m^r(i, j)$ most salient pixels in the query feature x_q as the set of salient parts $P_m = \{P_m^n\}_{n=1}^{N_m}$, which are discriminative local parts of the object for segmentation. The captured salient object parts are further associated together by a fully connected graph to learn rich query contextual information. Specifically, we treat each pixel in the query feature as a node and build the graph $G_0 = (V_0, E_0)$, where V_0 and E_0 are the node set the edge set, respectively. For the connecting edges, We adopt fully-connected edges between N_m salient parts specified in P_m . To achieve better generalization over novel objects, the edge weights are designed to be the similarity score between two connected nodes. Thus, the adjacency matrix $A_0 \in \mathbb{R}^{N \times N}$, $N = HW$, for the graph G_0 is constructed as:

$$A_0(i, j) = \begin{cases} \frac{v_i^T v_j}{\|v_i\| \|v_j\|}, & \text{if edge } < i, j > \in \mathcal{E} \\ 0, & \text{else} \end{cases}, \quad (2)$$

where $v_i \in \mathbb{R}^C$ and $v_j \in \mathbb{R}^C$ are two node vectors in the connection, \mathcal{E} is the edge connection set of all the salient parts feature vectors.

Graph Updating. To update graphs by aggregating the salient parts' information, we propose a support-induced graph convolution network (SiGCN) to achieve effective information aggregation with the guidance of support foreground prototypes. For the query graph V_0 , we first construct a two-layer GCN to perform the node state updating. Formally, the l -th layer graph convolution is formulated as:

$$X^{l+1} = \sigma^l(\hat{A}X^lW^l), \quad (3)$$

where $X^l \in \mathbb{R}^{N \times d_l}$ and $X^{l+1} \in \mathbb{R}^{N \times d_{l+1}}$ are the input and output feature of all nodes at the l -th layer, respectively, d_l and d_{l+1} are the corresponding feature dimensions, d_0 is equal to the node feature dimension C . $\hat{A} = \tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{\frac{1}{2}}$ is the normalized adjacency matrix [14], where $\tilde{A} = A + I$ and \tilde{D} is the diagonal degree matrix of \tilde{A} . $W^l \in \mathbb{R}^{d_l \times d_{l+1}}$ and $\sigma^l(\cdot)$ denote the state updating matrix and activation function at the layer l , respectively.

Normally, the state updating matrix W is optimized with the model training proceeds. Chen et. al [14] propose to adopt a 1D convolution as state updating function. Inspired by

this, we implement the state updating matrix as a 1D convolution with fixed kernel generated from the support foreground vectors.

$$X^{l+1} = \sigma^l(\hat{A}\mathcal{F}_c^l(X^l|\theta^l)), \quad (4)$$

where $\mathcal{F}_c^l(\cdot)$ is the 1D convolution with fixed kernel parameters θ^l at the l -th layer. $\theta \in \mathbb{R}^{k \times C}$, where k is the number of support prototypes (i.e., kernel size of 1D convolution), is generated by applying an average pooling to the support foreground feature. Given the high-level activation map A_h^r , we can similarly select salient query object parts, construct the query graph $G_1 = (V_1, E_1)$, and implement support-induced graph updating. Finally, we can obtain two updated query graphs, which aggregate rich query contextual representation with the guidance of support prototypes.

3.4 Instance Association

Although rich context has been aggregated into query graphs, the complementary inter-graph context still remains unexploited. Therefore, the updated query graphs are further transformed into query instance features $v_q^0 \in \mathbb{R}^{C \times H \times W}$ and $v_q^1 \in \mathbb{R}^{C \times H \times W}$, respectively. Meanwhile, we can also adopt the 1D average pooling and the reshape operation over the support foreground prototypes \tilde{x}_s to get the support instance feature $v_s \in \mathbb{R}^{C \times s \times s}$, where s is the spatial size. Then, the the query instance features and support instance feature are associated together by the instance association module to learn instance-level contextual information. For the query instance feature v_q^0 , the information gathered from the support instance feature v_s and the query instance feature v_q^1 can be formulated as:

$$m_0 = \mathcal{R}(v_s)\mathcal{R}(v_s)^T\mathcal{R}(v_q^0), m_{10} = \mathcal{R}(v_q^1)\mathcal{R}(v_q^1)^T\mathcal{R}(v_q^0), \quad (5)$$

where $m_0 \in \mathbb{R}^{C \times HW}$ and $m_{10} \in \mathbb{R}^{C \times HW}$ are the instance-level contextual information from support instance and query instance, respectively. $\mathcal{R}(\cdot)$ is the vectorization operation. Then, the query instance feature can be updated as \tilde{v}_q^0 by:

$$\tilde{v}_q^0 = \frac{1}{2}(v_q^0 + (\alpha m_0 + \beta m_{10})), \quad (6)$$

where α and β are two hyperparameters to weight the contribution of the contextual information from the support and query, respectively. Similarly, the other query instance feature v_q^1 can be updated as \tilde{v}_q^1 . Compared with previous state, \tilde{v}_q^0 and \tilde{v}_q^1 aggregate more instance-level context information from both support and query, thus being more robust to the appearance variations.

3.5 Mask Prediction and Model Training

Finally, we feed the updated query instance features as well as activation maps specified in Section 3.2 to make pixel-wise prediction for query.

$$\hat{M}_q = \mathcal{F}_{decoder}(\mathcal{C}(\tilde{v}_q^0, \tilde{v}_q^1, A_m^p, A_m^r, A_h^p, A_h^r)), \quad (7)$$

where $\mathcal{C}(\cdot)$ is the channel-wise concatenation. $\mathcal{F}_{decoder}(\cdot)$ is the decoder network, which consists of a Atrous Spatial Pyramid Pooling (ASPP) module [2] and three consecutive convolution layers with residual connection. For the activation maps A_m^p and A_h^p , the superscript

Table 1: Comparison with the state-of-the-art methods on PASCAL-5ⁱ dataset under both 1-shot and 5-shot settings. The mIoU of each fold, and the averaged mIoU & FB-IoU of all folds are reported.

Methods	Backbone	1-shot						5-shot					
		Fold-0	Fold-1	Fold-2	Fold-3	Mean	FB-IoU	Fold-0	Fold-1	Fold-2	Fold-3	Mean	FB-IoU
OSLSM (BMVC'17) [40]	VGG16	33.6	55.3	40.9	33.5	40.8	61.3	35.9	58.1	42.7	39.1	43.9	61.5
co-FCN (ICLRW'18) [41]	VGG16	36.7	50.6	44.9	32.4	41.1	60.1	37.5	50.0	44.1	33.9	41.4	60.2
PFENet (TPAMI'20) [27]	VGG16	56.9	68.2	54.4	52.4	58.0	72.0	59.0	69.1	54.8	52.9	59.0	72.3
HSNet (ICCV'21) [42]	VGG16	59.6	65.7	59.6	54.0	59.7	73.4	64.9	69.0	64.1	58.6	64.1	76.6
PFENet (TPAMI'20) [27]	ResNet50	61.7	69.5	55.4	56.3	60.8	73.3	63.1	70.7	55.8	57.9	61.9	73.9
RePRI (CVPR'21) [43]	ResNet50	59.8	68.3	62.1	48.5	59.7	-	64.6	71.4	71.1	59.3	66.6	-
SAGNN (CVPR'21) [44]	ResNet50	64.7	69.6	57.0	57.3	62.1	73.2	64.9	70.0	57.0	59.3	62.8	73.3
MMNet (ICCV'21) [45]	ResNet50	62.7	70.2	57.3	57.0	61.8	-	62.2	71.5	57.5	62.4	63.4	-
HSNet (ICCV'21) [42]	ResNet50	64.3	70.7	60.3	60.5	64.0	76.7	70.3	73.2	67.4	67.1	69.5	80.6
CyCTR (NIPS'21) [46]	ResNet50	67.8	72.8	58.0	58.0	64.2	-	71.1	73.2	60.5	57.5	65.6	-
Baseline	ResNet50	62.5	69.4	58.9	55.9	61.7	72.0	63.2	70.5	60.1	57.6	62.9	74.4
SiGCN	ResNet50	65.1	70.1	65.2	60.8	65.3	77.5	68.9	72.6	66.8	65.8	68.5	78.3

Table 2: Comparison with the state-of-the-art methods on COCO-20ⁱ dataset under both 1-shot and 5-shot settings. The mIoU of each fold, and the averaged mIoU & FB-IoU of all folds are reported.

Methods	Backbone	1-shot						5-shot					
		Fold-0	Fold-1	Fold-2	Fold-3	Mean	FB-IoU	Fold-0	Fold-1	Fold-2	Fold-3	Mean	FB-IoU
FWB (ICCV'19) [47]	VGG16	18.4	16.7	19.6	25.4	20.0	-	20.9	19.2	21.9	28.4	22.6	-
PFENet (TPAMI'20) [27]	VGG16	33.4	36.0	34.1	32.8	34.1	60.0	35.9	40.7	38.1	36.1	37.7	61.6
SAGNN (CVPR'21) [44]	VGG16	35.0	40.5	37.6	36.0	37.3	61.2	37.2	45.2	40.4	40.0	40.7	63.1
RePRI (CVPR'21) [43]	ResNet50	31.2	38.1	33.3	33.0	34.0	-	38.5	46.2	40.0	43.6	42.1	-
SAGNN (CVPR'21) [44]	ResNet101	36.1	41.0	38.2	33.5	37.2	60.9	40.9	48.3	42.6	38.9	42.7	63.4
MMNet (ICCV'21) [45]	ResNet50	34.9	41.0	37.2	37.0	37.5	-	37.0	40.3	39.3	36.0	38.2	-
HSNet (ICCV'21) [42]	ResNet50	36.3	43.1	38.7	38.7	39.2	68.2	43.3	51.3	48.2	45.0	46.9	70.7
CyCTR (NIPS'21) [46]	ResNet50	38.9	43.0	39.6	39.8	40.3	-	41.1	48.9	45.2	47.0	45.6	-
Baseline	ResNet50	33.6	39.2	36.5	34.2	35.9	60.7	35.3	43.1	38.4	38.7	38.9	62.0
SiGCN	ResNet50	38.7	46.3	43.1	37.5	41.4	62.7	44.9	54.5	46.5	45.9	48.0	66.2

indicates that the activation map comes from pixel-to-pixel matching (p) [48] between support and query features. To generate the final segmentation map, we upsample \hat{M}_q by bilinear interpolation. Finally, we adopt the binary cross entropy loss to train our model on all the episodes in the training set \mathcal{D}_{tr} .

4 Experiments

4.1 Implementation Details

Following previous methods [49, 48], we conduct experiments on PASCAL-5ⁱ and COCO-20ⁱ. Classes in each dataset are divided into four base-novel class folds for cross validation. Our model is based on PFENet [27] and we also adopt the same training and evaluation setting as PFENet. We train our model with SGD optimizer on PASCAL-5ⁱ for 200 epochs and COCO-20ⁱ for 50 epochs. The learning rate is initialized as 0.005 with batch size 8 and decays according to the poly learning rate scheduler. During evaluation, we adopt the mIoU as main metric and also report FB-IoU. For the K-shot setting, we fuse foreground prototypes from each support image-mask pair to generate θ in Eq. (4) and the support instance feature v_s in the instance association module.

4.2 Comparison with the State-of-the-art methods

PASCAL-5ⁱ. In Table 1, we show the performance comparison between the proposed SiGCN and the leading methods on the PASCAL-5ⁱ dataset. With the ResNet50 as backbone, we can conclude that the proposed SiGCN achieves the state-of-the-art performance



Figure 3: Qualitative results on the PASCAL-5ⁱ (left four columns) and COCO-20ⁱ (right four columns) under the 1-shot setting. Our SiGCN can make accurate segmentation prediction, even though the support and query objects exhibit larger appearance variations.

under the 1-shot setting and competitive result under the 5-shot setting. Specifically, SiGCN surpasses previous SOTA (i.e., CyCTR) by 1.1% in the 1-shot setting, while achieving competitive results as previous SOTA HSNet (68.5 vs 69.5) in the 5-shot setting. Additionally, compared with the baseline model, which is constructed by removing the SiGR and IA modules from SiGCN (see complementary materials for more details), our SiGCN achieves significant improvement, i.e., 3.6% and 6.4% under the 1-shot and 5-shot settings, respectively. Considering the comparison with previous SOTA methods using the FB-IoU, SiGCN also obtain competitive results (77.5 and 78.3 under the 1-shot and 5-shot setting, respectively). **COCO-20ⁱ.** Table 2 shows the comparison between previous SOTA methods and our SiGCN on the COCO-20ⁱ dataset. Although the COCO-20ⁱ is more challenging (larger appearance and scale variations than PASCAL-5ⁱ), SiGCN also achieves state-of-the-art results (41.4 and 48.0 under 1-shot and 5-shot settings, respectively). It demonstrates the superiority of our proposed SiGCN in addressing the appearance variation problem. Furthermore, SiGCN with ResNet50 backbone achieves significant mIoU improvement (i.e., 5.5% and 9.1% under the 1-shot and 5-shot settings, respectively) over the baseline model.

Qualitative Results. In Figure 3, we report some typical qualitative results generated by our SiGCN. As can be seen, the support and query object exhibits large appearance variations, but our approach can also produce accurate mask prediction for the query image, which thanks to the query context mining by our method. More experimental and quantitative results can be found in the supplementary material.

Table 3: Effect of key modules in SiGCN.

SiGR	IA	1-shot mIoU					FB-IoU
		5 ⁰	5 ¹	5 ²	5 ³	mean	
		62.5	69.4	58.9	55.9	61.7	72.0
	✓	62.0	70.0	58.5	59.4	62.7	73.5
✓		64.6	69.8	64.1	60.8	64.8	77.4
✓	✓	65.1	70.1	65.2	60.8	65.3	77.5

Table 4: Effect of different GCN variants.

GCN Variants	1-shot mIoU					FB-IoU
	5 ⁰	5 ¹	5 ²	5 ³	mean	
baseline	62.0	70.0	58.5	59.4	62.7	73.5
GCN [10]	63.0	69.4	59.7	57.1	62.3	73.3
SiGCN-1	63.4	69.6	62.2	60.0	63.8	75.4
SiGCN-2	65.1	70.1	65.2	60.8	65.3	77.5

4.3 Ablation Study

Component analysis. There are two main components in our model, namely the support-induced graph reasoning (SiGR) module and the instance association (IA) module. We validate the effectiveness of each component and present the results in Table 3. It can be summarized that the SiGR module plays the most important role in the performance improvement while the IA module is indispensable. With SiGR and IA module, our SiGCN




Figure 4: Ablation studies of (a) support instance size s and (b) mask threshold t on PASCAL-5ⁱ under the 1-shot setting.

achieves state-of-the-art performance.

Effect of different GCN variants. To investigate the effect of different GCN variants in the target association module, we adopt three GCN variants, i.e., plain GCN, SiGCN with one and two layers (denoted as SiGCN-1 and SiGCN-2, respectively). As shown in Table 4, our model with plain GCN achieves slightly better performance than the baseline model, while our model with SiGCN makes substantial performance improvement as the number of the SiGCN layer increasing from one to two.

Table 5: Effect of k prototypes.

k prototypes	1-shot mIoU					FB-IoU
	s^0	s^1	s^2	s^3	mean	
1	64.8	69.7	63.5	59.0	64.3	75.9
3	64.1	70.2	64.3	59.5	64.5	76.2
5	65.1	70.1	65.2	60.8	65.3	77.5
7	63.7	69.8	63.8	60.2	64.4	75.5

Table 6: Effect of different backbones.

Backbone	1-shot mIoU					FB-IoU
	s^0	s^1	s^2	s^3	mean	
VGG16	58.0	68.0	62.9	54.2	60.8	73.5
ResNet-50	65.1	70.1	65.2	60.8	65.3	77.5
ResNet-101	65.4	71.0	64.6	61.7	65.7	78.3

Influence of k support prototypes. The number of support prototypes k in Eq. (4) is a key parameter in SiGCN, which determines the quality of support foreground used for guiding the query graph updating. We vary k in $\{1, 3, 5, 7\}$ to observe the performance of our model. In Table 5, $k = 5$ achieves best performance, which indicates that SiGCN with 5 support foreground prototypes provides better guidance to query graphs for capturing contextual information.

Influence of support instance size s . The support instance feature size s in the instance association module (section 3.4) determines the context information in the support instance. The value of s is taken from $\{1, 3, 5, 10, 20\}$ to observe the performance variations. As in Figure 4(a), SiGCN with $s = 10$ exhibits the best performance, thus we set $s = 10$ in all of our experiments.

Influence of mask threshold t . The mask threshold t in Eq.(1) is essential for selecting discriminative parts, i.e., targets in the query graph. Thus we take its value from 0, 0.1, 0.3, 0.5, 0.7, 0.9 to investigate the performance of our SiGCN. As can be seen in Figure 4(b), SiGCN with $t = 0.7$ performs best. Additionally, SiGCN without mask threshold in Eq.(5) (i.e., $t = 0$) also shows competitive results.

Effect of different backbones. To evaluate the performance of our model with different backbones, we adopt three backbone networks (i.e., VGG16, ResNet50, and ResNet101) to implement experiments. As shown in Table 6, our SiGCN with ResNet101 as backbone yields superior performance.

5 Conclusion

In this work, we propose a support-induced graph convolutional network (SiGCN) to tackle the challenging appearance variation problem for the few-shot semantic segmentation (FSS) task. Our main idea is to enrich query context representation with the discriminative parts in the query objects and support objects. To achieve this, we propose a support-induced graph reasoning (SiGR) module to associate salient query object parts with the graph convolution network, in which support prototypes is introduced to guide the graph updating process. Additionally, an instance association (IA) module is designed to capture high-order context from the support instances and complementary query instances simultaneously. Extensive experiments on two FSS benchmarks prove that our SiGCN achieves state-of-the-art performance under both 1-shot and 5-shot settings. We believe that the idea of explicitly mining query context will shed light for future research in this field.

Acknowledgement This work was partially funded by Elekta Oncology Systems AB and a RVO public-private partnership grant (PPS2102).

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.
- [2] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Re-thinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- [3] Xiaokang Chen, Yuhui Yuan, Gang Zeng, and Jingdong Wang. Semi-supervised semantic segmentation with cross pseudo supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2613–2622, 2021.
- [4] Yunpeng Chen, Marcus Rohrbach, Zhicheng Yan, Yan Shuicheng, Jiashi Feng, and Yannis Kalantidis. Graph-based global reasoning networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 433–442, 2019.
- [5] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 764–773, 2017.
- [6] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015.
- [7] Guangxing Han, Yicheng He, Shiyuan Huang, Jiawei Ma, and Shih-Fu Chang. Query adaptive few-shot object detection with heterogeneous graph convolutional networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3263–3272, 2021.

- [8] 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*, pages 770–778, 2016.
- [9] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [10] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [11] Jungbeom Lee, Jihun Yi, Chaehun Shin, and Sungroh Yoon. Bbam: Bounding box attribution map for weakly supervised semantic and instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2643–2652, 2021.
- [12] 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 (CVPR)*, pages 8334–8343, June 2021.
- [13] Yongfei Liu, Xiangyi Zhang, Songyang Zhang, and Xuming He. Part-aware prototype network for few-shot semantic segmentation. In *European Conference on Computer Vision*, pages 142–158. Springer, 2020.
- [14] 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.
- [15] 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.
- [16] Michael Meyer, Georg Kuschik, and Sven Tomforde. Graph convolutional networks for 3d object detection on radar data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3060–3069, 2021.
- [17] Juhong Min, Dahyun Kang, and Minsu Cho. Hypercorrelation squeeze for few-shot segmentation. *arXiv preprint arXiv:2104.01538*, 2021.
- [18] Megha Nawhal and Greg Mori. Activity graph transformer for temporal action localization. *arXiv preprint arXiv:2101.08540*, 2021.
- [19] 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.
- [20] Kate Rakelly, Evan Shelhamer, Trevor Darrell, Alyosha Efros, and Sergey Levine. Conditional networks for few-shot semantic segmentation. 2018.
- [21] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.

- [22] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [23] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots. One-shot learning for semantic segmentation. *arXiv preprint arXiv:1709.03410*, 2017.
- [24] Wataru Shimoda and Keiji Yanai. Self-supervised difference detection for weakly-supervised semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5208–5217, 2019.
- [25] Mennatullah Siam, Boris N Oreshkin, and Martin Jagersand. Amp: Adaptive masked proxies for few-shot segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5249–5258, 2019.
- [26] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [27] 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 & Machine Intelligence*, (01):1–1, 2020.
- [28] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *Advances in neural information processing systems*, 29:3630–3638, 2016.
- [29] 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.
- [30] Zhonghua Wu, Xiangxi Shi, Guosheng Lin, and Jianfei Cai. Learning meta-class memory for few-shot semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 517–526, 2021.
- [31] Guo-Sen Xie, Jie Liu, Huan Xiong, and Ling Shao. Scale-aware graph neural network for few-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5475–5484, 2021.
- [32] Shipeng Yan, Songyang Zhang, and Xuming He. Budget-aware few-shot learning via graph convolutional network. *arXiv preprint arXiv:2201.02304*, 2022.
- [33] 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.
- [34] Lihe Yang, Wei Zhuo, Lei Qi, Yinghuan Shi, and Yang Gao. Mining latent classes for few-shot segmentation. *arXiv preprint arXiv:2103.15402*, 2021.
- [35] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015.

- [36] Runhao Zeng, Wenbing Huang, Mingkui Tan, Yu Rong, Peilin Zhao, Junzhou Huang, and Chuang Gan. Graph convolutional networks for temporal action localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7094–7103, 2019.
- [37] Bingfeng Zhang, Jimin Xiao, and Terry Qin. Self-guided and cross-guided learning for few-shot segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8312–8321, 2021.
- [38] 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.
- [39] 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 (CVPR)*, June 2019.
- [40] Gengwei Zhang, Guoliang Kang, Yi Yang, and Yunchao Wei. Few-shot segmentation via cycle-consistent transformer. *Advances in Neural Information Processing Systems*, 34, 2021.
- [41] Jianhong Zhang, Manli Zhang, Zhiwu Lu, and Tao Xiang. Adargcn: adaptive aggregation gcnn for few-shot learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3482–3491, 2021.
- [42] Li Zhang, Xiangtai Li, Anurag Arnab, Kuiyuan Yang, Yunhai Tong, and Philip HS Torr. Dual graph convolutional network for semantic segmentation. *arXiv preprint arXiv:1909.06121*, 2019.
- [43] Wenwei Zhang, Jiangmiao Pang, Kai Chen, and Chen Change Loy. K-net: Towards unified image segmentation. *Advances in Neural Information Processing Systems*, 34, 2021.
- [44] 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.
- [45] 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.

6 Supplementary Materials

In this supplementary material, we firstly present extensive implementation details about our experiments, and then give two more additional ablation studies and analysis about different level features and matching methods for activation maps. Finally, we provide more qualitative visualization results for PASCAL-5ⁱ and COCO-20ⁱ benchmarks under the large object appearance and scale variations scenarios.

Implementation Details In our experiments, we employ ResNet-50 (VGG-16 and ResNet-101) pre-trained on ImageNet as our backbone network. For ResNet-50 and ResNet-101, the dilation convolution is introduced to ensure that the feature receptive fields of layer2, layer3, and layer4 preserve the same spatial resolution. The backbone weights are frozen except for layer4, which is required to learn more robust activation maps. The proposed model is validated on PASCAL-5ⁱ and COCO-20ⁱ benchmarks, which are widely used in the few-shot semantic segmentation for cross validations. The detailed split of testing classes for each cross validation (fold) is shown in Table 7 and Table 8, respectively.

The model is trained with a SGD optimizer for 200 and 50 epochs on the PASCAL-5ⁱ and the COCO-20ⁱ benchmarks, respectively. The learning rates are initialized as 0.005 and 0.002 with a poly learning rate schedule in PASCAL-5ⁱ and COCO-20ⁱ, respectively. Our entire network is trained with the same learning rate during each epoch, while layer4 of the backbone network should be ensured a lower learning rate for fine-tuning, thus its parameters starts back-propagation after training for multiple epochs. And the batch size is set as 8 on PASCAL-5ⁱ and 32 on COCO-20ⁱ. Data augmentation strategies like random rotation and flip are adopted in the training stage, and all images are cropped to 473×473 patches for two benchmarks. Besides, no post-processing is used on PASCAL-5ⁱ, while for COCO-20ⁱ, we adopt the multi-scale testing strategy for the model evaluation due to extremely large object appearance and scale variations. By the way, the original Groudtruth of the query image without any resize operations is adopted for the model evaluation.

In addition, the support instance size s in the IA module are set as 10, and mask threshold t are set as 0.7. For the SiGR module, we set the number of support prototypes $k = 5$ in each SiGCN layer. Finally, we implement our model with PyTorch 1.10.0 and conduct all the experiments with Nvidia Tesla A100 GPUs and CUDA11.3.

Table 7: Testing classes split for each fold in PASCAL-5ⁱ dataset.

Fold	Testing (novel) classes
Fold-0	Aeroplane, Bicycle, Bird, Boat, Bottle
Fold-1	Bus, Car, Cat, Chair, Cow
Fold-2	Diningtable, Dog, Horse, Motorbike, Person
Fold-3	Potted plant, Sheep, Sofa, Train, Tvmonitor

Effect of activation maps from different feature levels. To evaluate the quality of activation maps generated from different feature levels, we adopt three different-level features, i.e., block3 (middle-level), block4 (high-level), and block3&4 (middle-level& high-level). Support and query features from each level are matched by pixel-to-pixel matching to generate corresponding activation map. As shown in Table 9, our model with the activation map generated from block3&4 yields superior performance, while our model with the activation map generated from block3 or block4 also achieves competitive results.

Table 8: Testing classes split for each fold in COCO-20ⁱ dataset.

Fold	Testing (novel) classes
Fold-0	Person, Airplane, Boat, Parking meter, Dog, Elephant, Backpack, Suitcase, Sports Ball, Skateboard, Wine glass, Spoon, Sandwich, Hot dog, Chair, Dining table, Mouse, Microwave, Scissors
Fold-1	Bicycle, Bus, Traffic light, Bench, Horse, Bear, Umbrella, Frisbee, Kite, Surfboard, Cup, Bowl, Orange, Pizza, Couch, Toilet, Remote, Oven, Book, Teddy bear
Fold-2	Car, Train, Fire hydrant, Bird, Sheep, Zebra, Handbag, Skis, Baseball bat, Tennis racket, Fork, Banana, Broccoli, Donut, Potted plant, Tv, Keyboard, Sink, Toaster, Clock, Hair drier
Fold-3	Motorcycle, Truck, Stop sign, Cat, Cow, Giraffe, Tie, Snowboard, Baseball glove, Bottle, Knife, Apple, Carrot, Cake, Bed, Laptop, Cell phone, Refrigerator, Vase, Toothbrush

Effect of region matching. We append the activation maps generated by region-region matching to evaluate the performance of our model with region matching. As shown in Table 9, we can conclude that the activation maps generated from region-to-region matching bring overall performance improvement. Compared with pixel-to-pixel matching, region-to-region matching can capture more contextual information, thus leading to improved performance.

Table 9: Ablation study on the different activation maps on the PASCAL-5ⁱ dataset.

activation maps	1-shot mIoU					FB-IoU
	s^0	s^1	s^2	s^3	mean	
block3	64.7	70.3	61.1	60.6	64.2	75.0
block4	63.3	70.0	64.0	60.2	64.5	75.4
block3&4	64.5	70.1	64.3	60.4	64.8	75.9
block3+region	64.5	69.4	64.0	60.6	64.6	75.4
block4+region	62.9	69.5	65.1	60.1	64.4	75.2
block3&4+region	65.1	70.1	65.2	60.8	65.3	77.5

Additional Qualitative Results. We show more quantitative results of the proposed SiGCN in this section to further demonstrate its few-shot semantic segmentation performance. SiGCN aims to address the appearance variations problem in the few-shot semantic segmentation task, thus we show some examples sampled from PASCAL-5ⁱ and COCO-20ⁱ benchmarks with large appearance variations in Figure 5 and Figure 7, respectively. Furthermore, to demonstrate that our model can effectively address the scale variations problem, we give some examples from PASCAL-5ⁱ and COCO-20ⁱ benchmarks in Figure 6 and Figure 8, respectively. According to the quantitative results, we can conclude that our proposed SiGCN can effectively address the appearance and scale variations problem.




Figure 5: Qualitative results of the proposed SiGCN on PASCAL-5ⁱ benchmark with **large object appearance variations**. From top to bottom: Support Image & Ground truth, Query Image & Ground truth, and Query Prediction. Zoom in for details.




Figure 6: Qualitative results of the proposed SiGCN on PASCAL-5ⁱ benchmark with **large object scale variations**. From top to bottom: Support Image & Ground truth, Query Image & Ground truth, and Query Prediction. Zoom in for details.




Figure 7: Qualitative results of the proposed SiGCN on COCO-20ⁱ benchmark with **large object appearance variations**. From top to bottom: Support Image & Ground truth, Query Image & Ground truth, and Query Prediction. Zoom in for details.




Figure 8: Qualitative results of the proposed SiGCN on COCO-20ⁱ benchmark with **large object scale variations**. From top to bottom: Support Image & Ground truth, Query Image & Ground truth, and Query Prediction. Zoom in for details.