

SAM-Aware Graph Prompt Reasoning Network for Cross-Domain Few-Shot Segmentation

Shi-Feng Peng^{1*}, Guolei Sun^{2*}, Yong Li¹, Hongsong Wang³, Guo-Sen Xie^{1†}

¹School of Computer Science and Engineering, Nanjing University of Science and Technology, Nanjing, China

²Computer Vision Laboratory, ETH Zurich

³School of Computer Science and Engineering, Southeast University, Nanjing, China

psf.nerver@gmail.com; guolei.sun@vision.ee.ethz.ch; yong.li@njust.edu.cn; gsxiehm@gmail.com

Abstract

The primary challenge of cross-domain few-shot segmentation (CD-FSS) is the domain disparity between the training and inference phases, which can exist in either the input data or the target classes. Previous models struggle to learn feature representations that generalize to various unknown domains from limited training domain samples. In contrast, the large-scale visual model SAM, pre-trained on tens of millions of images from various domains and classes, possesses excellent generalizability. In this work, we propose a SAM-aware graph prompt reasoning network (GPRN) that fully leverages SAM to guide CD-FSS feature representation learning and improve prediction accuracy. Specifically, we propose a SAM-aware prompt initialization module (SPI) to transform the masks generated by SAM into visual prompts enriched with high-level semantic information. Since SAM tends to divide an object into many sub-regions, this may lead to visual prompts representing the same semantic object having inconsistent or fragmented features. We further propose a graph prompt reasoning (GPR) module that constructs a graph among visual prompts to reason about their interrelationships and enable each visual prompt to aggregate information from similar prompts, thus achieving global semantic consistency. Subsequently, each visual prompt embeds its semantic information into the corresponding mask region to assist in feature representation learning. To refine the segmentation masks during testing, we also design a non-parameter adaptive point selection module (APS) to select representative point prompts from query predictions and feed them back to SAM to refine inaccurate segmentation results. Experiments on four standard CD-FSS datasets demonstrate that our method establishes new state-of-the-art results.

Code — <https://github.com/CVL-hub/GPRN>.

Introduction

With the development of deep convolutional neural networks trained on large-scale datasets (Cordts et al. 2016; Zhou et al. 2017), numerous tasks (Li et al. 2018; Li, Zeng, and Shan 2020; Sun et al. 2022) have made promising progress. However, data annotation is a resource-intensive

*These authors contributed equally.

†Corresponding Author.




Figure 1: **Comparison of existing visual prompting methods and ours.** (a) Existing methods typically input randomly initialized visual prompts into the network alongside image tokens, which lack prior semantic information and spatial information. (b) Our approach leverages SAM to initialize task-specific visual prompts and constructs a graph convolutional network (GCN) to reason about their inherent relationships. Zoom in for details.

and time-consuming process, particularly when it involves dense pixel-wise labeling for segmentation tasks (Sun et al. 2022). Moreover, models trained on one category typically struggle to generalize to unseen categories, making it even more difficult to scale annotation efforts for diverse classes. In light of this, researchers propose few-shot segmentation (FSS), where the model can adapt to new tasks using a small number of samples (Xie et al. 2021; Peng et al. 2024).

Existing FSS methods can be broadly classified into two categories: prototype-based (Wang et al. 2019a; Liu et al. 2022; Yang et al. 2020) and correlation-based (Peng et al. 2023; Hong et al. 2022; Moon et al. 2023). Prototype-based methods extract prototypes from the support set and generate predictions by comparing the query features with these prototypes (Tian et al. 2020; Yang et al. 2020). Correlation-based methods model the pixel-to-pixel relationship between the query and support features to activate the query pixels that are similar to the support foreground pixels (Min, Kang, and Cho 2021; Shi et al. 2022). However, these methods are limited in scenarios where the base and novel classes are within the same domain.

To overcome the limitations of traditional FSS, Lei et al. (2022) extend the FSS task to cross-domain few-shot segmentation (CD-FSS) task, where the training and test-

ing datasets belong to different domains. Under this setting, many subsequent works (Su et al. 2024; Huang, Zhu, and Chen 2023) design sophisticated training strategies to learn generalized feature representations from limited data. However, these efforts are constrained by the feature representation capabilities of backbone networks (He et al. 2024) pre-trained on ImageNet (Deng et al. 2009). To address these limitations, visual prompt tuning (VPT) (Jia et al. 2022) emerges as a promising solution. VPT introduces learnable prompts into the input space, guiding the model to focus on the most relevant features for the specific task. However, as in Fig. 1 (a), existing VPT methods still face several challenges in CD-FSS. First, they lack prior semantic information during initialization, without any knowledge of the task-specific semantic structures, which can lead to sub-optimal feature adaptation. Moreover, these prompts lack spatial information (Wu et al. 2022) and are restricted to transformer architectures. Given that CD-FSS is a dense and semantically-aware task, directly applying this form of visual prompts is sub-optimal.

To address the aforementioned challenges, we resort to the large vision model, i.e., SAM (Kirillov et al. 2023), for tackling CD-FSS, which has intrinsic cross-domain generalization ability that provides rich prior knowledge for low-data regimes (Liu et al. 2023; Zhang et al. 2023). Specifically, we propose a **Graph Prompt Reasoning Network (GPRN)** which incorporates SAM to generate *semantic-aware* and *spatial-aware* learnable visual prompts to assist effective feature representation learning for CD-FSS. Our visual prompting scheme is shown in Fig. 1 (b), we first transform the masks generated by SAM into visual prompts enriched with prior knowledge. These masks, which segment objects and capture specific semantic information, enrich each prompt with high-level semantic details for the current task. This leads to our **SAM-aware Prompt Initialization module (SPI)**. However, SAM prefers to divide an object into many sub-regions, as seen in the aircraft’s wing and fuselage in Fig. 1 (b). This may lead to two issues: i) prompts belonging to the same category could be inconsistent, and ii) prompts from small masks contain fragmented information. As such, we propose a **Graph Prompt Reasoning module (GPR)** to mine inter-prompt relationships by constructing a graph among these prompts, enabling prompt interactions and ensuring global semantic consistency among prompts. Subsequently, each visual prompt embeds its semantic information into the corresponding mask region to assist in feature representation learning by the proposed reverse masked average pooling operation. Additionally, to address incomplete segmentation and background interference, we introduce an **Adaptive Point Selection module (APS)** that selects positive and negative point prompts from the initial predictions and feeds them back to the SAM for a final prediction. To sum up, our contributions are as follows:

- We propose a SAM-aware graph prompt reasoning network (GPRN), which incorporates SAM to learn visual prompts that integrate both semantic and spatial information to tackle CD-FSS. We further leverage a graph convolutional network (GCN) to reason about the inter-prompt relationships in pursuit of global semantic consistency.

To the best of our knowledge, this is the first work to do this in CD-FSS.

- We propose an APS module that selects positive and negative prompts from initial predictions to address incomplete foreground segmentation and background interference, thereby improving segmentation accuracy.
- Our method achieves state-of-the-art performance on four standard CD-FSS benchmarks.

Related Work

Cross-Domain Few-Shot Segmentation. The CD-FSS setting entails the distribution shifts of data between the base classes during the training phase and the novel classes during the testing phase. Lei et al. (2022) establish a widely used benchmark and design a pyramid transformation module to map features into a domain-invariant space, thus enabling the effective tackling of distribution shifts. Huang et al. (2023) build on their previous work and observe that the intra-domain relationship between the transformed support and query features can be preserved through the use of residual connections. They resort to learning generalized feature representations from limited source data, but they are constrained by the weak feature representation capabilities of the backbone network. Notably, He et al. (2024) develop an end-to-end segmentation framework based on SAM for CD-FSS, but they directly use SAM as a strong feature extractor and focus on learning high-quality prompts for SAM to avoid the instability associated with manual prompts. In contrast, GPRN incorporates SAM to generate visual prompts that assist with feature representation learning and refine coarse predictions under its guidance.

Prompt Learning. Prompt learning is first proposed in the field of natural language processing (NLP) as a strategy to adapt pre-trained language models to downstream tasks (Brown et al. 2020; Lee and Toutanova 2018). This advanced paradigm is also introduced into visual tasks (Du et al. 2022; Sandler et al. 2022). Since manual prompt design requires specialized knowledge, recent works (Jia et al. 2022; Wu et al. 2022) resort to optimizing prompts through gradient back-propagation, known as visual prompt tuning. This approach witnesses tremendous success in few-shot scenarios (Chen et al. 2023; Hossain et al. 2024), as it allows the model to adapt to new tasks. However, compared to textual prompts in NLP, visual prompts inherently lack high-level semantic information. Additionally, most visual prompt tuning methods are restricted to transformer architectures (Jia et al. 2022; Hossain et al. 2024), where visual prompts interact with image tokens to capture their long-range spatial dependencies. To address these challenges, we incorporate SAM to design semantic-aware and spatial-aware visual prompts, aiming to effectively guide the model for CD-FSS tasks.

Graph Convolutional Networks. Graph convolutional networks (GCNs) excel at handling irregular graph-structured data. They represent each input data point as a node and the adjacency relationships between data points as edges. Through information propagation, each node aggregates the features of its neighboring nodes. This enables GCNs to

extract richer and more comprehensive contextual information. GCNs have applications in tasks such as video understanding (Wang and Gupta 2018; Wang et al. 2019b) and zero-shot learning (Xie et al. 2019, 2020). In this work, we use a variant of GCN called graph attention network (GAT) (Velickovic et al. 2017). The key difference between GAT and GCN is that in GCN, the weights of the relationships between neighboring nodes are precomputed, whereas GAT dynamically assigns weights to each neighboring node. This dynamic weighting allows GAT to better handle complex relationships between nodes. In this way, the inter-prompt relationships can be fully captured to produce high-quality visual prompts for CD-FSS tasks.

SAM-Aware Graph Prompt Reasoning Network

Task Definition

The task definition for CD-FSS (Lei et al. 2022) is summarized as follows. We have a source domain (X_s, Y_s) and a target domain (X_t, Y_t) , where X_* represents the data distribution and Y_* denotes the label space. Both the images and the label spaces are disjoint and satisfy $X_s \neq X_t, Y_s \cap Y_t = \emptyset$. In this setup, we construct the training set D_{train} from (X_s, Y_s) and test set D_{test} from (X_t, Y_t) . Each set consists of multiple episodes. Each episode includes a support set $S = \{(I_k^s, M_k^s)\}_{k=1}^K$ and a query set $Q = \{(I^q, M^q)\}$, where I^* and M^* represent the input image and its corresponding ground-truth mask, respectively. Here, K represents the number of support samples. During the training phase, the model is trained on D_{train} to update its parameters. During the test phase, the model’s parameters are frozen to evaluate its performance on D_{test} .

Method Overview

Our proposed GPRN is illustrated in Fig. 2. The overall architecture of GPRN consists of three main modules: SAM-aware prompt initialization (SPI), graph prompt reasoning (GPR), and adaptive point selection (APS). SPI and GPR are plug-and-play and can be used in the training and/or fine-tuning phase, while APS operates exclusively during testing. Since the operations performed on the support and query in the SPI and GPR are identical, for simplicity, we use F (\bar{F} , \hat{F}) to represent the feature maps, omitting the superscripts (s or q), and denote the masks generated by SAM as $\{m_i\}_{i=1}^l$, where l is the number of the masks.

SPI is responsible for generating task-specific visual prompts. It takes as input the features F extracted by CNN and masks $\{m_i\}_{i=1}^l$ generated by SAM, and then produces initialized visual prompts, each of which represents the semantic information of its corresponding masked region. Since SAM tends to segment a single object into multiple sub-regions, there may be inconsistencies in semantics between these sub-regions. To address this, GPR constructs a fully connected graph among the visual prompts, where each visual prompt is treated as a node. Through message passing within the graph, each visual prompt aggregates the semantic information of similar nodes, thereby achieving global

semantic consistency. Then, a reverse masked average pooling (RMAP) operation is employed to restore the spatial information of the visual prompts. This makes the visual prompts have the same spatial resolution as the feature maps, enabling easy interaction with features. Next, the fused feature map \hat{F} along with support mask M^s are fed into SSP (Fan et al. 2022) for prototype-based prediction, resulting in initial query prediction \bar{M}^q . The APS is designed to further enhance the final segmentation results. It selects representative geometry points from the initial prediction’s foreground and background regions as point prompts for SAM. SAM then generates predictions specific to these point prompts. Finally, we fuse the initial segmentation result \bar{M}^q with SAM’s prediction \hat{M}^q to obtain the final prediction \tilde{M}^q . In the following, we specifically explain the proposed modules: SPI, GPR, and APS.

SAM-aware Prompt Initialization

In previous visual prompt tuning methods (Jia et al. 2022; Wu et al. 2022), visual prompts are typically initialized randomly and do not contain prior information, making it difficult for them to adapt to the current task. To address this issue, we convert object masks produced by SAM into visual prompts with prior semantic information. However, although the masks $\{m_i\}_{i=1}^l$ can cover different semantic objects, there are inevitably overlapping areas among them. We need to eliminate these overlapping regions to ensure that the information within each mask remains consistent and unconfused (Li et al. 2024). To resolve this issue, we first sort the masks by area, assigning the smallest index to the mask with the largest area and ensuring that each mask has a unique index. Then, for any overlapping region, we identify the mask with the largest index that includes this region, denoted as j . The overlapping part can be eliminated by:

$$m_i(x, y) = \begin{cases} 1, & \text{if } i = j \\ 0, & \text{if } i < j \end{cases}, \quad (1)$$

where (x, y) denotes the spatial coordinates of overlapping region. Eq. (1) ensures that the overlapping region is assigned to the mask with the largest index (i.e., smallest area), as it has finest information granularity. Meanwhile, it also removes the overlapping region from the larger mask. Next, we use the masked average pooling (MAP) operation to generate the initial visual prompt v_i w.r.t. m_i :

$$v_i = \text{MAP}(F, m_i) \in \mathbb{R}^{1 \times c}, \quad (2)$$

where c is the channel dimension. Visual prompts $\{v_i\}_{i=1}^l$ are obtained by repeatedly applying Eq. (2) to $\{m_i\}_{i=1}^l$.

Graph Prompt Reasoning

Since SAM tends to over-segment images, i.e., dividing a single object into many sub-regions. This leads prompts belonging to the same category to be inconsistent and prompts from small masks to contain fragmented information. Therefore, it is crucial to explore relationships between visual prompts to allow each prompt to aggregate information from other similar prompts. We use a graph attention network (GAT) to model relationships between visual prompts,

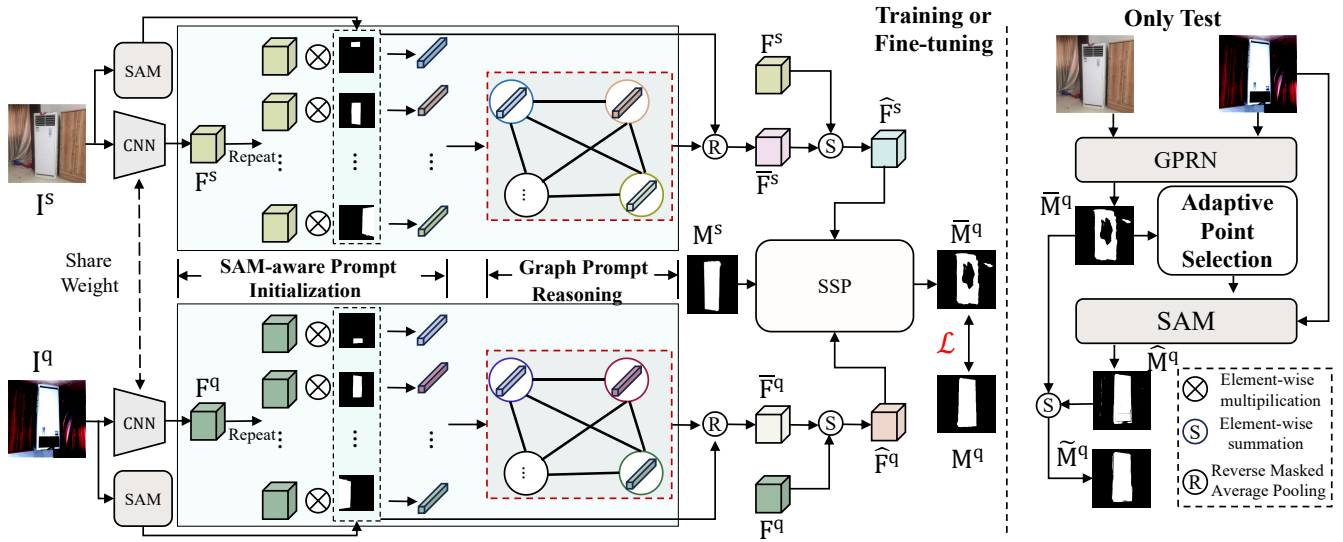


Figure 2: **Overall architecture of our method.** In the training or fine-tuning phase, support and query features along with their corresponding masks generated by SAM are first fed into the SAM-aware initialization module to create visual prompts. These prompts are then processed through a constructed graph to reason about their inter-prompt relationships. Finally, SSP (Fan et al. 2022) is employed to segment the query image. In the testing phase, the proposed adaptive point selection module allows for the generation of more accurate segmentation results.

where each visual prompt is represented as a node in the graph. GAT’s flexible weighting mechanism allows for dynamic adjustment of the relationships between nodes. We first apply learnable matrix $W \in \mathbb{R}^{c \times c}$ to each node to obtain new feature representation, i.e., $\{Wv_i\}_{i=1}^l$. Next, we calculate the importance of node j w.r.t. node i :

$$\phi_{ij} = \frac{d(Wv_i, Wv_j)}{\sum_{j=1}^l d(Wv_i, Wv_j)}, \quad (3)$$

where $d(\cdot, \cdot)$ is the cosine similarity function and ϕ_{ij} is the edge weight between node i and j , with larger values indicating a stronger correlation between the two nodes. Next, GAT aggregates the weighted average of all neighboring nodes of node i and applies a residual connection to form the new feature representation of node i :

$$\bar{v}_i = v_i + \alpha \sum_{j=1}^l \phi_{ij} Wv_j, \quad (4)$$

where α is a scaling parameter. By repeatedly applying Eq. (4), we obtained the desirable visual prompts $\{\bar{v}_i\}_{i=1}^l$ with global semantic consistency.

Self-support Prototype Matching

Inspired by (Li et al. 2024), we first utilize the reverse process of MAP, referred to as RMAP, to restore the spatial information of visual prompts. This allows us to overcome the limitations of the transformer architecture, enabling visual prompts to interact with features without relying on the attention mechanism:

$$\bar{f}_i = \text{RMAP}(\bar{v}_i, m_i) = \frac{\sum_{x,y}^{h,w} \bar{v}_i^\top \times m_i(x,y)}{\sum_{x,y}^{h,w} m_i(x,y)}, \quad (5)$$

where \times denotes matrix multiplication, \bar{v}_i^\top is reshaped by \bar{v}_i , $\bar{f}_i \in \mathbb{R}^{c \times h \times w}$ represents the feature map w.r.t. the i -th visual prompt. Since Eq. (1) has eliminated the overlapping regions, the resulting feature map $\bar{F} = \sum_i \bar{f}_i$ aggregates clear information from all visual prompts. Finally, we directly summarize \bar{F} and F , and the resulting new feature map \hat{F} is used to adapt to the new task. Since the feature representation capability of the CNN is insufficient, \bar{F} provides task-specific feature adaptation. Next, for clarity, we reuse \hat{F}^s and \hat{F}^q to represent new support and query features. SSP (Fan et al. 2022) is a prototype-based method. Its main process involves using support prototypes to estimate the query mask and then using query prototypes to predict query features. Since features of the same object are more similar than those between different objects, it helps mitigate the intra-class variation problem. The above process can be formulated as:

$$P^s = \text{MAP}(\hat{F}^s, M^s) \in \mathbb{R}^{2 \times c}, \quad (6)$$

$$\bar{M}^q = \text{SSP}(\hat{F}^q, P^s) \in \mathbb{R}^{2 \times h \times w}, \quad (7)$$

where P^s is the support prototypes, \bar{M}^q is the query prediction, h and w represent height and width, respectively.

Adaptive Point Selection

After obtaining the initial prediction \bar{M}^q , its accuracy is often suboptimal. Therefore, we aim to refine \bar{M}^q using SAM. Our method involves selecting positive and negative point prompts from both foreground and background mask regions as inputs for the SAM prompt mode. This directs SAM to enhance the segmentation of partially segmented target regions and reduce background interference. Since the

process for selecting prompts from both the foreground and background are identical, we only describe the procedure for choosing positive prompts.

To make the prompts more representative w.r.t. mask regions (i.e., to cover the broadest possible area with the fewest numbers), inspired by maskSLIC (Irving 2016), we start by using a bounding box to enclose the irregular mask, resulting in region R . We denote the set of points within the mask as E and those outside the mask as O , where $R = E \cup O$, and initialize $Z = \{\}$ to store the selected points. For each $e \in E$, its distance transform $DT(e)$ can be formulated as:

$$DT(e) = \min_{o \in O} \left(\sum_{i=1}^n (e_i - o_i)^2 \right)^{\frac{1}{2}}, \quad (8)$$

where, n denotes the feature dimension, and in this case, $n = 2$ because e and o represent two-dimensional coordinates. Next, we apply Gaussian smoothing to $DT(E)$ to reduce the impact of edge noise:

$$DT(E) = \text{Gaussian}(DT(E), \gamma), \quad (9)$$

where γ is a scalar parameter that controls the extent of smoothing, and we set it to 0.1 according to (Irving 2016). The point e that maximizes $DT(e)$ is the current candidate point we have found:

$$e_* = \arg \max_e DT(e). \quad (10)$$

We remove e_* from E , O becomes $O \cup \{e_*\}$, and Z becomes $Z \cup \{e_*\}$ for the next iteration. At the end of the iteration, we obtain the desirable positive point prompts Z . Points located at the edges of the mask exhibit significant instability. Eq. (9) assigns greater weights to points situated near the center, thereby ensuring that the e_* is positioned away from the edges. Simultaneously, Eq.(10) ensures that the candidate points in Z are as spatially distant from each other as possible, thereby maximizing the coverage area of each point. Negative prompts N can be derived from the background mask in \hat{M}^q using Eq. (8) to (10).

Subsequently, we input Z and N into the SAM along with the query image I^q . In prompt mode, the SAM outputs the segmentation logits $\hat{M}_{lgt}^q \in \mathbb{R}^{1 \times h \times w}$. The predicted probability map \hat{M}^q is obtained by using the following formulas:

$$\hat{M}^q = (\mathbf{1} - \sigma(\hat{M}_{lgt}^q)) \oplus \sigma(\hat{M}_{lgt}^q) \in \mathbb{R}^{2 \times h \times w}. \quad (11)$$

Here, σ represents the sigmoid function, and $\mathbf{1}$ denotes a matrix of ones with the same shape as \hat{M}_{lgt}^q . The symbol \oplus indicates concatenation along the channel dimension. Finally, we fuse \bar{M}^q with \hat{M}^q to obtain the final result \tilde{M}^q :

$$\tilde{M}^q = \beta \hat{M}^q + (1 - \beta) \bar{M}^q, \quad (12)$$

where β is a scaling parameter.

Loss Function

Since the APS module is used only during testing, the final prediction of the model during training and/or fine-tuning is \bar{M}^q . Using the binary cross entropy (BCE) function, the model parameters are updated by optimizing the following loss:

$$\mathcal{L} = \text{BCE}(\bar{M}^q, M^q). \quad (13)$$

Experiments

Datasets and Implementation Details

Our model can be optionally trained on the base data, i.e., PASCAL VOC (Everingham et al. 2010) augmented with SDS (Hariharan et al. 2011). We fine-tune and test it on four datasets: Deepglobe (Demir et al. 2018), ISIC (Codella et al. 2019), Chest X-Ray (Candemir et al. 2013), and FSS-1000 (Li et al. 2020). We use mean-IoU as the performance evaluation metric. It is calculated by averaging the intersection-over-unions (IoUs) across different foreground classes.

Our baseline model is SSP (Fan et al. 2022), utilizing the ResNet-50 backbone pre-trained on ImageNet (Deng et al. 2009). All images are initially resized to 400×400 for input into the CNN and subsequently upsampled to 1024×1024 to meet SAM’s input requirements. The number of masks l is fixed as 40. If SAM generates more than 40 masks, we discard those with the smallest areas; if fewer than 40 masks are generated, we pad them with zero masks. The parameter α is set to 0.1. The numbers of positive and negative point prompts, $|Z|$ and $|N|$, are both set to 20, and β is set to 0.5. During training, our model is optimized with 0.9 momentum and an initial learning rate of $1e-3$ and training with 5 epochs. During fine-tuning, we used 120, 520, 60, and 20 samples from the Deepglobe, FSS-1000, ISIC, and Chest X-ray datasets, respectively. We employed the SGD optimizer with a learning rate of $5e-4$, momentum of 0.9, and weight decay of $5e-4$. For evaluation, we randomly sample 600 episodes from ISIC, Deepglobe, and Chest X-Ray, and 2400 episodes from FSS-1000, averaging the results over five random seeds.

Comparisons with State-of-the-arts

Tab. 1 presents a comparison of GPRN with previous state-of-the-art CD-FSS methods under 1-shot and 5-shot settings. To ensure a fair comparison, we categorize the methods based on whether they include a fine-tuning phase, as fine-tuning typically enhances performance. GPRN demonstrates a significant advantage, outperforming the best non-fine-tuning method by 10.4% in the 1-shot setting and 9.9% in the 5-shot setting in terms of the mean-IoU. Even compared to methods that include a fine-tuning phase, GPRN surpasses the top-performing fine-tuning method (IFA) by 3.9% in both 1-shot and 5-shot settings. This establishes GPRN as the new state-of-the-art in various CD-FSS benchmarks. Additionally, when compared to other SAM-based methods such as PerSAM (Zhang et al. 2023) and APSEg (He et al. 2024) for the 1-shot setting, GPRN achieves significant improvements of 34.1% and 10.4%, respectively. These results highlight the effectiveness of the proposed modules in fully leveraging SAM’s potential in CD-FSS.

Our model can be flexibly extended to the K -shot scenario by averaging the prototypes from the support set, while all subsequent steps remain consistent with the 1-shot setting. We find that GPRN achieves good results even when fine-tuning using only a small amount of data, so we report only the fine-tuning results in Tab. 1.

Methods	Backbone	Publication	Deepglobe		ISIC		Chest X-Ray		FSS-1000		Average	
			1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
Methods without Fine-tuning Phase												
AMP (2019)	Vgg-16	ICCV	37.6	40.6	28.4	30.4	51.2	53.0	57.2	59.2	43.6	45.8
RestNet (2023)	ResNet-50	BMVC	-	-	42.3	51.1	70.4	73.7	<u>81.5</u>	<u>84.9</u>	-	-
PerSAM (2023)	ViT-base	ICLR	30.0	30.1	23.3	25.3	61.0	66.5	36.1	40.7	37.6	40.6
ABCDFSS (2024)	ResNet-50	CVPR	42.6	49.0	45.7	53.3	79.8	81.4	74.6	76.2	60.7	65.0
PMNet (2024)	ResNet-50	WACV	37.1	41.6	51.2	54.5	70.4	74.0	84.6	86.3	60.8	64.1
DRANet (2024)	ResNet-50	CVPR	41.3	50.1	40.8	48.8	82.4	82.3	79.1	80.4	60.9	65.4
APSeg (2024)	ViT-base	CVPR	36.0	40.0	45.4	54.0	<u>84.1</u>	84.5	79.7	81.9	61.3	65.1
Methods with Fine-tuning Phase												
PATNet (2022)	ResNet-50	ECCV	37.9	43.0	41.2	53.6	66.6	70.2	78.6	81.2	56.1	62.0
DARNet (2023)	ResNet-50	Arxiv	44.6	54.1	47.8	60.5	81.2	89.7	76.4	83.2	62.5	<u>71.9</u>
IFA (2024)	ResNet-50	CVPR	<u>50.6</u>	<u>58.8</u>	<u>66.3</u>	<u>69.8</u>	74.0	74.6	80.1	82.4	<u>67.8</u>	71.4
DMTNet (2024)	ResNet-50	IJCAI	40.1	51.2	43.6	52.3	73.8	77.0	81.5	83.3	59.7	66.0
GPRN (ours)	ResNet-50	-	51.7	59.3	66.8	72.2	87.0	<u>87.1</u>	81.1	82.6	71.7	75.3

Table 1: Comparison of model performance in 1-shot and 5-shot settings with and without fine-tuning. The best and second-best methods are highlighted in **bold** and underlined, respectively.




Figure 3: Qualitative analysis: I^q and M^q are the query image and ground truth mask, respectively. F^q , \bar{F}^q , and \hat{F}^q denote the feature map from the backbone, visual prompts, and task-adapted feature map. \bar{M}^q , \hat{M}^q , and \tilde{M}^q represent the model's prediction, SAM's prediction, and the final refined segmentation result, respectively.

Ablation Study and Visualizations

Effects of different modules. Tab. 2a presents the impact of GPRN's key components on accuracy for the FSS-1000

dataset. The first row represents our baseline. It can be observed that the SPI module provides a significant mean-IoU improvement of 4.1%. Adding the GPR module further en-

SPI	GPR	APS	1-shot mean-IoU
			73.8
✓			77.9
		✓	77.2
✓	✓		78.4
✓		✓	80.7
✓	✓	✓	81.1

(a)

α	1-shot mean-IoU
0.1	81.1
0.3	80.9
0.5	80.5

(b)

Number	1-shot mean-IoU
10	76.2
20	78.4
40	81.1
60	81.3

(c)

$ Z $	$ N $	1-shot mean-IoU
10	0	79.2
10	10	80.5
20	20	81.1
30	30	81.0

(d)

Strategy	1-shot mean-IoU
(Zhang et al. 2023)	77.9
Box	75.3
APS (ours)	81.1
APS + Box	81.1

(e)

β	1-shot mean-IoU
0.1	78.5
0.3	80.2
0.5	81.1
0.6	80.2

(f)

Table 2: **Ablation study on FSS-1000.** (a) Effects of different modules. (b) Effects of different α . (c) Effects of number of masks l . (d) Effects of number of point prompts $|Z|$ and $|N|$. (e) Effects of different prompts selection strategies. (f) Effects of different β .

hances performance, achieving a mean-IoU of 78.4%. This improvement is attributed to the GPR module’s ability to enhance global semantic consistency among prompts. Additionally, the parameter-free APS module contributes a 3.4% mean-IoU increase. Combining all three modules results in a mean-IoU of 81.1% (+7.3%). This indicates that the three modules are complementary. It can be interpreted that SPI and GPR enhance the accuracy of the initial predictions, ensuring that APS selects more representative point prompts, thereby further improving segmentation accuracy. **Effects of different α .** We set α to 0.1, 0.3, and 0.5 to investigate its impact on the results. As shown in Tab. 2b, the model achieves the best performance when α is 0.1.

Effects of number of masks l . The number of masks l directly affects the model’s performance. When l is small, the masks are coarse-grained, while increasing l results in finer-grained masks. We set l to 10, 20, 40, and 60, with corresponding results shown in Tab. 2c. When $l = 40$ and $l = 60$, the model’s mean-IoU accuracy is similar. Considering both performance and computational efficiency, we set l to 40.

Effects of number of point prompts $|Z|$ and $|N|$. The number of point prompts affects both performance and computational efficiency. Comparing the first and second rows in Tab. 2d shows that selecting negative point prompts from the background also improves performance. When the number of point prompts from both the foreground and background is set to 20, the model’s performance is saturated. To reduce computational burden, we ultimately select 20 point prompts from both the foreground and background.

Effects of different prompts selection strategies. To further validate the necessity of the point prompt selection strategy in APS, in Tab. 2e, we compare the APS with the selection principle proposed in PerSAM (Zhang et al. 2023), which selects the points with the highest foreground or background probabilities from the prediction results as candidate point prompts, and the prompt form that directly specifies the target location using a box. We similarly select 20 points

for both the foreground and background. Our strategy outperforms the one proposed by PerSAM by a clear margin of 3.2%. We also observe that using a box along with APS does not bring any performance improvement. So our final prompt solution is to use APS only.

Effects of different β . The selection of the β determines the balance between incorporating correction information from SAM and the original prediction data. We conduct ablation experiments with β of 0.1, 0.3, 0.5, and 0.6 in Tab. 2f. The model achieves the best performance when the β value is set to 0.5. As the beta value increases from 0.1 to 0.5, the model’s performance improves significantly. However, further increasing the beta value to 0.6 results in a decline in performance. This indicates that there is a complementary relationship between the model’s predictions and SAM’s predictions.

Visualizations Fig.3 provides some visualizations. It can be found that the feature map F^q extracted by the backbone contains significant aliasing artifacts, while the visual prompts \bar{F}^q exhibit clear boundaries. The combined feature map \hat{F}^q which incorporates information from \bar{F}^q significantly reduces the impact of noise. The initial prediction \bar{M}^q is often inaccurate, and SAM’s prediction \hat{M}^q also contains many artifacts. However, \tilde{M}^q , which fuses both predictions, achieves the best segmentation performance.

Conclusion

In this work, we propose a SAM-aware graph prompt network (GPRN) to tackle the CD-FSS task. Specifically, we transform the masks generated by SAM into several visual prompts, followed by a graph attention network to reason about inter-prompt relationships. We also design a novel parameter-free test-time refine module, which can enhance the model’s performance with a clear margin. GPRN achieves state-of-the-art performance on various CD-FSS datasets.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (Nos. 62276134, 62476124, and 62302093).

References

- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Candemir, S.; Jaeger, S.; Palaniappan, K.; Musco, J. P.; Singh, R. K.; Xue, Z.; Karargyris, A.; Antani, S.; Thoma, G.; and McDonald, C. J. 2013. Lung segmentation in chest radiographs using anatomical atlases with nonrigid registration. *IEEE transactions on medical imaging*, 33(2): 577–590.
- Chen, H.; Dong, Y.; Lu, Z.; Yu, Y.; and Han, J. 2024. Pixel matching network for cross-domain few-shot segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 978–987.
- Chen, J.; Quan, R.; and Qin, J. 2024. Cross-Domain Few-Shot Semantic Segmentation via Doubly Matching Transformation. In Larson, K., ed., *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, 641–649. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Chen, W.; Si, C.; Zhang, Z.; Wang, L.; Wang, Z.; and Tan, T. 2023. Semantic prompt for few-shot image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23581–23591.
- Codella, N.; Rotemberg, V.; Tschandl, P.; Celebi, M. E.; Dusza, S.; Gutman, D.; Helba, B.; Kalloo, A.; Liopyris, K.; Marchetti, M.; et al. 2019. Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic). *arXiv preprint arXiv:1902.03368*.
- Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; and Schiele, B. 2016. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3213–3223.
- Demir, I.; Koperski, K.; Lindenbaum, D.; Pang, G.; Huang, J.; Basu, S.; Hughes, F.; Tuia, D.; and Raskar, R. D. 2018. A challenge to parse the earth through satellite images. *arXiv preprint arXiv:1805.06561*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Du, Y.; Wei, F.; Zhang, Z.; Shi, M.; Gao, Y.; and Li, G. 2022. Learning to prompt for open-vocabulary object detection with vision-language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14084–14093.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88: 303–338.
- Fan, H.; Fan, Q.; Pagnucco, M.; and Song, Y. 2023. DAR-Net: Bridging Domain Gaps in Cross-Domain Few-Shot Segmentation with Dynamic Adaptation. *arXiv preprint arXiv:2312.04813*.
- Fan, Q.; Pei, W.; Tai, Y.-W.; and Tang, C.-K. 2022. Self-support few-shot semantic segmentation. In *European Conference on Computer Vision*, 701–719. Springer.
- Hariharan, B.; Arbeláez, P.; Bourdev, L.; Maji, S.; and Malik, J. 2011. Semantic contours from inverse detectors. In *2011 international conference on computer vision*, 991–998. IEEE.
- He, W.; Zhang, Y.; Zhuo, W.; Shen, L.; Yang, J.; Deng, S.; and Sun, L. 2024. APSeg: Auto-Prompt Network for Cross-Domain Few-Shot Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23762–23772.
- Herzog, J. 2024. Adapt Before Comparison: A New Perspective on Cross-Domain Few-Shot Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23605–23615.
- Hong, S.; Cho, S.; Nam, J.; Lin, S.; and Kim, S. 2022. Cost aggregation with 4d convolutional swin transformer for few-shot segmentation. In *European Conference on Computer Vision*, 108–126. Springer.
- Hossain, M. R. I.; Siam, M.; Sigal, L.; and Little, J. J. 2024. Visual Prompting for Generalized Few-shot Segmentation: A Multi-scale Approach. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23470–23480.
- Huang, X.; Zhu, C.; and Chen, W. 2023. Restnet: Boosting cross-domain few-shot segmentation with residual transformation network. *arXiv preprint arXiv:2308.13469*.
- Irving, B. 2016. maskSLIC: regional superpixel generation with application to local pathology characterisation in medical images. *arXiv preprint arXiv:1606.09518*.
- Jia, M.; Tang, L.; Chen, B.-C.; Cardie, C.; Belongie, S.; Hariharan, B.; and Lim, S.-N. 2022. Visual prompt tuning. In *European Conference on Computer Vision*, 709–727. Springer.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment Anything. *arXiv e-prints*, arXiv–2304.
- Lee, J.; and Toutanova, K. 2018. Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 3(8).
- Lei, S.; Zhang, X.; He, J.; Chen, F.; Du, B.; and Lu, C.-T. 2022. Cross-domain few-shot semantic segmentation. In *European Conference on Computer Vision*, 73–90. Springer.
- Li, S.; Liu, F.; Jiao, L.; Liu, X.; Chen, P.; and Li, L. 2024. Mask-guided correlation learning for few-shot segmentation in remote sensing imagery. *IEEE Transactions on Geoscience and Remote Sensing*.

- Li, X.; Wei, T.; Chen, Y. P.; Tai, Y.-W.; and Tang, C.-K. 2020. Fss-1000: A 1000-class dataset for few-shot segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2869–2878.
- Li, Y.; Zeng, J.; and Shan, S. 2020. Learning representations for facial actions from unlabeled videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1): 302–317.
- Li, Y.; Zeng, J.; Shan, S.; and Chen, X. 2018. Occlusion aware facial expression recognition using CNN with attention mechanism. *IEEE Transactions on Image Processing*, 28(5): 2439–2450.
- Liu, J.; Bao, Y.; Xie, G.-S.; Xiong, H.; Sonke, J.-J.; and Gavves, E. 2022. Dynamic prototype convolution network for few-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11553–11562.
- Liu, Y.; Zhu, M.; Li, H.; Chen, H.; Wang, X.; and Shen, C. 2023. Matcher: Segment anything with one shot using all-purpose feature matching. *arXiv preprint arXiv:2305.13310*.
- Min, J.; Kang, D.; and Cho, M. 2021. Hypercorrelation squeeze for few-shot segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6941–6952.
- Moon, S.; Sohn, S. S.; Zhou, H.; Yoon, S.; Pavlovic, V.; Khan, M. H.; and Kapadia, M. 2023. Msi: Maximize support-set information for few-shot segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19266–19276.
- Nie, J.; Xing, Y.; Zhang, G.; Yan, P.; Xiao, A.; Tan, Y.-P.; Kot, A. C.; and Lu, S. 2024. Cross-Domain Few-Shot Segmentation via Iterative Support-Query Correspondence Mining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3380–3390.
- Peng, B.; Tian, Z.; Wu, X.; Wang, C.; Liu, S.; Su, J.; and Jia, J. 2023. Hierarchical dense correlation distillation for few-shot segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23641–23651.
- Peng, S.-F.; Xie, G.-S.; Zhao, F.; Shu, X.; and Liu, Q. 2024. Multi-Granularity Aggregation Network for Remote Sensing Few-Shot Segmentation. *IEEE Transactions on Geoscience and Remote Sensing*.
- Sandler, M.; Zhmoginov, A.; Vladymyrov, M.; and Jackson, A. 2022. Fine-tuning image transformers using learnable memory. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12155–12164.
- Shi, X.; Wei, D.; Zhang, Y.; Lu, D.; Ning, M.; Chen, J.; Ma, K.; and Zheng, Y. 2022. Dense cross-query-and-support attention weighted mask aggregation for few-shot segmentation. In *European Conference on Computer Vision*, 151–168. Springer.
- Siam, M.; Oreshkin, B. N.; and Jagersand, M. 2019. Amp: Adaptive masked proxies for few-shot segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5249–5258.
- Su, J.; Fan, Q.; Pei, W.; Lu, G.; and Chen, F. 2024. Domain-Rectifying Adapter for Cross-Domain Few-Shot Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24036–24045.
- Sun, G.; Liu, Y.; Ding, H.; Probst, T.; and Van Gool, L. 2022. Coarse-to-fine feature mining for video semantic segmentation. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3126–3137.
- Tian, Z.; Zhao, H.; Shu, M.; Yang, Z.; Li, R.; and Jia, J. 2020. Prior guided feature enrichment network for few-shot segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 44(2): 1050–1065.
- Velickovic, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; Bengio, Y.; et al. 2017. Graph attention networks. *stat*, 1050(20): 10–48550.
- Wang, K.; Liew, J. H.; Zou, Y.; Zhou, D.; and Feng, J. 2019a. Panet: Few-shot image semantic segmentation with prototype alignment. In *proceedings of the IEEE/CVF international conference on computer vision*, 9197–9206.
- Wang, W.; Lu, X.; Shen, J.; Crandall, D. J.; and Shao, L. 2019b. Zero-shot video object segmentation via attentive graph neural networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9236–9245.
- Wang, X.; and Gupta, A. 2018. Videos as space-time region graphs. In *Proceedings of the European conference on computer vision (ECCV)*, 399–417.
- Wu, J.; Li, X.; Wei, C.; Wang, H.; Yuille, A.; Zhou, Y.; and Xie, C. 2022. Unleashing the power of visual prompting at the pixel level. *arXiv preprint arXiv:2212.10556*.
- Xie, G.-S.; Liu, J.; Xiong, H.; and Shao, L. 2021. Scale-aware graph neural network for few-shot semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5475–5484.
- Xie, G.-S.; Liu, L.; Jin, X.; Zhu, F.; Zhang, Z.; Qin, J.; Yao, Y.; and Shao, L. 2019. Attentive region embedding network for zero-shot learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9384–9393.
- Xie, G.-S.; Liu, L.; Zhu, F.; Zhao, F.; Zhang, Z.; Yao, Y.; Qin, J.; and Shao, L. 2020. Region graph embedding network for zero-shot learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, 562–580. Springer.
- Yang, B.; Liu, C.; Li, B.; Jiao, J.; and Ye, Q. 2020. Prototype mixture models for few-shot semantic segmentation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VIII 16*, 763–778. Springer.
- Zhang, R.; Jiang, Z.; Guo, Z.; Yan, S.; Pan, J.; Ma, X.; Dong, H.; Gao, P.; and Li, H. 2023. Personalize segment anything model with one shot. *arXiv preprint arXiv:2305.03048*.
- Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 633–641.