

Advancing Medical Image Segmentation via Self-supervised Instance-adaptive Prototype Learning

Guoyan Liang¹, Qin Zhou², Jingyuan Chen¹, Zhe Wang^{2*} and Chang Yao^{1*}

¹Zhejiang University, Hangzhou, China

²East China University of Science and Technology, Shanghai, China

{guoyanl, jingyuanchen, changy}@zju.edu.cn, {sunniezq, wangzhe}@ecust.edu.cn

Abstract

Medical Image Segmentation (MIS) plays a crucial role in medical therapy planning and robot navigation. Prototype learning methods in MIS focus on generating segmentation masks through pixel-to-prototype comparison. However, current approaches often overlook sample diversity by using a fixed prototype per semantic class and neglect intra-class variation within each input. In this paper, we propose to generate instance-adaptive prototypes for MIS, which integrates a common prototype proposal (CPP) capturing common visual patterns and an instance-specific prototype proposal (IPP) tailored to each input. To further account for the intra-class variation, we propose to guide the IPP generation by re-weighting the intermediate feature map according to their confidence scores. These confidence scores are hierarchically generated using a transformer decoder. Additionally we introduce a novel self-supervised filtering strategy to prioritize the foreground pixels during the training of the transformer decoder. Extensive experiments demonstrate favorable performance of our method.

1 Introduction

Medical image segmentation (MIS) aims to divide an image into meaningful parts, representing different semantic classes. It serves as a fundamental requirement for various tasks such as cancer diagnosis, radiation therapy dosage control and surgical robot navigation. Recognizing the significant resemblance between MIS and the task of clustering pixels into semantic segments, recent researchers have focused on prototype-based methods. By learning representative prototype features, which serve as cluster centers for each semantic class, the MIS objective is reformulated as pixel-to-prototype matching. Existing prototype-based MIS methods primarily concentrate on learning representative prototypes, which remain fixed upon completion of the training process [Zhou *et al.*, 2022]. In order to capture the

intra-class variation within each semantic class, prior studies have explored the use of local prototypes [Yu *et al.*, 2021; Wang *et al.*, 2022]. However, these prototypes are predominantly generated based on spatial grids or pre-segmentation techniques such as superpixels. Recently, the introduction of mask transformers [Cheng *et al.*, 2022; Yu *et al.*, 2022; Yuan *et al.*, 2023] has allowed for the direct learning of object queries using a transformer decoder. By considering these object queries as cluster centers, we can obtain the segmentation mask by assigning pixels to clusters and mapping clusters to their corresponding semantic classes. However, as mentioned in [Cheng *et al.*, 2022; Yu *et al.*, 2022; Yuan *et al.*, 2023], without proper guidance, existing mask transformers may introduce a bias towards focusing on the background.

To tackle the aforementioned challenges in prototype-based segmentation, we introduce a novel framework called Self-supervised Instance-adaptive Prototype Learning (referred as SIPL) for MIS. Specifically, we propose an Instance-adaptive Prototype Learning (referred as IPL) scheme that adaptively learns prototypes tailored to different inputs. In IPL, the final class-specific prototypes are generated using a combination of common prototype proposals (CPPs), which capture the common visual patterns of each class, and instance-specific prototype proposals (IPPs), which dynamically emphasize the instance-specific visual characteristics of the current input.

To further account for the intra-class variation during prototype generation, we propose a novel hierarchical Self-supervised Mask Generation (referred as SMG) module to hierarchically generate pseudo masks with confidence scores to guide the generation of IPPs. These confidence scores are used to re-weight the intermediate feature map, indicating the pixel-wise importance in contributing to the IPPs for each class. In SMG, a transformer decoder is employed to generate query embeddings from the multi-scale features obtained from the pixel decoder. To mitigate background bias in the SMG module, we incorporate a self-supervised filtering strategy that prioritizes foreground pixels during the learning of query embeddings. To improve the accuracy of the confidence scores, We further introduce an auxiliary loss to supervise the training process of the transformer decoder.

The overall contributions of our method can be summarized as follows.

*Corresponding Authors.

- We propose a novel framework to learn instance-adaptive prototypes for MIS, which incorporates common and instance-specific prototype proposals to capture the common and instance-specific visual patterns for each class.
- We propose a self-supervised mask generation module to guide instance-specific prototype proposal generation by re-weighting intermediate features, where we incorporate a novel self-supervised filtering strategy to prioritize foreground pixels during pseudo mask generation.
- Extensive experiments across three challenging MIS tasks demonstrate favorable performance of our method.

2 Related Works

2.1 Medical Image Segmentation

Early methods in medical image segmentation (MIS) primarily rely on contour-based and traditional machine learning algorithms [Tsai *et al.*, 2003]. In contrast, contemporary MIS approaches predominantly leverage deep CNNs and transformer models. U-Net initially introduced for MIS based on CNN, stands out for its simplicity and superior performance [Ronneberger *et al.*, 2015]. Subsequently, various U-Net variants have emerged for 2D MIS such as ResUNet [Xiao *et al.*, 2018], U-Net++ [Zhou *et al.*, 2018], and UNet3+ [Huang *et al.*, 2020]. Transformer-based methods have emerged due to the impressive speed-accuracy performance of the Vision Transformer (ViT) in various visual tasks [Dosovitskiy *et al.*, 2020]. Swin Transformer has notably excelled by introducing an efficient and effective hierarchical vision transformer that leverages the shifted windows mechanism [Liu *et al.*, 2021]. Building upon this paradigm, Swin-unet incorporates a U-shaped Encoder-Decoder architecture with skip connections, further advancing the field [Cao *et al.*, 2022]. Segtran proposes a novel positional encoding scheme for transformers, which imposes a continuity inductive bias for images and achieves high segmentation accuracy with good cross-domain generalization capabilities [Li *et al.*, 2021]. Some recent works explore hybrid approaches that demonstrate improved segmentation performance compared to pure CNNs and transformers-based counterparts. UNETR seamlessly integrates the long-range spatial dependencies of transformers with the inductive bias of CNNs within a “U-shaped” encoder-decoder architecture [Hatamizadeh *et al.*, 2022]. nnFormer strategically combines CNN and transformer-based blocks, interleaving them to harness the complementary strengths of each [Zhou *et al.*, 2021]. UNETR++ introduces an innovative Efficient Paired Attention (EPA) block, skillfully capturing enriched interdependent spatial and channel features by employing both spatial and channel attention in two distinct branches [Abdelrahman Shaker *et al.*, 2022].

2.2 Prototype Based Methods

Prototype learning assigns query embeddings to class centers in the network’s output space, typically defined as the mean vector of the training set for each class [Mensink *et al.*, 2013]. Prototype-based networks exhibit a straightforward output layout [Wen *et al.*, 2016] and demonstrate fast

generalization to new classes. As a result, they often concentrate on the few-shot regime, [Boney and Ilin, 2017; Snell *et al.*, 2017] or update the prototypes online based on mini-batches [Guerriero *et al.*, 2018]. Matching Network is proposed to leverage learned prototypes for one-shot learning, capturing class features during training and demonstrating robust classification during testing [Vinyals *et al.*, 2016]. Prototypical Network utilizes prototype spaces to compute distances between input samples and class prototypes for few-shot learning, proving effective in scenarios with limited labeled data. [Santoro *et al.*, 2016] incorporates memory augmentation to store task-specific prototype information, facilitating rapid adaptation to new tasks. Relation Networks [Sung *et al.*, 2018] emphasizes modeling relationships between input samples and class prototypes for enhanced few-shot learning performance. Additionally, [Chen *et al.*, 2019] refines prototype-based methods and explore improved strategies to leverage limited training data for more accurate classification. While existing prototype-based methods have made significant progress, they often lack the flexibility to dynamically adapt to different inputs [Zhou *et al.*, 2022]. In contrast, our method focuses on learning instance-adaptive prototypes, resulting in improved performance in MIS.

3 Method

In this section, we will provide an overview of our proposed method. We start by discussing the background and limitations of existing prototype-based segmentation methods in Sec. 3.1. Then, we present the Self-supervised Instance-adaptive Prototype Learning (SIPL) framework in Sec. 3.2. We further explain the Instance-adaptive Prototype Learning (IPL) strategy in Sec. 3.3 and the hierarchical Self-supervised Mask Generation (SMG) module in Sec. 3.4.

3.1 Revisiting Prototype Based Segmentation

Prototype learning methods in Medical Image Segmentation (MIS) focus on generating segmentation masks through pixel-to-prototype comparison. These approaches typically learn one prototype feature for each class. Denote the learned prototypes as $\mathbf{G} = \{\mathbf{g}_k \in \mathbb{R}^d\}, k \in \{1, \dots, K\}$, the i -th pixel embedding as $\mathbf{f}_i \in \mathbb{R}^d$, where K is the number of classes and d is the feature dimension. Then the probability that the i -th pixel belongs to the k -th class can be calculated as follows:

$$p(k|\mathbf{f}_i) = \frac{\exp(-\langle \mathbf{f}_i, \mathbf{g}_k \rangle)}{\sum_{k'=1}^K \exp(-\langle \mathbf{f}_i, \mathbf{g}_{k'} \rangle)}, \quad (1)$$

where \langle, \rangle denotes the cosine similarity between two items. The existing prototype-based segmentation methods encounter challenges in two aspects: 1) *Fixed prototypes during inference that lack dynamic adaptation to different inputs.* 2) *Prototypes generated by averaging feature embeddings of each class, overlooking intra-class variation.* To overcome these limitations, we propose a novel framework that generates instance-adaptive prototypes for each input. In addition, we address intra-class variation by hierarchically generating pseudo masks with confidence scores, which serve as the pixel-wise importance in the instance-specific prototype




Figure 1: Overview of our proposed SIPL framework.

proposal generation for each class. The subsequent section provides a detailed description of the network architecture.

3.2 Overall Architecture

Figure 1 illustrates the overall architecture of our method. Given the input 3D volume-mask pair (\mathbf{V}, Y) , where $\mathbf{V} \in \mathbb{R}^{H \times W \times Z}$ is the 3D image patch, and $Y \in \mathbb{R}^{H \times W \times Z}$ refers to the corresponding segmentation mask. We utilize a 3D backbone comprising an encoder branch $\psi_e(\cdot)$ and a pixel decoder branch $\psi_d(\cdot)$ for segmentation mask prediction. Concretely, the input \mathbf{V} is first fed into the encoder to obtain the intermediate feature map $\mathbf{F} \in \mathbb{R}^{\frac{H}{S} \times \frac{W}{S} \times \frac{Z}{S} \times d_i}$, where S is the stride of the feature map, d_i is the feature dimension of the intermediate feature map. Then \mathbf{F} is fed into the pixel decoder branch $\psi_d(\cdot)$ to generate the output per-pixel embeddings $\mathbf{F}_o \in \mathbb{R}^{H \times W \times Z \times d}$. The final segmentation mask is predicated as the pixel-to-prototype comparison between \mathbf{F}_o and the learned prototypes \mathbf{G} .

To facilitate dynamic adaptation of learned prototypes to current input, we introduce the Instance-adaptive Prototype Learning (IPL) strategy, which combines Common Prototype Proposals (CPPs) and Instance-specific Prototype Proposals (IPPs) generation. The CPPs are learnable parameters obtained from the entire dataset, while IPPs are generated by reweighting the intermediate feature map \mathbf{F} using the pixel-wise confidence scores. The final prototype proposals $\mathbf{P} = \{\mathbf{p}_k\}, k \in \{1, \dots, K\}$ are obtained by concatenating the CPPs and IPPs along the feature channel dimension. These proposals are then processed through a few fully connected layers (MLP) to obtain the final prototypes \mathbf{G} .

In order to account for the intra-class variation during the generation of IPPs, we introduce the Self-supervised Mask Generation (SMG) module to generate pixel-wise confidence scores for each class. This module utilizes a mask transformer to generate pseudo masks with confidence scores, where the pseudo masks are hierarchically aggregated from pixel-to-cluster and cluster-to-class assignments. To provide the IPL

module with more accurate pseudo masks, we introduce the auxiliary loss (L_{aux}) between pseudo masks M and ground truth masks Y .

3.3 Instance-adaptive Prototype Learning

The overall workflow of IPL is illustrated in Figure 2. Formally, given the intermediate feature map $\mathbf{F} \in \mathbb{R}^{\frac{H}{S} \times \frac{W}{S} \times \frac{Z}{S} \times d_i}$ and pseudo masks from the final transformer decoder layer $M = \{M_k \in \mathbb{R}^{H \times W \times Z \times 1}, k \in \{1, 2, \dots, K+1\}\}$, where $M_k \in [0, 1]$ denotes the predicted confidence scores of the k -th class. Then the instance-specific prototype proposal (IPP) for the k -th class $\mathbf{p}_k^i \in \mathbb{R}^{1 \times d_i}$ is generated as,

$$\mathbf{p}_k^i = GAP(M_k \otimes \mathbf{F}), \quad (2)$$

where \otimes is the pixel-wise multiplication, $GAP()$ denotes global average pooling. In our method, pseudo masks M are generated hierarchically from the self-supervised mask generation module (please refer to Eq. (9) for details). Please note that the pseudo mask M_k are reshaped to match the size of \mathbf{F} in Eq. (2).

In addition to the IPPs, our method includes a learnable common prototype proposal (CPP) for each class. This shared proposal captures the common visual dynamics of each class present in the entire training set. Concretely, denote the CPPs as $\mathbf{P}^c = \{\mathbf{p}_k^c\}, k \in \{1, 2, \dots, K\}$, where $\mathbf{p}_k^c \in \mathbb{R}^{1 \times d_i}$ denotes the prototype proposal of the k -th class. These CPPs are optimized by the overall learning objective and are randomly initialized in our approach. The IPPs and CPPs are then fused together to generate the final prototype \mathbf{G} . Formally, the final prototype belonging to the k -th class is obtained as,

$$\mathbf{g}_k = MLP(\mathbf{p}_k^i \oplus \mathbf{p}_k^c), \quad (3)$$

where \oplus denotes the concatenation operation, $MLP()$ refers to a few fully connected layers. Given the output per-pixel embeddings $\mathbf{F}_o \in \mathbb{R}^{H \times W \times Z \times d}$ from the pixel decoder

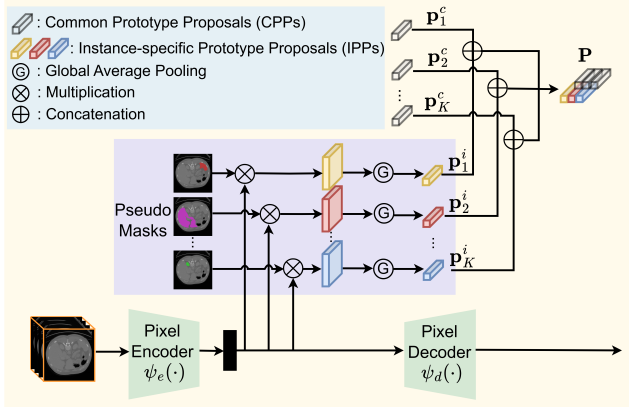


Figure 2: Illustration of our instance-adaptive prototype learning (IPL) module. The IPL module consists of two components: common prototype proposal generation and instance-specific prototype proposal generation. The pseudo masks with confidence scores are hierarchically generated in a self-supervised manner to account for the intra-class variation, as described in Sec 3.4.

branch and the learned prototype $\mathbf{g}_k \in \mathbb{R}^{d \times 1}$, the final segmentation prediction for the k -th class is formulated as,

$$\widehat{Y}_k = \sigma(\mathbf{F}_o \otimes \mathbf{g}_k), \quad (4)$$

where \otimes denotes the matrix multiplication, $\sigma(\cdot)$ denotes the Sigmoid operation. We adopt soft dice loss and binary cross-entropy to calculate the segmentation loss, which is mathematically formulated as,

$$L_{seg} = 1 - \sum_{k=1}^K \left(\sum_i \frac{Y_{k,i} \widehat{Y}_{k,i}}{Y_{k,i}^2 + \widehat{Y}_{k,i}^2} + \sum_i Y_{k,i} \log \widehat{Y}_{k,i} \right). \quad (5)$$

Here, $Y_{k,i}, \widehat{Y}_{k,i}$ correspond to the i -th pixel of the segmentation map for class k in the ground truth and predictions, respectively.

3.4 Self-supervised Mask Generation

Starting from DETR [Carion *et al.*, 2020] and MaX-DeepLab [Wang *et al.*, 2021a], the object queries, updated by multiple transformer decoders, are employed as mask embedding vectors. In our approach, we employ learnable object queries to capture intra-class variation, as shown in Figure 3. These object queries, which can be seen as clusters, are mapped onto the semantic class to generate pseudo masks with confidence scores. This process facilitates instance-adaptive prototype learning. Mathematically, denote the learnable object queries as $\mathbf{Q} \in \mathbb{R}^{N \times d_q}$, where N ($N > K$) refers to the number of object queries, d_q denotes the feature dimension of the object queries. Then \mathbf{Q} is updated by kMax decoders [Yu *et al.*, 2022] using the multi-scale features from the pixel decoder. Denote the reshaped feature map from the l -th layer of the pixel decoder as $\mathbf{F}_d^l \in \mathbb{R}^{H_l \times W_l \times Z_l \times d_q}$, the randomly initialized query embeddings as \mathbf{Q}^0 , then the updated query embeddings \mathbf{Q}^l is calculated as,

$$\begin{aligned} \mathbf{A} &= \underset{N}{\operatorname{argmax}}(\mathbf{Q}^{l-1} \otimes (\mathbf{F}_d^l)^T), \\ \mathbf{Q}^l &= \operatorname{MHSA}(\mathbf{Q}^{l-1}) + \mathbf{A} \otimes \mathbf{F}_d^l, \end{aligned} \quad (6)$$

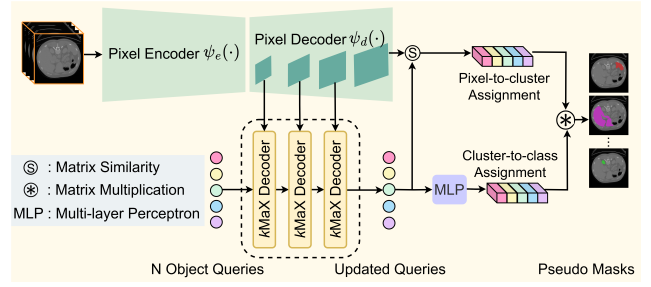


Figure 3: Illustration of the self-supervised mask generation (SMG) module. In SMG, object queries are learned from a transformer decoder using multi-scale features from the pixel decoder. Pseudo masks are generated by hierarchically assigning pixels to clusters and clusters to classes.

where \mathbf{Q}^{l-1} denotes the query embeddings obtained from the $l-1$ layer, $\operatorname{MHSA}(\cdot)$ denotes the multi-head self-attention. Based on the updated query embeddings \mathbf{Q}^l from layer l , we can obtain the corresponding pixel-to-cluster assignment map M_{pc}^l as,

$$M_{pc}^l = \operatorname{sim}(\mathbf{F}_d^l, \mathbf{Q}^l), \quad (7)$$

where $\operatorname{sim}(\cdot, \cdot)$ calculates the feature similarity between each pixel of \mathbf{F}_d^l and each query embedding of \mathbf{Q}^l . The generated pixel-to-cluster assignment map M_{pc}^l is of size $H_l \times W_l \times Z_l \times N$. Since the IPL module requires class-specific pseudo masks, we further introduce a cluster-to-class projection head to generate cluster-to-class mapping matrix $M_{cc}^l \in \mathbb{R}^{N \times (K+1)}$,

$$M_{cc}^l = \operatorname{MLP}(\mathbf{Q}^l). \quad (8)$$

We then aggregate the pixel-to-cluster assignments M_{pc}^l and the cluster-to-class mapping matrix M_{cc}^l to obtain the pseudo masks for the l -th layer as,

$$M^l = M_{pc}^l \otimes M_{cc}^l. \quad (9)$$

M^l undergoes softmax processing to get the normalized confidence scores. It should be noted that during the generation of M^l , the background class (i.e., $k = K+1$) is considered to obtain reasonable confidence scores using softmax.

Self-supervised Filtering. As mentioned in [Yuan *et al.*, 2023], generating object queries without proper guidance can introduce a bias towards focusing on the background. To mitigate this issue and prioritize foreground features during training, we introduce a novel self-supervised filtering strategy. This strategy involves comparing the overlapping ratios between the activated regions of each object cluster and the foreground regions of each class. The comparison is performed between the pseudo masks generated by the previous layer M^{l-1} and the pixel-to-cluster assignments at current layer M_{pc}^l , which is formulated as,

$$\operatorname{Overlap}(M^{l-1}, M_{pc}^l) = \underset{k}{\operatorname{argmax}} \frac{\sum (M_k^{l-1} \cap M_{pc}^l)}{\sum M_{pc}^l}, \quad (10)$$

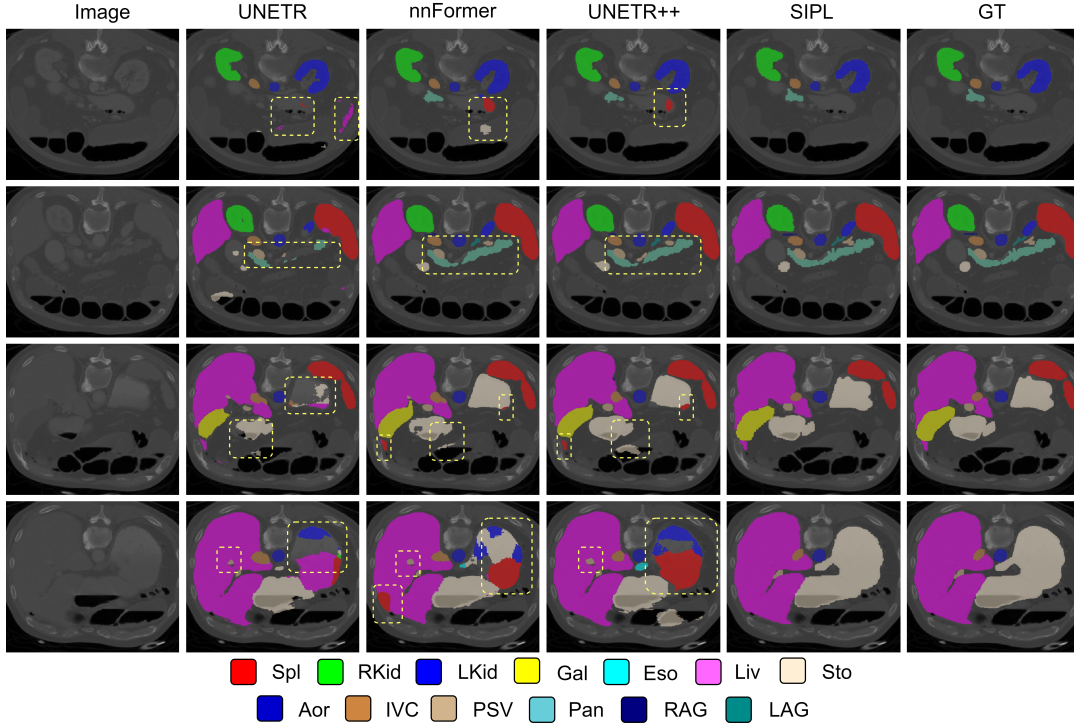


Figure 4: Qualitative comparison between SIPL and SOTA models on BTCV dataset. Current methods face challenges in accurately segmenting various organs (highlighted within yellow dashed boxes). In contrast, our proposed framework exhibits commendable segmentation performance.

where $M_{pc}^{l,n}$ denotes the binarized mask prediction of the n -th cluster in M_{pc}^l , M_k^{l-1} refers to the binarized mask prediction of the k -th semantic class in M^{l-1} . The numerator of $Overlap(\cdot)$ calculates the intersection between the mask of the current cluster and the masks of each semantic class, and $Overlap(\cdot)$ returns the maximum ratio between the cluster-to-class intersection and the area of the related cluster. The underlying assumption is that each cluster represents a specific part of a particular semantic class. Therefore, only clusters with an overlapping ratio higher than the threshold τ are chosen for query embedding update. In our method, the value of τ is progressively increased from 0.1 to 0.5 in a linear manner and remains constant after epoch 50, as shown below,

$$\tau = \min\left(\frac{epoch}{50}, 1.0\right) * 0.4 + 0.1, \quad (11)$$

where $epoch$ denotes the index of current training epoch, $\min(\cdot)$ returns the smaller value. In order to enhance the quality of the pseudo masks, we introduce an auxiliary loss,

$$L_{aux}^l = 1 - \sum_{k=1}^{K+1} \left(\sum_i \frac{Y_{k,i} M_{k,i}^l}{Y_{k,i}^2 + M_{k,i}^l} + \sum_i Y_{k,i} \log M_{k,i}^l \right), \quad (12)$$

where $Y_{k,i}$, $M_{k,i}^l$ refer to the results of the i -th pixel. The overall objective for training is formally defined as,

$$L = L_{seg} + \alpha \sum_l L_{aux}^l, \quad (13)$$

where α is empirically set as 0.05 in our experiments.

4 Experiments & Results

Datasets and Evaluation. We conduct experiments on three challenging datasets, including the **BTCV** dataset [Landman *et al.*, 2015], the **Lungs** dataset [Simpson *et al.*, 2019] and the **BraTS** dataset [Menze *et al.*, 2014]. The **BTCV** dataset consists of abdominal CT scans from 30 subjects for training and 20 subjects for testing. It includes 13 organs: *spleen, right kidney, left kidney, gallbladder, liver, stomach, aorta, pancreas, esophagus, inferior vena cava, portal and splenic veins, right and left adrenal gland*, respectively. The **Lungs** dataset consists of 63 CT volumes for a binary classification task, specifically focused on delineating lung cancer from the background. The dataset is divided into the training and testing set, with a split ratio of 80:20. The **BraTS** dataset comprises 484 MRI images, each with FLAIR, T1w, T1gd, and T2w channels. The dataset is divided into training, validation, and testing set, with a split ratio of 80:5:15. The target categories for segmentation are whole tumor (WT), enhancing tumor (ET), and tumor core (TC). We utilize the Dice Similarity Coefficient (DSC) as the performance metric for evaluation in all our experiments.

Implementation Details. we run all experiments based on Python 3.9, PyTorch 2.1.1 and Ubuntu 22.04. All training procedures have been performed on a single A800 GPU with 80GB memory. We report results with $96 \times 96 \times 96$ input size and patch resolution of (4, 4, 4) for BTCV dataset,

Methods	Spl	RKid	LKid	Gal	Eso	Liv	Sto	Aor	IVC	PSV	Pan	RAG	LAG	Avg
UNETR	90.48	82.51	86.05	58.23	71.21	94.64	72.06	86.57	76.51	70.37	66.06	66.25	63.04	76.00
Swin-UNETR	94.59	88.97	92.39	65.37	75.43	95.61	75.57	88.28	81.61	76.30	74.52	68.23	66.02	80.44
TransBTS	94.55	89.20	90.97	68.38	75.61	96.44	83.52	88.55	82.48	74.21	76.02	67.23	67.03	81.31
nnFormer	94.58	88.62	93.68	65.29	76.22	96.17	83.59	89.09	80.80	75.97	77.87	70.20	66.05	81.62
UNETR++	94.94	91.90	93.62	70.75	77.18	95.95	85.15	89.28	83.14	76.91	77.42	72.56	68.17	83.28
SIPL (ours)	94.90	94.85	95.23	85.37	81.54	97.14	92.99	89.98	85.77	80.79	83.92	71.38	77.15	87.00

Table 1: The comparison results with the state-of-the-art models on BTCV test set. The abbreviations are as follows: Spl: spleen, RKid: right kidney, LKid: left kidney, Gal: gallbladder, Eso: esophagus, Liv: liver, Sto: stomach, Aor: aorta, IVC: the inferior vena cava, PSV: portal and splenic veins, Pan: pancreas, RAG: right adrenal gland, LAG: left adrenal gland.

Model	DSC (%)	p-values
UNETR	73.29	< 0.0001
nnUNet	74.31	< 0.0001
Swin-UNETR	75.55	< 0.0001
nnFormer	77.95	< 0.0001
UNETR++	80.68	0.0014
SIPL (ours)	81.50	--

Table 2: Comparison results with SOTA methods on the Lungs dataset.

Model	DSC (%)	p-values
UNETR	81.2	0.0003
Swin-UNETR	81.5	< 0.0001
nnFormer	82.3	0.0025
UNETR++	82.8	0.0297
SIPL (ours)	83.0	--

Table 3: Comparison results with SOTA methods on the BraTS dataset.

where the models are trained for 2k epochs with the default initial learning rate of $4e^{-4}$, momentum of 0.9 and decay of $1e^{-5}$. For Lungs and BraTS datasets, consistent with [Abdelrahman Shaker *et al.*, 2022], the models are trained with input 3D patches of size $128 \times 128 \times 64$ for 1k epochs and the learning rate is 0.01 and weight decay is $3e^{-5}$. The transformer decoder consists of six kMax decoders evenly distributed across feature map with spatial resolutions of $\frac{1}{32}$, $\frac{1}{16}$, and $\frac{1}{8}$. Two kMax decoders are deployed for each resolution. Additionally, the number of object queries is empirically set to 32, and the threshold τ for query update linearly varies from 0.1 to 0.5 until epoch 50, and remains constant thereafter.

4.1 Comparison with State-of-the-Art Methods

We compare the segmentation performance between our proposed SIPL framework and the state-of-the-art (SOTA) methods, including nnUNet [Isensee *et al.*, 2021], UNETR [Hatamizadeh *et al.*, 2022], Swin-UNETR [Liu *et al.*, 2021], TransBTS [Wang *et al.*, 2021b], nnFormer [Zhou *et al.*, 2021] and UNETR++ [Abdelrahman Shaker *et al.*, 2022]. For fair comparison, all SOTA models were trained/tested on the same dataset splits. The base parameter settings are kept the




Figure 5: SIPL ablation study. (a) The influence of different object query numbers. (b) The influence of different threshold τ .

same as ours.

Multi-organ Segmentation on the BTCV Dataset. We provide a detailed comparison with other methods in the testing set of BTCV in Table 1. In comparison, SIPL outperforms the second-best method by 3.72% in terms of mean DSC scores, demonstrating a significant improvement. Specifically, SIPL achieves the highest DSC scores in 11 organs. Figure 4 shows a qualitative comparison of our method with existing SOTA approaches on BTCV dataset, where the inaccurate regions are marked with yellow dashed boxes. The visualization also demonstrates that SIPL obtains favorable segmentation performance than the existing models, achieving the SOTA on the BTCV dataset. For instance, In the first row, the *spleen* and *pancreas* are not visible in the slice, but were identified by UNETR, nnFormer, and UNETR++. In the second row, these models do not perform well in segmenting the pancreas. In contrast, SIPL accurately segments the *pancreas* with a superior mean DSC score of 83.92%. In addition, UNETR exhibits lower accuracy in segmenting the *right kidney*. Similarly, in the third row, we observe that UNETR, nnFormer, and UNETR++ struggle to accurately segment the *stomach*, while SIPL achieves the best segmentation performance with a DSC score of 92.99%. Lastly, the listed SOTA methods tend to interfere between adjacent tissues such as the *spleen* when segmenting the *liver* and *stomach* in the fourth row. These results demonstrate that SIPL outperforms listed SOTA methods with improved boundary delineation in these examples.

Evaluation Results on the Lungs and BraTS Datasets. Table 2 demonstrates the quantitative segmentation performance on the Lungs dataset. As shown, SIPL outperforms the second best method (UNETR++) by a margin of 0.82% in terms of mean DSC. The p-values from paired t-tests between

Model	Spl	RKid	LKid	Gal	Eso	Liv	Sto	Aor	IVC	PSV	Pan	RAG	LAG	Avg
- IPL	90.40	83.99	88.85	66.18	65.39	93.61	88.53	80.27	77.00	58.07	72.36	39.02	47.04	73.13
- SMG	93.91	91.80	90.38	80.17	79.63	96.69	91.23	89.96	84.07	77.76	80.17	65.59	72.73	84.16
SIPL (ours)	94.90	94.85	95.23	85.37	81.54	97.14	92.99	89.98	85.77	80.79	83.92	71.38	77.15	87.00

Table 4: Illustration on the effectiveness of each component in our method. We remove one component at a time. ‘- IPL’ means removing the IPL module and ‘- SMG’ means removing the SMG module.

Embedding	Spl	RKid	LKid	Gal	Eso	Liv	Sto	Aor	IVC	PSV	Pan	RAG	LAG	Avg
one-hot	86.16	92.56	90.98	65.38	74.49	95.46	84.41	86.63	76.70	73.58	74.34	47.13	8.26	73.54
BioBERT	90.78	92.48	89.95	56.27	63.16	96.58	91.16	88.58	82.46	76.29	80.73	65.48	10.68	75.74
CLIP	94.57	94.07	95.05	83.06	77.83	96.99	91.04	90.06	85.75	78.86	83.44	70.86	72.77	85.72
RanInit (ours)	94.90	94.85	95.23	85.37	81.54	97.14	92.99	89.98	85.77	80.79	83.92	71.38	77.15	87.00

Table 5: The ablation study on different prototype initialization strategies. Here RanInit represents random initialization.

our method and the listed state-of-the-art (SOTA) methods are all less than 0.05, indicating a statistically significant performance difference between our method and the listed SOTA methods.

Table 3 displays the comparison results on the BraTS dataset, with the DSC averaged across the three target regions (whole tumor, enhancing tumor, and tumor core). Our method achieves the best segmentation performance, as demonstrated. The p-values (all less than 0.05) further validate that the performance improvements between our method and the listed SOTA methods are statistically significant.

4.2 Ablation Study

We now analyze SIPL through a series of ablation studies and all experiments are performed on the BTCV dataset. The DSC scores are leveraged as the default evaluation metric.

Effectiveness of Each Component. We evaluate the effectiveness of each component by systematically removing them individually. Table 4 illustrates the results of these ablation experiments. When the IPL module is absent, our SIPL framework experiences a reduction of 13.87% in overall performance. Similarly, the absence of the SMG module leads to a 2.84% reduction in the average DSC score. These findings highlight the significant impact of the IPL module, which enables the learned prototypes to dynamically adapt to different inputs.

Evaluation on the Influence of Different Prototype Initializations. We further show the influence of different initialization strategies of the common prototype proposals (CPPs), including one-hot encoding [Zhang *et al.*, 2021], BioBERT embedding [Yasunaga *et al.*, 2022], CLIP embedding [Radford *et al.*, 2021] and random initialization embedding (referred as RanInit), as shown in Table 5. The performance comparison reveals that CLIP embedding, which utilizes image-text pre-training, significantly improves performance compared to traditional one-hot encoding and text-only pre-trained embedding (BioBERT). Surprisingly, random initialization embedding (RanInit) achieves an even larger improvement of 1.28% compared to CLIP embedding. This unexpected result may be attributed to the fact that predefined

fixed embeddings could potentially restrict the learning capabilities of the CPPs in our method.

Assessment of the Impact of Object Query Numbers.

The impact of varying object query numbers is depicted in Figure 5 (a). It is observed that when the number of object queries is set to 32, the best segmentation performance with a mean DSC of 87.00% is achieved. Results deteriorate when the number of object queries is either too high or too low. A low object query count fails to capture intra-class variation, while a high object query count increases the complexity of pseudo mask generation and may interfere with the main task learning.

Evaluation of the Different Threshold τ . We conduct an in-depth analysis of the constant threshold values of τ in Figure 5 (b). The results demonstrate that our framework achieves the best performance when the highest threshold is set to 0.5. Too large or too small threshold values yield inferior outcomes. This is because a large threshold filters out a significant portion of informative pixels during training, hampering the update of query embeddings. Conversely, if the threshold is too small, the learned object embeddings tend to focus on background regions, negatively affecting the segmentation performance.

5 Conclusion

In this paper, we propose a novel Self-supervised Instance-adaptive Prototype Learning (SIPL) method for MIS. Our SIPL framework introduces an Instance-adaptive Prototype Learning (IPL) scheme that adaptively learns prototypes tailored to different inputs. Moreover, we also design a Self-supervised Mask Generation (SMG) module to hierarchically generate pseudo masks with confidence scores. These pseudo masks have been validated to improve the segmentation performance by generating more discriminative instance-specific prototype proposals. Experiments and ablation studies validate the effectiveness of each component and demonstrate favorable performance compared to other SOTA methods.

Acknowledgements

This study was supported under the Key Research and Development Program of Zhejiang Province (Grant No. 2023C03192). It was also funded by the National Science Foundation of China (Grant No. 62201341).

Contribution Statement

Guoyan Liang and Qin Zhou made equal contributions. All the authors participated in designing research, performing research, analyzing data, and writing the paper.

References

- [Abdelrahman Shaker *et al.*, 2022] Muhammad Maaz Abdelrahman Shaker, Hanoona Rasheed, Salman Khan, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Unetr++: Delving into efficient and accurate 3d medical image segmentation. *arXiv:2212.04497*, 1, 2022.
- [Boney and Ilin, 2017] Rinu Boney and Alexander Ilin. Semi-supervised few-shot learning with prototypical networks. *CoRR abs/1711.10856*, 2017.
- [Cao *et al.*, 2022] Hu Cao, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiaopeng Zhang, Qi Tian, and Manning Wang. Swin-unet: Unet-like pure transformer for medical image segmentation. In *European conference on computer vision*, pages 205–218. Springer, 2022.
- [Carion *et al.*, 2020] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020.
- [Chen *et al.*, 2019] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. *arXiv:1904.04232*, 2019.
- [Cheng *et al.*, 2022] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1290–1299, June 2022.
- [Dosovitskiy *et al.*, 2020] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv:2010.11929*, 2020.
- [Guerriero *et al.*, 2018] Samantha Guerriero, Barbara Caputo, and Thomas Mensink. Deepncm: Deep nearest class mean classifiers. In *International Conference on Learning Representations - Workshop*, 2018.
- [Hatamizadeh *et al.*, 2022] Ali Hatamizadeh, Yucheng Tang, Vishwesh Nath, Dong Yang, Andriy Myronenko, Bennett Landman, Holger R Roth, and Daguang Xu. Unetr: Transformers for 3d medical image segmentation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 574–584, 2022.
- [Huang *et al.*, 2020] Huimin Huang, Lanfen Lin, Ruofeng Tong, Hongjie Hu, Qiaowei Zhang, Yutaro Iwamoto, Xi-anhua Han, Yen-Wei Chen, and Jian Wu. Unet 3+: A full-scale connected unet for medical image segmentation. In *ICASSP 2020-2020 IEEE international conference on acoustics, speech and signal processing*, pages 1055–1059. IEEE, 2020.
- [Isensee *et al.*, 2021] Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, 18:203–211, 2021.
- [Landman *et al.*, 2015] Bennett Landman, Zhoubing Xu, J Igelsias, Martin Styner, Thomas Langerak, and Arno Klein. Miccai multi-atlas labeling beyond the cranial vault—workshop and challenge. In *Proc. MICCAI Multi-Atlas Labeling Beyond Cranial Vault—Workshop Challenge*, volume 5, page 12, 2015.
- [Li *et al.*, 2021] Shaohua Li, Xiuchao Sui, Xiangde Luo, Xinxing Xu, Yong Liu, and Rick Goh. Medical image segmentation using squeeze-and-expansion transformers. *arXiv:2105.09511*, 2021.
- [Liu *et al.*, 2021] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021.
- [Mensink *et al.*, 2013] Thomas Mensink, Jakob Verbeek, Florent Perronnin, and Gabriela Csurka. Distance-based image classification: Generalizing to new classes at near-zero cost. *IEEE transactions on pattern analysis and machine intelligence*, 35(11):2624–2637, 2013.
- [Menze *et al.*, 2014] Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34(10):1993–2024, 2014.
- [Radford *et al.*, 2021] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [Ronneberger *et al.*, 2015] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015.
- [Santoro *et al.*, 2016] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *International conference on machine learning*, pages 1842–1850. PMLR, 2016.

- [Simpson *et al.*, 2019] Amber L Simpson, Michela Antonelli, Spyridon Bakas, Michel Bilello, Keyvan Farahani, Bram Van Ginneken, Annette Kopp-Schneider, Bennett A Landman, Geert Litjens, Bjoern Menze, et al. A large annotated medical image dataset for the development and evaluation of segmentation algorithms. *arXiv:1902.09063*, 2019.
- [Snell *et al.*, 2017] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30, 2017.
- [Sung *et al.*, 2018] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1199–1208, 2018.
- [Tsai *et al.*, 2003] Andy Tsai, Anthony Yezzi, William Wells, Clare Tempany, Dewey Tucker, Ayres Fan, W Eric Grimson, and Alan Willsky. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE transactions on medical imaging*, 22(2):137–154, 2003.
- [Vinyals *et al.*, 2016] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *Advances in neural information processing systems*, 29, 2016.
- [Wang *et al.*, 2021a] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Max-deeplab: End-to-end panoptic segmentation with mask transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5463–5474, 2021.
- [Wang *et al.*, 2021b] Wenxuan Wang, Chen Chen, Meng Ding, Hong Yu, Sen Zha, and Jiangyun Li. Transbts: Multimodal brain tumor segmentation using transformer. In *Medical Image Computing and Computer Assisted Intervention*, pages 109–119. Springer, 2021.
- [Wang *et al.*, 2022] Runze Wang, Qin Zhou, and Guoyan Zheng. Few-shot medical image segmentation regularized with self-reference and contrastive learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 514–523. Springer, 2022.
- [Wen *et al.*, 2016] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *Computer Vision—ECCV 2016*, pages 499–515. Springer, 2016.
- [Xiao *et al.*, 2018] Xiao Xiao, Shen Lian, Zhiming Luo, and Shaozi Li. Weighted res-unet for high-quality retina vessel segmentation. In *2018 9th international conference on information technology in medicine and education*, pages 327–331. IEEE, 2018.
- [Yasunaga *et al.*, 2022] Michihiro Yasunaga, Jure Leskovec, and Percy Liang. Linkbert: Pretraining language models with document links. *arXiv:2203.15827*, 2022.
- [Yu *et al.*, 2021] Qinji Yu, Kang Dang, Nima Tajbakhsh, Demetri Terzopoulos, and Xiaowei Ding. A location-sensitive local prototype network for few-shot medical image segmentation. In *2021 IEEE 18th international symposium on biomedical imaging (ISBI)*, pages 262–266. IEEE, 2021.
- [Yu *et al.*, 2022] Qihang Yu, Huiyu Wang, Siyuan Qiao, Maxwell Collins, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. k-means mask transformer. In *European Conference on Computer Vision*, pages 288–307. Springer, 2022.
- [Yuan *et al.*, 2023] Mingze Yuan, Yingda Xia, Hexin Dong, Zifan Chen, Jiawen Yao, Mingyan Qiu, Ke Yan, Xiaoli Yin, Yu Shi, Xin Chen, et al. Devil is in the queries: Advancing mask transformers for real-world medical image segmentation and out-of-distribution localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23879–23889, 2023.
- [Zhang *et al.*, 2021] Jianpeng Zhang, Yutong Xie, Yong Xia, and Chunhua Shen. Dodnet: Learning to segment multi-organ and tumors from multiple partially labeled datasets. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1195–1204, 2021.
- [Zhou *et al.*, 2018] Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. Unet++: A nested u-net architecture for medical image segmentation. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pages 3–11. Springer, 2018.
- [Zhou *et al.*, 2021] Hong-Yu Zhou, Jiansen Guo, Yinghao Zhang, Lequan Yu, Liansheng Wang, and Yizhou Yu. nn-former: Interleaved transformer for volumetric segmentation. *arXiv:2109.03201*, 2021.
- [Zhou *et al.*, 2022] Tianfei Zhou, Wenguan Wang, Ender Konukoglu, and Luc Van Gool. Rethinking semantic segmentation: A prototype view. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2582–2593, 2022.