

Polyper: Boundary Sensitive Polyp Segmentation

Hao Shao¹, Yang Zhang², Qibin Hou^{1*}

¹VCIP, School of Computer Science, Nankai University

²Department of Genetics and Cell Biology, College of Life Sciences, Nankai University
shaoh@mail.nankai.edu.cn, houqb@nankai.edu.cn

Abstract

We present a new boundary sensitive framework for polyp segmentation, called Polyper. Our method is motivated by a clinical approach that seasoned medical practitioners often leverage the inherent features of interior polyp regions to tackle blurred boundaries. Inspired by this, we propose explicitly leveraging polyp regions to bolster the model’s boundary discrimination capability while minimizing computation. Our approach first extracts boundary and polyp regions from the initial segmentation map through morphological operators. Then, we design the boundary sensitive attention that concentrates on augmenting the features near the boundary regions using the interior polyp regions’s characteristics to generate good segmentation results. Our proposed method can be seamlessly integrated with classical encoder networks, like ResNet-50, MiT-B1, and Swin Transformer. To evaluate the effectiveness of Polyper, we conduct experiments on five publicly available challenging datasets, and receive state-of-the-art performance on all of them. Code is available at https://github.com/haoshao-nku/medical_seg.git.

Introduction

Colon polyps are protruding growths within the colon mucosa, exhibiting considerable variability in shape, texture, and color (Pooler et al. 2023). Importantly, colon polyps are recognized as precancerous lesions closely associated with the development of colon cancer (Djinbachian et al. 2020). Consequently, there is a pressing need to enhance both the efficiency of early detection and the accuracy of polyp contour segmentation. Polyps pose diagnostic challenges during colonoscopy due to the inconspicuous borders and low contrast. In their initial stages, polyps often manifest smaller dimensions, resulting in less defined margins that exacerbate detection difficulties.

To address this challenge, one of the current research trends is to maximize the integration of features at various scales to preserve as many boundary details as possible. Typically, Wu et al. (Wu et al. 2021) introduced semantic calibration and refinement techniques to bridge the semantic gap between feature maps at different levels. SwinE-Net (Park and Lee 2022) refines the multi-level features extracted from

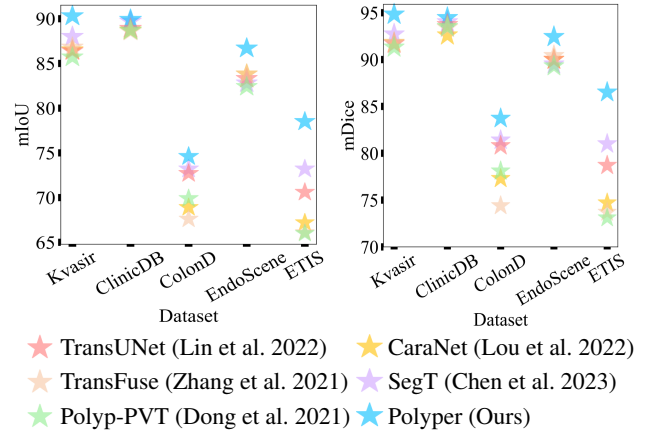


Figure 1: Comparisons with other methods on five datasets.

both CNN and Swin Transformer architectures using multi-dilation convolutions and multi-feature aggregation blocks. While this approach can perform well in the mid to late stages of a well-margined lesion, it struggles to effectively handle early polyps with lower edge contrast.

Another popular strategy is to generate a polyp mask to coarsely localize the polyp and then enhance the semantic features around the potential boundaries. As an early attempt, SegT (Chen, Ma, and Zhang 2023) highlights the boundaries of the polyp area by assessing the disparity between the foreground and background regions. However, accurately capturing the polyp boundaries poses a challenge due to the ambiguity caused by blending polyp boundaries with the surrounding mucosa. Relying solely on the synergistic effect between foreground and background information may not lead to accurate polyp segmentation. CaraNet (Lou et al. 2022) leverages the long-distance interaction of axial attention to calculate the pairwise affinity from a global perspective. The roughly estimated polyp regions are removed from the deep features, and then features at different scales are utilized to supplement the boundary details to produce good segmentation results. However, axial attention may not consistently benefit the network because it risks inadvertently excluding crucial boundary information (Thanh Duc et al. 2022). Hence, the methods above cannot sufficiently address the issue of blurred boundaries.

*Corresponding author.

In endoscopic screening, skilled medical practitioners often utilize the polyp characteristics of non-boundary regions to address the challenge of boundary blurring. Drawing inspiration from this observation, our method first generates an initial segmentation map and employs morphological operators to partition the polyp into boundary and non-boundary regions. We then leverage the semantic features extracted from the non-boundary regions to refine the boundary regions with a novel boundary sensitive attention module, which can take advantage of both global and local features to identify the real polyp boundaries.

Our method, called Polyper, is simple, easy to follow, and suitable for practical medical scenarios. We conduct a series of experiments on five widely used datasets to evaluate Polyper. As shown in Fig. 1, Polyper outperforms previous methods in all datasets in terms of mDice and mIoU scores. Typically, because of the heterogeneity of polyp features, previous works mostly do not perform well for small polyps in the early growth stage. In the subsequent ablation study, we verify that our method performs well on small polyps.

Our main contributions can be summarized as follows:

- We present a novel module, named boundary sensitive attention, which can model the relationships between the boundary regions and interior regions of polyps to augment the features near the boundary regions by capitalizing on the inherent characteristics of the interior regions.
- We design a novel decoder for polyp segmentation composed of two distinct stages: potential boundary extraction and boundary sensitive refinement. This decoder helps us effectively identify the real polyp boundaries, resolving the challenge of boundary blurring in endoscopy.
- We evaluate the proposed Polyper on five popular polyp segmentation datasets and set new records on almost all the benchmarks.

Related Work

Architectures for Medical Image Segmentation. CNN is one of the most widely used deep neural network architectures in medical image segmentation. A typical example should be U-Net (Ronneberger, Fischer, and Brox 2015), one of the most classic networks. Attention U-Net (Oktay et al. 2018) introduces a novel attention gate mechanism that empowers the model to selectively focus on targets of diverse shapes and sizes. Res-UNet (Xiao et al. 2018) incorporates a weighted attention mechanism to improve segmentation performance. R2U-Net (Alom et al. 2018) ingeniously merges the strengths of residual networks (He et al. 2016) and U-Net. KiU-Net (Valanarasu et al. 2020) proposes an innovative structure that leverages under-complete and super-complete features to enhance the segmentation of lesion regions with small anatomical structures. AttResDU-Net (Khan et al. 2023) incorporates attention gates on the skip connections and residual connections in the convolutional blocks. Recently, there has been a surge of interests in utilizing Transformers (Huang et al. 2022) for polyp segmentation. SwinMM (Wang et al. 2023) develops a cross-viewpoint decoder that aggregates multi-viewpoint information through cross-attention blocks. MPU-Net (Yu and Han

2023) aims to achieve precise localization by combining image serialization with a positional attention module, enabling the model to comprehend deeper contextual dependencies effectively. Segtran (Li et al. 2021) presents a novel Squeeze-and-Expansion transformer.

Refinement Methods. One avenue for refinement is to maximize the utilization of features at different scales. Wu et al. (Wu et al. 2021) employed semantic calibration and refinement techniques to bridge the semantic gap between different levels of feature mapping. SwinE-Net (Park and Lee 2022) refines and enhances the multi-level features extracted from CNN and Swin Transformer through multi-dilation convolutions and multi-feature aggregation blocks. FTMF-Net (Liu et al. 2023) presents a Fourier transform multi-scale feature fusion network for segmenting small polyp objects. Another kind of approaches involves targeting specific areas for refinement. PraNet (Fan et al. 2020) and CaraNet (Lou et al. 2022) both integrate the Reverse Attention module (Chen et al. 2018), a specialized component that accentuates the boundaries between polyps and their surroundings. Xie et al. (Xie et al. 2020) introduced assisted boundary supervision as a guiding mechanism for refining glass segmentation, aiding in predicting uncertain regions around the boundaries. In contrast to direct boundary feature enhancement, He et al. (He et al. 2021) advocated supervising the non-edge portion through a residual approach to attain finer edges. Zhang et al. (Zhang et al. 2020) proposed the Local Context Attention module to pass local context features from the encoder layer to the decoder layer, enhancing the focus on hard regions. RFENet (Fan et al. 2023) introduces a structure focus refinement module to facilitate fine-grained feature refinement of fuzzy points around the boundaries. EAMNet (Sun, Jiang, and Qi 2023) considers edge detection and camouflaged object segmentation as an interlinked cross-refinement process.

Method

Accurately recognizing polyp boundaries from the surrounding mucosa is challenging due to the low contrast between the polyp and surrounding tissues. To address this issue, a prevalent strategy is to enhance the quality of the segmentation map by refining the semantic features near the potential boundaries. However, the ambiguous nature of polyp boundaries mixed with the surrounding tissues often hinders accurate prediction of the polyp boundary regions. In the context of endoscopic screening, seasoned medical practitioners often leverage the inherent features of polyps within non-boundary regions to effectively tackle the issue of blurred boundaries. Motivated by this clinical approach, we present Polyper.

Fig. 2 provides an overview of Polyper. Like most previous works for polyp segmentation, we employ the classical encoder-decoder architecture. The encoder plays a crucial role in extracting features at different scales and levels, enabling the model to capture coarse and fine details. We utilize the Swin-T from Swin Transformer (Liu et al. 2021) as our encoder. The decoder comprises two distinctive stages: potential boundary extraction and boundary sen-

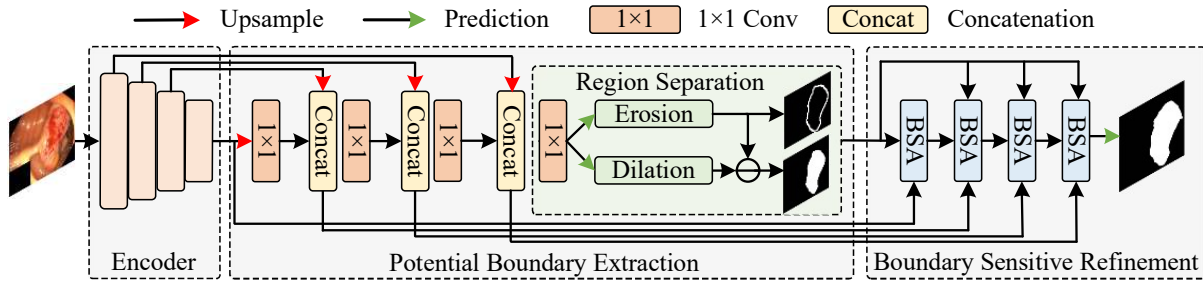


Figure 2: Overall architecture of Polyper. The decoder is divided into two main stages. The first potential boundary extraction (PBE) stage aims to capture multi-scale features from the encoder, which are then aggregated to generate the initial segmentation results. Next, we extract the predicted polyps’ potential boundary and interior regions using morphology operators. In the second boundary sensitive refinement (BSR) stage, we model the relationships between the potential boundary and interior regions to generate better segmentation results.

sitive refinement. In the potential boundary extraction stage, the encoder’s multi-scale features are aggregated to generate an initial prediction, which is used to extract the predicted polyps’ potential boundaries and interior regions. The boundary sensitive stage leverages the distinctive characteristics of the interior regions to enhance the model’s precision by modeling the relationships between the potential boundary regions and the interior regions. In what follows, we will describe these two stages in detail.

Potential Boundary Extraction

An overview of the potential boundary extraction stage is depicted in Fig. 2. We predict the segmentation map using 1×1 convolution and employ morphology operators to extract the boundaries and interior polyp regions from the initial segmentation result. This stage can be separated into feature aggregation and region separation.

Feature Aggregation. In the feature aggregation part, we use 1×1 convolution and the concatenation operation to aggregate features of different scales. Given the features from the four stages of the encoder, denoted as E_0, E_1, E_2 , and E_3^1 , we first build a feature pyramid following (Lin et al. 2017). Specifically, we resize feature maps E_1, E_2 , and E_3 to ensure they share the same size as E_0 through linear interpolation, yielding E'_1, E'_2 , and E'_3 . The formula for calculating feature map D_i of the intermediate layer in each stage of the feature aggregation process is:

$$D_i = [\text{Conv}_{1 \times 1}(D_{i+1}), E'_i], \quad (1)$$

where $i \in \{0, 1, 2, 3\}$ and $[\dots]$ means the concatenation operation. Here, D_3 is equivalent to E'_3 and E'_0 is equivalent to E_0 . $\text{Conv}_{1 \times 1}$ means 1×1 convolution.

Region Separation. We introduce a region separation module to separate the boundaries and the interior polyp regions from the initial segmentation map. Specifically, given the output D_0 of the last stage of the feature aggregation, the initial segmentation mask f_m is obtained by a 1×1 convolution. Then, we utilize the erosion operator (E) and the

¹The resolutions are denoted as $\frac{H}{4} \times \frac{W}{4}$, $\frac{H}{8} \times \frac{W}{8}$, $\frac{H}{16} \times \frac{W}{16}$, and $\frac{H}{32} \times \frac{W}{32}$, respectively. H and W are the height and width of the input image, respectively.

dilation operator (D)² to separate the boundary and interior regions from the initial segmentation mask. At each iteration, the mask edge erodes or expands by one pixel. These regions offer essential guidance for the subsequent refinement process. The separation process can be written as:

$$P_{CR} = E(f_m) \times T, \quad (2)$$

$$P_{BR} = (D(f_m) \times T - P_{CR}), \quad (3)$$

where P_{CR} is the interior regions while P_{BR} denotes the potential boundary regions. Here, T is the number of operations for operator D . It is noteworthy that the total number of operations of operator D and operator E is the same.

Boundary Sensitive Refinement

As mentioned in the introduction section, the diverse characteristics of polyps in different growth stages, including their shape, texture, and color, lead to significant challenges to the robustness of polyp segmentation methods. To address this, we present the boundary sensitive refinement stage to refine the feature of boundary regions based on P_{CR} and P_{BR} .

The boundary sensitive refinement stage can be illustrated in the right part of Fig. 2, which includes boundary sensitive attention module and full-stage sensitivity strategy. We first extract features from the boundary, interior polyp, and background regions, respectively. Then, we leverage the cross-attention mechanism to model the relationships between the boundary region and the interior polyp regions as well as the relationships between the interior polyp regions and the background regions. This enables simultaneous encoding of global and local features to improve the quality of segmentation results with accurate boundaries. After the above process is done, the features corresponding to different regions will be restored to their initial positions. The goal is to ensure the efficient use of hardware resources. To realize the process above, we introduce a novel boundary sensitive attention module. In addition, deep features excel at capturing and conveying semantic information, while low-level features are good at representing complex geometric details.

²Please refer to Chapter 9 of the Digital image processing (Gonzales and Woods 1987) for more details.

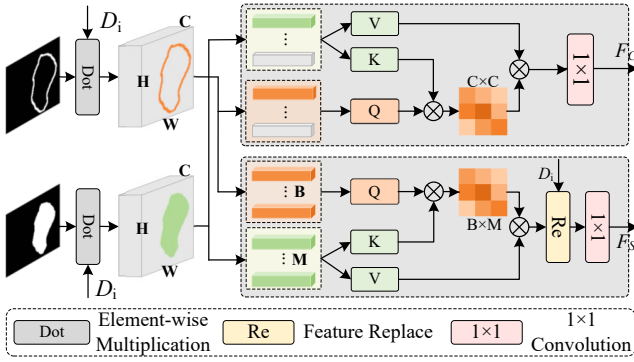


Figure 3: Detailed structure of boundary sensitive attention (BSA) module. This process is separated into two parallel branches, which systematically capitalize on the distinctive attributes of polyps at various growth stages, both in terms of spatial and channel characteristics. ‘B’ and ‘M’ indicate the number of pixels in the boundary and interior polyp regions within an input of size $H \times W$ and C channels.

We introduce a full-stage sensitivity strategy that tactically harnesses the strengths of both deep and shallow features.

Boundary Sensitive Attention. The structure of the boundary sensitive attention module is illustrated in Fig. 3. It comprises two branches thoroughly exploring polyps’ inherent characteristics by building spatial and channel attentions. We first describe the working mechanism of the boundary sensitive attention module on encoding spatial information.

Given the input D_i , the boundary region mask P_{BR} , and the interior polyp region mask P_{CR} , we first perform element-wise product operations on D_i over P_{BR} and P_{CR} to attain the features corresponding to the boundary region and the interior polyp region, denoted as F_{BR} and F_{CR} , respectively. To better discover the real polyp boundaries, we do not consider the background regions. We treat F_{BR} as the query matrix and F_{CR} as the key and value matrices and compute the cross-attention between them as follows:

$$F_S = \text{MHCA}_S(F_{BR}, F_{CR}, F_{CR}), \quad (4)$$

where $\text{MHCA}_S(\cdot, \cdot, \cdot)$ denotes multi-head cross-attention (Vaswani et al. 2017) along the spatial dimension. This operation aims to more accurately mine regions that are real polyps by leveraging the priors of the interior polyp regions based on our observation mentioned at the beginning of this section. In addition, as the cross-attention is computed over only F_{BR} and F_{CR} but not the background region, the computational cost is also low. The results will then be put back to the corresponding positions of D_i .

We also consider using the background region to capture the variations and correlations between the background and boundary regions and between the background and internal polyp regions. We use F'_{BR} and F'_{CR} to denote boundary region features and interior polyp region features that contain background information, respectively. We treat F'_{BR} as the query matrix and F'_{CR} as the key and value matrices and compute the cross-attention between them as follows:

$$F_C = \text{MHCA}_C(F'_{BR}, F'_{CR}, F'_{CR}), \quad (5)$$

where $\text{MHCA}_C(\cdot, \cdot, \cdot)$ denotes multi-head cross-attention along the channel dimension, following (Yin et al. 2022; Zamir et al. 2022) to save computations. The goal of this operation is to capture the consistency and correlation among the different regions from a global view.

The output of the proposed boundary sensitive attention module can be formulated as:

$$F_i = \text{Conv}_{1 \times 1}(F_S) + \text{Conv}_{1 \times 1}(F_C) + D_i, \quad (6)$$

where $\text{Conv}_{1 \times 1}$ is 1×1 convolution.

Full-Stage Sensitive Strategy. As depicted in the right part of Fig 2, the boundary sensitive attention module is used to leverage the features at various scales for refining the boundary regions progressively. The formulation of this full-stage sensitive strategy can be described as follows. At the deepest stage, the refinement process is calculated as follows:

$$F_3 = \text{BSA}(D_3, P_{BR}, P_{CR}), \quad (7)$$

where $\text{BSA}(\cdot, \cdot, \cdot)$ denotes the boundary sensitive attention module. Furthermore, the refinement process in the subsequent stages can be defined as:

$$F_i = \text{BSA}(\text{Conv}_{1 \times 1}(F_{i+1}) + D_i, P_{BR}, P_{CR}). \quad (8)$$

Finally, we add a 1×1 convolution to F_0 to generate the final predictions. From the above descriptions, we can see that our decoder primarily contains 1×1 convolutions and simple cross-attention. This makes our method efficient.

Experiments

Datasets

We report results on datasets used in Pranut (Fan et al. 2020), including: Kvasir-SEG, CVC-ClinicDB, CVC-ColonDB, EndoScene, and ETIS. The training set consists of 900 images from Kvasir-SEG and 550 images from ClinicDB. The test sets comprise 100 images from Kvasir-SEG, 62 images from CVC-ClinicDB, 380 images from CVC-ColonDB, 60 images from EndoScene, and 196 images from ETIS.

Implementation Details

We use PyTorch (Paszke et al. 2019) and mmsegmentation³ to implement Polyper. The input resolution during training is set to 224×224 , and the batch size is set to 6. The number of iterations during training is 80k. We employ the AdamW optimizer for training with an initial learning rate of 0.0002, a momentum of 0.9, and a weight decay of $1e-4$. All the experiments are conducted on one NVIDIA RTX 3090 GPU. Following the configuration of previous research (Lin et al. 2022), evaluation metrics mIoU and mDice are employed. The calculation of flops and parameters in the experiments is based on an input size of 512×512 , and the calculation method is from the mmsegmentation project.

Analysis of Experimental Results

Comparison with the State-of-the-Art Methods. We compare the segmentation performance of Polyper with other

³<https://github.com/open-mmlab/msegmentation>.

Methods	Params (M)	Flops (G)	Kvasir		ClinicDB		ColonDB		EndoScene		ETIS	
			mDice	mIoU	mDice	mIoU	mDice	mIoU	mDice	mIoU	mDice	mIoU
U-Net	24.56	38.26	81.80	74.60	82.30	75.50	51.20	44.40	71.00	62.60	39.80	33.50
UNet++	25.09	84.30	82.10	74.30	79.40	72.90	48.30	41.00	70.70	62.40	40.10	34.40
SFA	-	-	72.30	61.10	70.00	60.70	46.90	34.70	46.70	32.90	29.70	21.70
ACSNet	46.02	29.45	89.80	83.80	88.20	82.60	71.60	64.90	86.30	78.70	57.80	50.90
PraNet	32.50	221.90	89.80	84.00	89.90	84.90	70.90	64.00	87.10	79.70	62.80	56.70
SANet	-	-	90.40	86.40	93.70	88.90	75.30	67.00	88.80	81.50	75.00	65.40
TransFuse	115.59	38.73	91.80	86.80	93.40	88.60	74.40	67.60	90.40	83.80	73.70	66.10
TransUNet	105.28	24.66	91.30	85.70	93.50	88.70	78.10	69.90	89.30	82.40	73.10	66.00
HarDNet-MSEG	33.80	192.74	91.20	85.70	93.20	88.20	73.10	66.00	88.70	82.10	67.70	61.30
DS-TransUNet	177.44	30.97	93.40	88.80	93.80	89.10	79.80	71.70	88.20	81.00	77.20	69.80
SwinE-Net	-	-	92.00	87.00	93.80	89.20	80.40	72.50	90.60	84.20	75.80	68.70
Polyp-PVT	-	-	91.70	86.40	93.70	88.90	80.80	72.70	90.00	83.30	78.70	70.60
CaraNet	46.64	21.69	91.80	86.50	93.60	88.70	77.30	68.90	90.30	83.80	74.70	67.20
ColonFormer	52.94	22.94	92.40	87.60	93.20	88.40	81.10	73.30	90.60	84.20	80.10	72.20
SegT	-	-	92.70	88.00	94.00	89.70	81.40	73.20	89.50	82.80	81.00	73.20
Polyper w/o BSR	28.70	37.42	91.69	85.16	89.29	81.86	74.93	66.00	86.12	80.25	75.36	66.77
Polyper w/o RS	29.82	48.26	91.97	85.58	88.35	83.68	75.60	66.84	87.69	82.12	82.36	73.13
Polyper	29.11	43.54	94.82	90.36	94.45	89.85	83.72	74.55	92.43	86.72	86.51	78.51

Table 1: Comparisons with other methods. ‘Polyper w/o RS’ means no region separation is used in potential boundary extraction, and the entire foreground region is refined instead of the boundary region. ‘Polyper w/o BSR’ means not to refine the initial segmentation results.




Figure 4: Performance for small polyps on five dataset. ‘Proportion Size’ is the ratio of the polyp’s size to the entire image.

state-of-the-art models. Table 1 presents a comprehensive comparison with CNN-based methods (Ronneberger, Fischer, and Brox 2015; Zhou et al. 2018; Fang et al. 2019; Zhang et al. 2020; Fan et al. 2020; Wei et al. 2021), Transformer-based methods (Zhang, Liu, and Hu 2021; Huang, Wu, and Lin 2021; Lin et al. 2022; Dong et al. 2021; Park and Lee 2022), and refinement-based methods (Lou et al. 2022; Thanh Duc et al. 2022; Chen, Ma, and Zhang 2023). As depicted in Table 1, it is evident that our proposed boundary sensitive method, called Polyper, outperforms other listed methods. The visualization results are shown in Fig. 5, from which it can be observed that our proposed method outperforms the previous methods in boundary processing and the processing of small polyps.

Small Polyp Analysis. We also evaluate the performance of small polyps. This type of polyps tends to appear at the onset of the disease and has lower contrast (Antonelli et al. 2021). Specifically, we follow the approach presented in CaraNet (Lou et al. 2022) and focus on evaluating polyps that make up less than 6% of the entire image. In this experiment, we compare with a CNN-based method PraNet (Fan et al. 2020), a Transformer-based method DS-TransUNet (Lin et al. 2022), and a refinement-based method

CaraNet (Lou et al. 2022). The results are shown in Fig. 4. It can be observed that Polyper performs better in small polyps than other methods thanks to the boundary sensitive strategy. Notably, Polyper even outperforms CaraNet, a method specially designed for small polyp targets.

Ablation Study

We conduct extensive ablation experiments on the Kvasir dataset to analyze Polyper.

Ablations on Encoder. Initially, we conduct experiments to evaluate the impact of different encoders. We choose the commonly used ResNet-50 (He et al. 2016) and MiT-B1 (Xie et al. 2021) as the evaluation encoders. As shown in Table 2, when using our decoder and ResNet-50 as an encoder, a significant enhancement of 2.22 for mIoU and 1.16 for mDice is observed compared to not refining the initial segmentation results. Additionally, using the MiT-B1 as encoder brings significant enhancement. This demonstrates the broad applicability of Polyper to various encoders.

Ablations on Potential Boundary Extraction. We first evaluate the influence of different feature aggregation methods. We consider the Non-Local block (Wang et al. 2018)




Figure 5: Segmentation results of different methods on the Kvasir, ClinicDB and ColonDB datasets (from top to bottom).

Encoder	Decoder	mIoU	mDice
ResNet-50	Ours w/o BSR	82.75	90.13
ResNet-50	Ours w/ BSR	84.97 (2.22 ↑)	91.29 (1.16 ↑)
MiT-B1	Ours w/o BSR	87.74	93.28
MiT-B1	Ours w/ BSR	89.19 (1.45 ↑)	94.13 (0.85 ↑)
Swin-T	Non-Local w/o BSR	84.56	91.31
Swin-T	Non-Local w/ BSR	86.55 (1.99 ↑)	92.55 (1.24 ↑)
Swin-T	Hamburger w/o BSR	82.40	89.85
Swin-T	Hamburger w/ BSR	83.67 (1.27 ↑)	90.65 (0.80 ↑)
Swin-T	Ours w/o BSR	87.12	92.90
Swin-T	Ours w/ BSR	90.57 (3.45 ↑)	94.49 (1.59 ↑)

Table 2: Ablations on different encoders and feature aggregation methods. ‘Polyper w/o BSR’ means not to refine the initial segmentation results

and Hamburger (Geng et al. 2021). As Table 2 shows, our proposed boundary sensitive approach Polyper is compatible with different feature aggregation methods, reflecting our method’s generalization ability. Furthermore, it is noticeable that the Non-Local block and Hamburger do not surpass the performance of our proposed feature aggregation method when used for feature aggregation. We attribute this to the fact that these two methods are originally designed for semantic segmentation in natural images. Given the limited amount of medical data available for segmentation and the resulting challenges of network convergence, their performance in medical segmentation remains unsatisfactory.

Furthermore, we conduct experiments to analyze the effectiveness of Region Separation (RS). In the absence of RS, the subsequent refinement stage employs the full mask to encompass the entire foreground area. As presented in Table 3, refinement with RS demonstrates improvements of 4.40 on mIoU and 2.59 on mDice compared to refinement with the whole mask. This is because the boundary regions of the

RS	Number of iterations						mIoU	mDice
	1	2	3	4	5	6		
							86.55	92.57
✓	✓						87.14	92.91
✓		✓					87.95	93.40
✓			✓				88.79	93.90
✓					✓		89.66	94.42
✓						✓	88.79	93.30
✓				✓			90.57	94.49

Table 3: Ablations on region separation. ‘Number of iterations’: number of iterations applied by the erosion operator.

initial segmentation results contain unreliable features with low confidence. When the whole foreground is refined, it is affected by these unreliable features, which reduces the quality of the segmentation results.

Finally, we conduct experiments on the width of the boundary region, and the results are shown in Table 3. The width of the boundary region is determined by performing a subtraction operation between the mask after applying the erosion operator and the mask after applying the dilation operator, following (Zhu, Qiao, and Yang 2023) to calculate. For this experiment, our evaluation focuses on assessing the impact of the number of iterations executed by the erosion operator. From the table, it can be observed that the optimal number of iterations for the erosion operator is 4. We suggest that the readers use this number in their experiments.

Ablations on Boundary Sensitive Refinement. Here, we conduct experiments to validate the importance of boundary sensitive refinement stage. First, we evaluate the importance of the two branches, spatial attention and channel attention. Table 4 clearly illustrates the contributions of spatial attention and channel attention. This illustrates the effectiveness of fully utilizing the relationship between the boundary re-




Figure 6: Visual analysis of the feature maps produced by different versions of Polyper. All images are selected from the Kvasir dataset. (a) w/o Boundary Sensitive Refinement; (b) w/o Region Separation; (c) w/o Boundary Sensitive Attention; (d) Polyper w/o channel attention; (e) w/o spatial attention; (f) Polyper; (g) segmentation results w/o Boundary Sensitive ; (h) segmentation results w/o Region Separation ; (i) segmentation results w/ Polyper; (j) ground-truth annotations.

SA	CA	mIoU	mDice
		87.65	93.22
✓		89.45 (1.80 ↑)	94.28 (1.06 ↑)
	✓	89.16 (1.51 ↑)	94.12 (0.90 ↑)
✓	✓	90.57 (2.92 ↑)	94.49 (1.27 ↑)

Table 4: Ablations on the boundary sensitive attention module. ‘SA’: spatial attention. ‘CA’: channel attention.

D_3	D_2	D_1	D_0	mIoU	mDice
✓				85.43	91.88
✓	✓			86.26 (0.83 ↑)	92.16 (0.28 ↑)
✓	✓	✓		88.90 (2.64 ↑)	93.48 (1.32 ↑)
✓	✓	✓	✓	90.57 (1.67 ↑)	94.49 (1.01 ↑)

Table 5: Ablations on full-stage sensitive strategy. When all features are used, the performance is the best.

gion and the interior polyp region and the relationship between the interior polyp region and the background region for producing polyp regions with accurate boundaries.

Then, we evaluate the importance of our full-stage sensitive strategy. Table 5 illustrates the effectiveness of fully utilizing different levels of features to improve the quality of segmentation results with accurate boundaries. We can see that gradually incorporating more features from lower levels can continuously increase the model’s performance.

Visual Analysis. We use the method in (Komodakis and Zagoruyko 2017) to visualize the feature maps generated by different versions of Polyper. The visual results are presented in Fig. 6. From Fig. 6(b), it can be observed that when refining the initial results from a global perspective, this approach has limited effect and does not solve the problem of edge blurring due to the presence of interference in the boundary regions with low-confidence predictions. In




Figure 7: Failure cases of Polyper.

contrast, as can be seen from Fig. 6(f), our method is effective when augmented by modeling the relationship between the interior polyp region and the boundary region and the relationship between the interior polyp region and the background region to differentiate the boundaries of the lesion regions.

Limitations of Polyper. We show some failure cases of Polyper in Fig. 7. First, we assume polyp localization is accurate and cannot handle false positives or false negatives well. Second, we employ the fixed-width method to define the boundary width and extract boundary regions. This may not account for the diversity of polyp features. In the future, we will explore adaptive boundary width methods.

Conclusion

We present Polyper, a novel approach for polyp segmentation. We employ morphology operators to delineate boundary and interior polyp regions from the initial segmentation results. Then, we leverage the features of the interior polyp regions to enhance the features of boundary regions. Our experiments on five datasets demonstrate the remarkable performance of Polyper.

Acknowledgments

This research was supported by NSFC (No. 62276145), the Fundamental Research Funds for the Central Universities (Nankai University, 070-63223049), CAST through Young Elite Scientist Sponsorship Program (No. YESS20210377). Computations were supported by the Supercomputing Center of Nankai University (NKSC).

References

- Alom, M. Z.; Hasan, M.; Yakopcic, C.; Taha, T. M.; and Asari, V. K. 2018. Recurrent residual convolutional neural network based on u-net (r2u-net) for medical image segmentation. *arXiv preprint arXiv:1802.06955*.
- Antonelli, G.; Badalamenti, M.; Hassan, C.; and Repici, A. 2021. Impact of artificial intelligence on colorectal polyp detection. *Best Practice & Research Clinical Gastroenterology*, 52: 101713.
- Chen, F.; Ma, H.; and Zhang, W. 2023. SegT: A Novel Separated Edge-guidance Transformer Network for Polyp Segmentation. *arXiv preprint arXiv:2306.10773*.
- Chen, S.; Tan, X.; Wang, B.; and Hu, X. 2018. Reverse attention for salient object detection. In *Proceedings of the European conference on computer vision (ECCV)*, 234–250.
- Djinbachian, R.; Iratni, R.; Durand, M.; Marques, P.; and von Renteln, D. 2020. Rates of incomplete resection of 1- to 20-mm colorectal polyps: a systematic review and meta-analysis. *Gastroenterology*, 159(3): 904–914.
- Dong, B.; Wang, W.; Fan, D.-P.; Li, J.; Fu, H.; and Shao, L. 2021. Polyp-pvt: Polyp segmentation with pyramid vision transformers. *arXiv preprint arXiv:2108.06932*.
- Fan, D.-P.; Ji, G.-P.; Zhou, T.; Chen, G.; Fu, H.; Shen, J.; and Shao, L. 2020. Pranel: Parallel reverse attention network for polyp segmentation. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part VI 23*, 263–273. Springer.
- Fan, K.; Wang, C.; Wang, Y.; Wang, C.; Yi, R.; and Ma, L. 2023. RFENet: Towards Reciprocal Feature Evolution for Glass Segmentation. *arXiv preprint arXiv:2307.06099*.
- Fang, Y.; Chen, C.; Yuan, Y.; and Tong, K.-y. 2019. Selective feature aggregation network with area-boundary constraints for polyp segmentation. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part I 22*, 302–310. Springer.
- Geng, Z.; Guo, M.-H.; Chen, H.; Li, X.; Wei, K.; and Lin, Z. 2021. Is attention better than matrix decomposition? *arXiv preprint arXiv:2109.04553*.
- Gonzales, R. C.; and Woods, P. 1987. *Digital image processing*. Addison-Wesley Longman Publishing Co., Inc.
- He, H.; Li, X.; Cheng, G.; Shi, J.; Tong, Y.; Meng, G.; Prinnet, V.; and Weng, L. 2021. Enhanced boundary learning for glass-like object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15859–15868.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Huang, C.-H.; Wu, H.-Y.; and Lin, Y.-L. 2021. Hardnet-mseg: A simple encoder-decoder polyp segmentation neural network that achieves over 0.9 mean dice and 86 fps. *arXiv preprint arXiv:2101.07172*.
- Huang, S.; Li, J.; Xiao, Y.; Shen, N.; and Xu, T. 2022. RT-Net: relation transformer network for diabetic retinopathy multi-lesion segmentation. *IEEE Transactions on Medical Imaging*, 41(6): 1596–1607.
- Khan, A. M.; Ashrafee, A.; Khan, F. S.; Hasan, M. B.; and Kabir, M. H. 2023. AttResDU-Net: Medical Image Segmentation Using Attention-based Residual Double U-Net. *arXiv preprint arXiv:2306.14255*.
- Komodakis, N.; and Zagoruyko, S. 2017. Paying more attention to attention: improving the performance of convolutional neural networks via attention transfer. In *ICLR*.
- Li, S.; Sui, X.; Luo, X.; Xu, X.; Liu, Y.; and Goh, R. 2021. Medical image segmentation using squeeze-and-expansion transformers. *arXiv preprint arXiv:2105.09511*.
- Lin, A.; Chen, B.; Xu, J.; Zhang, Z.; Lu, G.; and Zhang, D. 2022. Ds-transunet: Dual swin transformer u-net for medical image segmentation. *IEEE Transactions on Instrumentation and Measurement*, 71: 1–15.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125.
- Liu, G.; Chen, Z.; Liu, D.; Chang, B.; and Dou, Z. 2023. FTMF-Net: A Fourier Transform-Multiscale Feature Fusion Network For Segmentation Of Small Polyp Objects. *IEEE Transactions on Instrumentation and Measurement*.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10012–10022.
- Lou, A.; Guan, S.; Ko, H.; and Loew, M. H. 2022. CaraNet: context axial reverse attention network for segmentation of small medical objects. In *Medical Imaging 2022: Image Processing*, volume 12032, 81–92. SPIE.
- Oktay, O.; Schlemper, J.; Folgoc, L. L.; Lee, M.; Heinrich, M.; Misawa, K.; Mori, K.; McDonagh, S.; Hammerla, N. Y.; Kainz, B.; et al. 2018. Attention u-net: Learning where to look for the pancreas. *arXiv preprint arXiv:1804.03999*.
- Park, K.-B.; and Lee, J. Y. 2022. SwinE-Net: hybrid deep learning approach to novel polyp segmentation using convolutional neural network and Swin Transformer. *Journal of Computational Design and Engineering*, 9(2): 616–632.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Pooler, B. D.; Kim, D. H.; Matkowskyj, K. A.; Newton, M. A.; Halberg, R. B.; Grady, W. M.; Hassan, C.; and Pickhardt, P. J. 2023. Growth rates and histopathological outcomes of small (6–9 mm) colorectal polyps based on CT colonography surveillance and endoscopic removal. *Gut*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 234–241. Springer.

- Sun, D.; Jiang, S.; and Qi, L. 2023. Edge-Aware Mirror Network for Camouflaged Object Detection. *arXiv preprint arXiv:2307.03932*.
- Thanh Duc, N.; Oanh, N. T.; Thuy, N. T.; Triet, T. M.; and Viet Sang, D. 2022. ColonFormer: An Efficient Transformer based Method for Colon Polyp Segmentation. *arXiv e-prints*, arXiv-2205.
- Valanarasu, J. M. J.; Sindagi, V. A.; Hacihaliloglu, I.; and Patel, V. M. 2020. Kiu-net: Towards accurate segmentation of biomedical images using over-complete representations. In *International conference on Medical Image Computing and Computer-Assisted Intervention*, 363–373. Springer.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7794–7803.
- Wang, Y.; Li, Z.; Mei, J.; Wei, Z.; Liu, L.; Wang, C.; Sang, S.; Yuille, A.; Xie, C.; and Zhou, Y. 2023. SwinMM: Masked Multi-view with Swin Transformers for 3D Medical Image Segmentation. *arXiv preprint arXiv:2307.12591*.
- Wei, J.; Hu, Y.; Zhang, R.; Li, Z.; Zhou, S. K.; and Cui, S. 2021. Shallow attention network for polyp segmentation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part I 24*, 699–708. Springer.
- Wu, H.; Zhong, J.; Wang, W.; Wen, Z.; and Qin, J. 2021. Precise yet efficient semantic calibration and refinement in convnets for real-time polyp segmentation from colonoscopy videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 2916–2924.
- Xiao, X.; Lian, S.; Luo, Z.; and Li, S. 2018. Weighted resunet for high-quality retina vessel segmentation. In *International Conference on Information Technology in Medicine and Education (ITME)*, 327–331. IEEE.
- Xie, E.; Wang, W.; Wang, W.; Ding, M.; Shen, C.; and Luo, P. 2020. Segmenting transparent objects in the wild. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*, 696–711. Springer.
- Xie, E.; Wang, W.; Yu, Z.; Anandkumar, A.; Alvarez, J. M.; and Luo, P. 2021. SegFormer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34: 12077–12090.
- Yin, B.; Zhang, X.; Hou, Q.; Sun, B.-Y.; Fan, D.-P.; and Van Gool, L. 2022. Camoformer: Masked separable attention for camouflaged object detection. *arXiv preprint arXiv:2212.06570*.
- Yu, Z.; and Han, S. 2023. 3D Medical Image Segmentation based on multi-scale MPU-Net. *arXiv preprint arXiv:2307.05799*.
- Zamir, S. W.; Arora, A.; Khan, S.; Hayat, M.; Khan, F. S.; and Yang, M.-H. 2022. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5728–5739.
- Zhang, R.; Li, G.; Li, Z.; Cui, S.; Qian, D.; and Yu, Y. 2020. Adaptive context selection for polyp segmentation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part VI 23*, 253–262. Springer.
- Zhang, Y.; Liu, H.; and Hu, Q. 2021. Transfuse: Fusing transformers and cnns for medical image segmentation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part I 24*, 14–24. Springer.
- Zhou, Z.; Rahman Siddiquee, M. M.; Tajbakhsh, N.; and Liang, J. 2018. Unet++: A nested u-net architecture for medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 3–11. Springer.
- Zhu, Y.; Qiao, Y.; and Yang, X. 2023. The optimal connection model for blood vessels segmentation and the MEA-Net. *arXiv preprint arXiv:2306.01808*.