



Adapting Segment Anything Model to Melanoma Segmentation in Microscopy Slide Images

Qingyuan Liu¹  and Avideh Zakhori¹ 

University of California, Berkeley
{pixelled, avz}@berkeley.edu

Abstract. Melanoma segmentation in Whole Slide Images (WSIs) is useful for prognosis and the measurement of crucial prognostic factors such as Breslow depth and primary invasive tumor size. In this paper, we present a novel approach that uses the Segment Anything Model (SAM) for automatic melanoma segmentation in microscopy slide images. Our method employs an initial semantic segmentation model to generate preliminary segmentation masks that are then used to prompt SAM. We design a dynamic prompting strategy that uses a combination of centroid and grid prompts to achieve optimal coverage of the super high-resolution slide images while maintaining the quality of generated prompts. To optimize for invasive melanoma segmentation, we further refine the prompt generation process by implementing in-situ melanoma detection and low-confidence region filtering. We select Segformer as the initial segmentation model and EfficientSAM as the segment anything model for parameter-efficient fine-tuning. Our experimental results demonstrate that this approach not only surpasses other state-of-the-art melanoma segmentation methods but also significantly outperforms the baseline Segformer by 9.1% in terms of IoU.

Keywords: Segment Anything Model, Melanoma Segmentation, Whole Slide Images

1 Introduction

Melanoma, one of the most serious forms of skin cancer, originates in melanocytes, the pigment-producing cells responsible for melanin production [22]. Based on its progression and location within the skin, melanoma can be categorized into two main types: in-situ melanoma and invasive melanoma. While in-situ melanoma represents cancerous melanocytes that are confined to the epidermis, invasive melanoma penetrates beyond the epidermis into the dermis, posing a significant risk of spreading to other vital organs. As invasive melanoma grows, it may invade blood vessels and lymphatic vessels, allowing cancer cells to detach from the primary tumor and cause metastatic cancer [1]. In this paper, we primarily focus on segmenting invasive melanoma to assist dermatologists in deciding on treatment options.

Early detection and accurate diagnosis of melanoma are crucial for improving the survival rate [3, 6]. The standard diagnosis practice begins with an initial

examination of dermatoscopic features to determine the types of melanocytic lesions. For suspicious and malignant lesions, a histopathologic analysis of skin biopsies stained with hematoxylin and eosin (H&E) is required [11]. While traditional approach requires examination of tissue specimens under a microscope, the advent of Whole Slide Images (WSIs) has revolutionized this process by digitizing tissue samples into high resolution images. These digital slide images enable pathologists to examine various characteristics such as the cellular architecture, breslow depth and complex histologic features that are crucial for determining the stage and aggressiveness of melanoma.

Recent studies [2, 16–18, 21, 24] have focused on utilizing deep learning technology for melanoma segmentation in whole slide images. Melanoma segmentation results with sufficient accuracy in slide images have proven highly beneficial for aiding diagnosis and assisting manual measurement of breslow depth and primary invasive tumor size, which are crucial prognostic factors. This shows the potential for a fully automatic diagnosis procedure. Phillips et al. [18] demonstrated the effectiveness of multi-scale FCN in segmentation the dermis, epidermis and tumor in whole slide images. Want et al. [26] utilized more advanced transformer models including Hierarchical Pyramid Transformers (HIPT) and Segformers to achieve accurate segmentation of invasive melanoma and the epidermis in microscopy slide images.

In this paper, we propose a novel method to use the Segment Anything Model (SAM) [13] for automatic melanoma segmentation in microscopy slide images. SAM has demonstrated great success in various computer vision tasks and has achieved state-of-the-art performance in a diverse range of image segmentation tasks [4, 13, 20]. One of the key features of SAM is a prompt encoder that allows the model to adapt to diverse downstream segmentation tasks with prompt engineering. Despite SAM’s strong zero-shot generalization abilities, several studies [8, 12, 19, 32] have shown that its accuracy is limited in segmentation tasks that require specific domain knowledge. Recent studies [15, 27, 31] have proposed methods to adapt SAM for medical image segmentation, such as augmenting data with SAM’s predictions or fine-tuning SAM for better performance. Despite showing impressive performance in segmenting medical images such as CT and MRI scans, these methods do not generalize well to microscopy images, especially for melanoma segmentation in microscopy slide images.

Our proposed method addresses these challenges by introducing an innovative framework that automatically generates prompts from an initial segmentation maps to guide SAM. We use Segformer, an effective semantic segmentation model for melanoma segmentation, to generate the initial segmentation map. We design a dynamic prompt strategy that uses a combination of centroid and grid prompts. Additionally, we incorporate in-situ melanoma detection and low-confidence region filtering to ensure precise prompt generation. Our experimental results demonstrate that this approach not only surpasses other state-of-the-art melanoma segmentation methods but also significantly improves upon the baseline performance of Segformer by over 9.1%.

2 Related Work

SAM in Medical Imaging. Given SAM’s impressive results on various natural image segmentation tasks [4, 13, 20], recent works have explored its application to medical image segmentation. However, several studies [8, 12, 19, 32] have shown that SAM underperforms in medical image segmentation. This is attributed to the model’s lack of domain-specific medical knowledge, the uncertain and complex object boundaries, intricate structures, and the wide-range of scales unique to medical objects [12]. Recent efforts have focused on adapting SAM for medical images, primarily through fine-tuning or adapting SAM to labeled medical dataset. Ma et al. [15] proposed to fine-tune SAM fully on labeled medical data, which is not cost-effective due to the vast number of parameters. Wu et al. [27] proposed to integrate adapter modules into SAM, allowing efficient fine-tuning by freezing all modules except for the adapters during training. These studies all focus on utilizing SAM as an interactive model and evaluate its performance by providing accurate prompts based on the ground-truth annotations. For microscopy slide image segmentation, SAM’s interactive feature can assist annotators but cannot fully automate the labor-intensive process of segmenting high-resolution whole slide images. Although SAM offers automatic mask generation, this feature is designed for segmenting everything on an entire image rather than targeting specific small objects, such as scattered melanoma cells in microscopy slide images. Zhang et al. [31] proposed to enhance medical images by adding semantic structures using SAM’s automatic mask generation. This approach combines generated masks, features and stability scores to help train other image segmentation models with enhanced data. The success of this method depends on SAM’s ability to generate useful structural information during the automatic mask generation process.

Melanoma Segmentation in Slide Images. Melanoma semantic segmentation can be classified into two categories: 1) skin lesion segmentation based on images captured at a macroscopic scale; 2) microscopy slide image segmentation that involves WSIs captured at a microscopic scale. While skin lesion segmentation only segments a large rounded melanoma component at a macroscopic level, microscopy slide image segmentation requires segmenting irregular melanoma scattered across high-resolution, detailed views of tissue samples at the cellular level, making the task significantly more challenging. Our work focuses exclusively on the latter.

Most studies [2, 16–18, 21, 24] focus on utilizing CNN-based methods for melanoma segmentation in microscopy slide images. Phillips et al. [18] devised a multi-stride fully convolutional network (FCN) that can effectively segment tumour, epidermis and dermis. Shah et al. [21] developed a two-stage method to segment invasive melanoma by leveraging the difference between in-situ melanoma and invasive melanoma. They used HRNet-OCR [30] and HookNet [23] as backbones and trained two models, one for epidermis segmentation and one for tumor segmentation, to predict two separate segmentation masks for epidermis and melanoma. By removing all predicted melanoma from the epidermis mask, they obtained segmentation masks for invasive melanoma.




Fig. 1: An overview of the proposed method. The initial mask X generated by Segformer is post-processed to generate the mask \hat{X}^m . We run SAM on the prompts generated from \hat{X}^m to generate the mask \hat{Y}^m . The two masks are combined to create the final mask Y .

Recent works have applied transformer-based architectures [25] to melanoma segmentation. Wang et al. [26] proposed to fine-tune the Hierarchical Pyramid Transformer (HIPT) [5] and the Segformer [28] for invasive melanoma segmentation in microscopy slide images. Segformer has been shown to outperform all other CNN-based methods and HIPT in segmenting invasive melanoma, achieving state-of-the-art IoU [26]. Our method is designed with Segformer [28] as the initial segmentation model to achieve the best results.

3 Method

3.1 Overview of the Proposed Method

Our proposed method comprises an initial semantic segmentation model and a segment anything model as illustrated in Fig. 1. To segment invasive melanoma automatically, we prompt SAM with prompts generated from the mask produced by the initial segmentation model. To optimize for our task of segmenting invasive melanoma, we select Segformer as the initial segmentation model due to its superior performance compared to other models [26]. We choose EfficientSAM [29] as the segment anything model for parameter-efficient fine-tuning. The entire framework is described as follows.

Step 1. Initial Mask Generation. We run Segformer [28] to generate the initial segmentation mask X . Then we perform in-situ melanoma detection to separate X into the estimated in-situ melanoma regions X^s and the remaining invasive melanoma regions X^v . We filter out low-confidence regions from X^s to obtain the filtered mask X_c^s and combine it with X^v to obtain the post-processed mask \hat{X}^m . The details are described in Secs. 3.2 and 3.3.

Step 2. Prompt Generation. We generate single point prompts from the post-processed mask \hat{X}^m . We determine the best prompt type for each connected

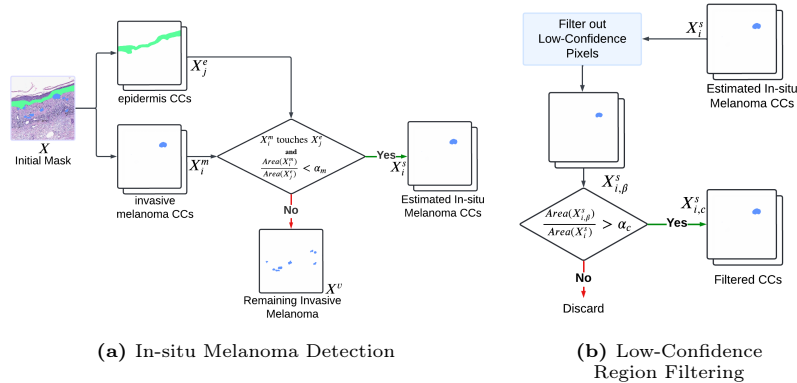


Fig. 2: In-situ melanoma detection finds the estimated in-situ melanoma regions X^s from the initial mask X . Low-confidence region filtering discards low-confidence connected components from X^s .

component based on its shape distribution and geometric characteristics. The details of the prompt generation strategy are described in Sec. 3.4.

Step 3. Final Mask Generation. We run SAM on the generated prompts to produce its own invasive melanoma mask \hat{Y}^m . We use SAM’s mask to refine the post-processed mask \hat{X}^m by combining the two masks together. This aims to enhance the overall accuracy and robustness of melanoma segmentation. The details are described in Sec. 3.5.

3.2 In-situ Melanoma Detection

We exclude low-confidence invasive melanoma predictions that touch the epidermis, since these are considered to be in-situ melanoma from a medical perspective. This process is shown in Fig. 2a. The initial mask X produced by Segformer contains both the epidermis mask X^e , and the invasive melanoma mask X^m . The invasive melanoma mask can be represented as a union of invasive melanoma connected components (CCs) $X^m = \bigcup X_i^m$, where i denotes the i th invasive melanoma connected component. Similarly, the mask for epidermis can be represented as $X^e = \bigcup X_j^e$, where j denotes the j th epidermis connected component. Since SAM heavily relies on prompts to specify the exact objects to segment, inaccurate and ambiguous prompts can significantly degrade its performance. This can lead SAM to produce more false positives compared to the initial segmentation mask from which prompts are generated. Therefore, we filter out low-confidence invasive melanoma components to improve the accuracy of prompts so that it is clicked on true positives in most cases.

To begin with, we iterate over all connected components X_i^m in the invasive melanoma mask X^m . For each connected component, we determine whether it touches any epidermis component X_j^e . If such a touch exists, we compute the ratio between the area of the melanoma component X_i^m and the area of the

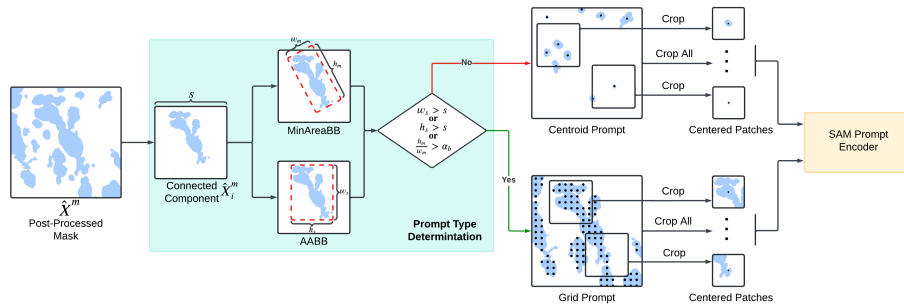


Fig. 3: Prompt Generation Strategy. We apply the centroid prompt and the grid prompt to based on the geometric characteristics of each connected component.

touched epidermis region X_j^e . If this ratio falls below a specified threshold α_m , we classify it as an estimated in-situ melanoma region. This is because careful inspection of Segformer’s predictions shows that in-situ melanoma that are misclassified as invasive melanoma tend to have a relatively small size compared to the epidermis. This process results in the estimated in-situ melanoma mask X^s and the remaining invasive melanoma regions X^v .

3.3 Low-Confidence Region Filtering

We further filter out low-confidence estimated in-situ melanoma regions from X^s . We keep high-confidence regions even if they touch the epidermis, since this could result from false predictions in the epidermis that make them touch each other. We determine the confidence level based on the probability map generated by Segformer. As shown in Fig. 2b, for each component X_i^s , we first find its high-confidence sub-component $X_{i,\beta}^s$ that has a probability exceeding the defined threshold β :

$$X_{i,\beta}^s = \{x \mid P(x) > \beta, x \in X_i^s\}. \quad (1)$$

We then determine the confidence level by computing the ratio between the area of $X_{i,\beta}^s$ and X_i^s . If the ratio fails to meet the confidence threshold α_c , we discard this component and do not generate prompts from it. Careful inspections of Segformer’s mask show that it generally predicts a larger proportion of low-probability areas for in-situ melanoma compared to invasive melanoma. Therefore, if we keep only regions with a high proportion of high probability areas, we can significantly reduce the number of false positives in X^m . Next, we combine the filtered melanoma regions X_c^s and the rest invasive melanoma regions X^v to obtain the post-processed invasive melanoma mask \hat{X}^m .

3.4 Prompt Generation

Melanoma can be in various forms in microscopy slide images, ranging from rounded clusters to jagged streaks or any combination of irregular shapes. To

Algorithm 1 Determine prompt type for a connected component

Require:

- A connected component of invasive melanoma: x
- The length of the side of a patch: s
- The threshold for the ratio between sides of a minimum bounding box: α_b

Ensure:

The generated prompts (a set of single point prompts): P

```

1: function GENERATEPROMPTS( $x$ )
2:    $P = \emptyset$ 
3:    $w_s, h_s = \text{FINDAABB}(x)$ 
4:    $w_m, h_m = \text{FINDMINAREABB}(x)$ 
5:   if  $w_s > s$  or  $h_s > s$  or  $\frac{h_m}{w_m} > \alpha_b$  then
6:      $P_{grid} = \text{GENERATEGRIDPROMPTS}(x)$ 
7:      $P = P \cup P_{grid}$ 
8:      $center = \text{FINDCENTROID}(x)$ 
9:     if  $center \in x$  then
10:        $P = P \cup \{center\}$ 
11:     end if
12:   else
13:      $center = \text{FINDCENTROID}(x)$ 
14:      $P = P \cup \{center\}$ 
15:   end if
16:   return  $P$ 
17: end function

```

facilitate accurate segmentation of diverse morphologies, we design two types of point prompts: centroid and grid. Centroid prompts consist of a single point placed at the centroid of a connected component. On the other hand, grid prompts comprise a grid of points distributed within a connected component. As shown in Fig. 3, we choose a prompt type for each connected component \hat{X}_i^m and serve each point as a single point prompt for the SAM prompt encoder. We feed a patch centered at the point to the SAM image encoder, since this ideally provides maximum context.

We choose prompts dynamically based on the characteristics of the connected components. Centroid prompts excel in providing optimal context for small, regularly shaped melanoma, but often fail to provide sufficient context for those covering large regions and those with irregular shapes. In cases where a connected component is too large to fit inside a patch, numerous melanoma predictions are lost in a centered patch and not enough context is provided for its surroundings. Furthermore, if the centroid falls outside the component, it turns out to be a point prompt not clicked on invasive melanoma, resulting in false positives.

For each connected component \hat{X}_i^m , we determine the prompt type based on its shape distribution and geometric attributes, as detailed in Algorithm 1. As shown in Fig. 3, we find the width w_s and height h_s of an axis aligned bounding box (AABB). If either dimension exceeds the patch length, we use grid prompts. Then we find the width w_m and height h_m of an arbitrarily oriented bounding box with minimum area and check whether the ratio between h_m and w_m exceeds a threshold α_b assuming h_m is the longer side. A high ratio implies that the shape of the component resembles a narrow rectangle or any other irregular shape where one dimension is larger than the other, in which case centroid prompt is ambiguous for segmenting the whole component. In such cases, we also choose to use grid prompts. Additionally, we include its centroid as part of the prompt if it is inside the component as a supplementary to the grid prompts. For all other cases, we use centroid prompt only.

If a connected component that meets the criteria for grid prompts has more than $\frac{1}{4}$ of the patch area, we apply grid prompts to all connected components in the slide image instead of evaluating each one individually. This is because neighboring melanomas often share similar shapes, even if they don't fit a single morphology. When a large melanoma with an irregular shape is present, nearby melanomas usually have similar shapes, making grid prompts the most effective for covering all components.

3.5 Final Mask Generation

After deriving the prompts, we run SAM to generate its own mask \hat{Y}^m as shown in Fig. 1. To obtain the final mask, we combine the post-processed Segformer's mask \hat{X}^m and SAM's mask \hat{Y}^m by performing a union between the two:

$$Y = \hat{Y}^m \cup \hat{X}^m \quad (2)$$

We choose to combine these two masks because single point prompts are sparse and may not comprehensively cover all regions. This sparsity can lead to ambiguity and result in missing some high-confidence predictions from Segformer. Therefore, we choose to include the post-processed mask \hat{X}^m to ensure that high-confidence regions are kept in the final mask Y .

3.6 Training Strategy

We devise a training strategy involving two stages that optimizes SAM for the way we use prompts. Since we only use single point prompts, we first train SAM on single point prompts for melanoma. For each patch sampled from a microscopy slide image, we first find all connected components and for each connected component we generate a single point prompt clicked at a random position inside it. This enables the model to learn to segment melanoma anywhere in a patch. After the initial training on random point prompts, we fine-tune our model on patches centered at each individual single point prompt. Each patch might contain multiple components, but we only set the component containing the point prompt as the target to reduce ambiguity. This allows the model to optimize for inference using our SAM-based method.

4 Experiments

4.1 Dataset

Our dataset comprises 101 microscopy slide images, with sizes ranging from 23700×21199 pixels to 1996×1679 pixels. These images are derived from skin biopsies stained with H&E [10] and captured under a microscope at 40x magnification. Detailed Annotations at this magnification level are provided by an expert dermatopathologist. Each image in the dataset is carefully selected and cropped from the original WSIs by the dermatopathologist. This is because whole slide images often contain multiple focal planes of the same tissue, leading to redundancy in the data. By selecting the most informative regions at a single focal plane, we reduce the redundancy and ensure that the dataset focuses on the most relevant tissue structures. The annotations for our microscopy images include seven classes: air, background cells, epidermis, invasive melanoma, inflamed tumor, fibrotic tumor, and uncertain tumor. It is important to note that invasive melanoma is the only type of melanoma precisely annotated. In-situ melanoma is labeled as epidermis in our dataset due to its confinement to the epidermis and due to our focus on segmenting invasive melanoma.

Segformer Dataset Generation. To generate the dataset for Segformer, We follow the approach described by Wang et al [26]. Given the extremely high resolution of microscopy slide images, we divide each slide into non-overlapping patches of 512×512 and 1024×1024 pixels. We re-categorize the original annotations into three distinct classes: invasive melanoma, epidermis, and others. This reclassification is essential since the ambiguous boundaries of fibrotic and inflamed tumor make them less suitable for effective segmentation model training. We under-sample the background by discarding patches with 97% of background cells and air to address class imbalance issues. The resulting dataset comprises 14885 patches of around 3.9 billion pixels for 512×512 resolution and 4326 patches of around 4.5 billion pixels for 1024×1024 resolution. This pre-processing approach ensures a balanced and representative dataset for training the Segformer model.

SAM Dataset Generation. We generate the SAM dataset corresponding to the two-stage training process described in Sec. 3.6. For the first stage, we divide each microscopy image into non-overlapping patches and generate a single point prompt within each connected component randomly. The ground truth for each prompt is the mask corresponding solely to that particular connected component, rather than all components in the patch. This aims to minimize the ambiguity of single point prompts as much as possible and facilitates model convergence. For the second stage, we use both the centroid and randomly sampled points as the training prompts. We use patches centered at each prompt to optimize SAM specifically for our method. This stage’s dataset comprises 8357 patches of 512×512 pixel resolution and 6170 patches of 1024×1024 pixel resolution.

4.2 Implementation Details

Model Settings. We use Segformer B0 and B1 [28] as the initial segmentation model and EfficientSAM-S [29] as the segment anything model in our approach. We choose EfficientSAM-S since it is the most efficient variant of SAM that reconstructs the image embeddings of ViT-H [9] in the original SAM. To further improve training efficiency, we integrate adapters into the ViT [9] image encoder following the approach described in Med-SA [27]. We use adapters with an input and output dimensionality of $d_a = 768$ and set the dimensionality of hidden layers to $d_h = 1024$. With adapters, we only fine-tune 6.7M parameters for EfficientSAM-S with a total of 29.3M parameters.

Training. For Segformer, we use AdamW [14] optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.999$, and a weight decay of 0.01. We use an initial learning rate of $5e - 4$ and a polynomial decay scheduler. We use the weights of Segformer pretrained on ImageNet [7] as our starting point and train the model for 150 epochs. For EfficientSAM, we use the Adam optimizer with an initial learning rate of $1e - 4$, an exponential decay rate of $\beta_1 = 0.9$ and $\beta_2 = 0.999$, and a weight decay of 0.05. We use the initial learning rate for the first 10 epochs and apply a learning rate decay factor of 0.5 every 10 epochs. We use a batch size of 8 for patch resolution 512×512 , and a batch size of 6 for patch resolution 1024×1024 . The model is first trained for 100 epochs on the first-stage dataset and then trained for 150 epochs on the second-stage dataset. The model is trained with a binary cross entropy loss function with equal weights for the invasive melanoma and the backgrounds. All experiments are implemented in PyTorch and executed on 4 NVIDIA Quadro RTX 8000 GPUs.

Inference. To generate the segmentation mask with Segformer, we process each microscopy slide image into patches. Specifically, we use a sliding window to create patches by shifting the window both horizontally and vertically with a step size of 128 pixels. In addition, we apply a 2D Gaussian kernel as a weighting mechanism for each pixel within a patch. The kernel has the same size as the patch and a standard deviation of $\frac{1}{4}$ of the patch’s side length. To generate the segmentation mask for EfficientSAM, we run our method with a fixed set of hyperparameters. We set the threshold for determining in-situ melanoma $\alpha_m = 0.1$, the probability threshold for high-confidence regions $\beta = 0.8$, and the threshold for excluding low-confidence regions $\alpha_c = 0.4$. Additionally, the threshold for the ratio between the sides of a minimum bounding box α_b , as described in Algorithm 1, is set to 3. The grid prompt employs a vertical and horizontal gap of 64 pixels between neighboring points.

4.3 Results

To evaluate the effectiveness of our proposed method, we compare it with state-of-the-art melanoma segmentation methods on our dataset. As shown in Tab. 1, we compare our method with melanoma segmentation methods including Multi-Scale FCN [18], HRNet & OCR [21], HIPT [26], and Segformer [28]. The results for Segformer B0 and B1 are reproduced following the methods described in

Table 1: The results of invasive melanoma segmentation on our dataset. We compare our proposed methods with other state-of-the-art segmentation methods.

Model	Resolution	IoU (%)	F1 (%)
Multi-Scale FCN [18]	512	13.0	14.0
HRNet & OCR [21]	788	29.1	44.0
HIPT [26]	512	40.1	57.3
Seg. B0 [28]	512	44.6	61.6
Seg. B1 [28]	512	42.0	59.2
Seg. B0 & EfficientSAM-S (ours)	512	54.1	70.2
Seg. B1 & EfficientSAM-S (ours)	512	47.5	64.4
HIPT [26]	1024	33.0	46.0
Seg. B0 [28]	1024	44.0	61.1
Seg. B1 [28]	1024	45.0	62.0
Seg. B0 & EfficientSAM-S (ours)	1024	49.5	66.2
Seg. B1 & EfficientSAM-S (ours)	1024	49.4	66.1

Wang et al [26]. All EfficientSAM methods shown in the table use a 1024×1024 patch size, which is different from the resolution field in Tab. 1 that represents the side length of a patch used for the initial segmentation model.

As shown in Tab. 1, our method outperforms all other baselines in terms of IoU and F1 score. Compared to using solely the Segformer, which is also used as the initial segmentation model in our method, our method achieves a gain of 9.1% IoU and 8.2% F1 over the state-of-the-art segmentation method. Fig. 4 presents representative qualitative results on our dataset, where our method significantly reduces errors in in-situ melanoma regions and enhances segmentation accuracy in areas where distinguishing between invasive melanoma and the epidermis is particularly challenging. Additionally, the best result is achieved with using EfficientSAM and Segformer B0 with 512×512 patches, even though this is not the best performing Segformer in terms of IoU and F1 score. A higher IoU for the initial segmentation mask does not necessarily indicate a higher improvement in IoU in the final mask of our method. The amount of improvement depends on the quality of sampled prompts and the room left for improvement for each prompt.

4.4 Ablation Studies

We now analyze our method through a series of ablation studies on each component of the framework.

In-situ Melanoma Detection We investigate the impact of in-situ melanoma detection described in Sec. 3.2 by filtering out all detected estimated in-situ melanoma regions regardless of their confidence levels. Tab. 2 shows that this consistently improves performance, with IoU gains ranging from 0.5% to 11.8%. The improvement varies based on the quality of the initial segmentation mask




Fig. 4: Qualitative results on our dataset. Red denotes invasive melanoma and green denotes the epidermis. Compared to Segformer, our method significantly improves the accuracy of predictions for invasive melanoma regions, especially in areas where distinguishing from the epidermis is challenging.

and the room left for improvement. In contrast, disabling in-situ melanoma detection results in a performance drop compared to the original mask produced by Segformer. This drop occurs because prompts generated from incorrect predictions lead SAM to make additional errors. Therefore, enabling in-situ melanoma detection not only mitigates errors from incorrect initial predictions but also leverages the strengths of the remaining high-confidence predictions to improve the final segmentation results.

Low-Confidence Region Filtering We study the impact of low-confidence region filtering presented in Sec. 3.3. As shown in Tab. 2, our method achieves a 3.9% IoU gain for Segformer B0 and a 4.6% IoU gain for Segformer B1 when using a 1024×1024 patch size. However, there is a slight drop of 0.1% in IoU for Segformer B0 with a 512×512 patch size. This suggests that some high-confidence invasive melanoma regions touch the epidermis, which are not filtered by the algorithm, are actually in-situ melanoma. Overall, this demonstrates that low-confidence region filtering improves the robustness of our segmentation results.

Table 2: Ablation study on in-situ melanoma detection and low-confidence region filtering in preparation for prompt generation, as described in Sec. 3.2 and Sec. 3.3.

Initial Model	Resolution	In-situ Melanoma Detection	Low-Confidence Region Filtering	IoU (%)
Seg. B0	512			42.4
Seg. B0	512	✓		54.2
Seg. B0	512	✓	✓	54.1
Seg. B0	1024			41.4
Seg. B0	1024	✓		45.6
Seg. B0	1024	✓	✓	49.5
Seg. B1	512			44.1
Seg. B1	512	✓		47.5
Seg. B1	512	✓	✓	47.5
Seg. B1	1024			44.3
Seg. B1	1024	✓		44.8
Seg. B1	1024	✓	✓	49.4

Prompt Types We study the effectiveness of different prompt types. We test prompts with centroid alone, grid alone and our proposed strategy that uses both as presented in Sec. 3.4. Tab. 3 shows that in most cases using both achieves the highest IoU in the final mask. Compared to grid prompts alone, using both achieves a gain as high as 3.7% IoU. Compared to centroid prompts alone, using both achieves gains up to 2.1% IoU with one case showing no improvement. It is noticeable that centroid prompts alone outperforms grid prompts alone in the final mask, but performs much worse in pre-merged mask. This shows that centroid prompts allow more accurate segmentation and thus complement the initial mask well, achieving high IoU in the final mask after merging. In contrast, grid prompts alone achieve full coverage over the initial mask, but not all points serve as effective prompts for accurate segmentation. This demonstrates that our method effectively leverages the strengths of both prompts by using grid prompts for large, irregular melanoma components and centroid prompts for small, regular melanoma components, maximizing the advantages of each prompt type.

Final Mask Generation We study the impact of final mask generation. As shown in Tab. 4, the final mask consistently achieves the highest IoU in all cases. EfficientSAM’s mask outperforms the post-processed mask in most cases, with the most significant improvement being a 5.5% increase for Segformer B0 using 512×512 patches. Merging two masks into a final mask significantly enhances accuracy. The post-processed mask shows gains ranging from 2.2% to 5.5% IoU, while the EfficientSAM’s mask gains improvements between 0.8% and 2.3% IoU. This demonstrates that EfficientSAM complements well the post-processed mask in generating the final mask. In addition, we evaluate the EfficientSAM’s performance with prompts generated from the ground truth. Its IoU performance is 10.8% higher than the best result of using EfficientSAM alone and 8.9% higher

Table 3: Ablation study on the effects of using different prompts. "Both" denotes the method that dynamically uses both centroid and grid prompts based on the shape distributions of melanoma components, as shown in Sec. 3.4.

Initial Segmentation Model	Prompt	Resolution	EfficientSAM Final Mask	
			IoU (%)	IoU (%)
Seg. B0	Centroid	512	43.3	53.1
Seg. B0	Grid	512	50.6	51.7
Seg. B0	Both	512	52.2	54.1
Seg. B1	Centroid	512	44.3	47.5
Seg. B1	Grid	512	46.4	46.5
Seg. B1	Both	512	46.1	47.5
Seg. B0	Centroid	1024	40.8	48.5
Seg. B0	Grid	1024	45.5	45.8
Seg. B0	Both	1024	48.7	49.5
Seg. B1	Centroid	1024	37.7	47.3
Seg. B1	Grid	1024	46.0	47.2
Seg. B1	Both	1024	47.1	49.4

Table 4: Ablation study on final mask generation. The last row represents the result when prompts are generated from the ground truth instead of the mask produced by Segformer.

Initial Segmentation Model	Resolution	Post-processed EfficientSAM Final		
		IoU (%)	IoU (%)	IoU (%)
Seg. B0	512	48.6	52.2	54.1
Seg. B1	512	43.4	46.1	47.5
Seg. B0	1024	46.7	48.7	49.5
Seg. B1	1024	47.2	47.1	49.4
GT	-	-	63.0	-

than the best final mask. This demonstrates the upper limit of the performance of our method when using fully accurate prompts.

5 Conclusion

We proposed a novel approach to explore the potential of SAM for melanoma segmentation in microscopy slide images. Our method utilizes Segformer to generate initial segmentation masks and subsequently prompts EfficientSAM using a dynamic selection of centroid and grid prompts for automatic invasive melanoma segmentation. Our experimental results demonstrate that this approach surpasses other state-of-the-art melanoma segmentation methods by a large margin of 9.1% in IoU.

Acknowledgements

We would like to express our gratitude to Michael Wang and Timothy McCalmont for providing the data, and to Shinwoo Choi for sharing valuable insights, without which this paper would not be possible.

References

1. Abelloff, M.D., Armitage, J.O., Niederhuber, J.E., Kastan, M.B., McKenna, W.G.: *Abelloff's Clinical Oncology E-Book*. Elsevier Health Sciences (2008)
2. Alheejawi, S., Mandal, M., Xu, H., Lu, C., Berendt, R., Jha, N.: Deep learning-based histopathological image analysis for automated detection and staging of melanoma. In: *Deep Learning Techniques for Biomedical and Health Informatics*, pp. 237–265. Elsevier (2020)
3. American Cancer Society: Cancer facts and figures 2024 (2024), <https://www.cancer.org/research/cancer-facts-statistics/all-cancer-facts-figures/2024-cancer-facts-figures.html>, accessed: 2024-06-08
4. Arbelaez, P., Maire, M., Fowlkes, C., Malik, J.: Contour detection and hierarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence* **33**(5), 898–916 (2010)
5. Chen, S., Ge, C., Tong, Z., Wang, J., Song, Y., Wang, J., Luo, P.: Adaptformer: Adapting vision transformers for scalable visual recognition. *Advances in Neural Information Processing Systems* **35**, 16664–16678 (2022)
6. Conic, R.Z., Cabrera, C.I., Khorana, A.A., Gastman, B.R.: Determination of the impact of melanoma surgical timing on survival using the national cancer database. *Journal of the American Academy of Dermatology* **78**(1), 40–46 (2018)
7. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *2009 IEEE conference on computer vision and pattern recognition*. pp. 248–255. Ieee (2009)
8. Deng, R., Cui, C., Liu, Q., Yao, T., Remedios, L.W., Bao, S., Landman, B.A., Wheless, L.E., Coburn, L.A., Wilson, K.T., et al.: Segment anything model (sam) for digital pathology: Assess zero-shot segmentation on whole slide imaging. *arXiv preprint arXiv:2304.04155* (2023)
9. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
10. Fischer, A.H., Jacobson, K.A., Rose, J., Zeller, R.: Hematoxylin and eosin staining of tissue and cell sections. *Cold spring harbor protocols* **2008**(5), pdb-prot4986 (2008)
11. Goodson, A.G., Grossman, D.: Strategies for early melanoma detection: Approaches to the patient with nevi. *Journal of the American Academy of Dermatology* **60**(5), 719–735 (2009)
12. Huang, Y., Yang, X., Liu, L., Zhou, H., Chang, A., Zhou, X., Chen, R., Yu, J., Chen, J., Chen, C., et al.: Segment anything model for medical images? *Medical Image Analysis* **92**, 103061 (2024)
13. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4015–4026 (2023)

14. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
15. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. *Nature Communications* **15**(1), 654 (2024)
16. Nofallah, S., Mokhtari, M., Wu, W., Mehta, S., Knezevich, S., May, C.J., Chang, O.H., Lee, A.C., Elmore, J.G., Shapiro, L.G.: Segmenting skin biopsy images with coarse and sparse annotations using u-net. *Journal of digital imaging* **35**(5), 1238–1249 (2022)
17. Oskal, K.R., Risdal, M., Janssen, E.A., Undersrud, E.S., Gulsrud, T.O.: A u-net based approach to epidermal tissue segmentation in whole slide histopathological images. *SN Applied Sciences* **1**, 1–12 (2019)
18. Phillips, A., Teo, I., Lang, J.: Segmentation of prognostic tissue structures in cutaneous melanoma using whole slide images. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops* (2019)
19. Roy, S., Wald, T., Koehler, G., Rokuss, M.R., Disch, N., Holzschuh, J., Zimmerer, D., Maier-Hein, K.H.: Sam. md: Zero-shot medical image segmentation capabilities of the segment anything model. arXiv preprint arXiv:2304.05396 (2023)
20. Van de Sande, K.E., Uijlings, J.R., Gevers, T., Smeulders, A.W.: Segmentation as selective search for object recognition. In: *2011 international conference on computer vision*. pp. 1879–1886. IEEE (2011)
21. Shah, A., Mehta, A., Wang, M., Neumann, N., Zakhor, A., McCalmont, T.: Deep learning segmentation of invasive melanoma. In: *Medical Imaging 2023: Digital and Computational Pathology*. vol. 12471, pp. 447–452. SPIE (2023)
22. Skin Cancer Foundation: Melanoma overview: A dangerous skin cancer (2024), <https://www.skincancer.org/skin-cancer-information/melanoma>, accessed: 2024-06-08
23. Van Rijthoven, M., Balkenhol, M., Siliņa, K., Van Der Laak, J., Ciompi, F.: Hooknet: Multi-resolution convolutional neural networks for semantic segmentation in histopathology whole-slide images. *Medical image analysis* **68**, 101890 (2021)
24. Van Zon, M., Stathonikos, N., Blokk, W.A., Komina, S., Maas, S.L., Pluim, J.P., Van Diest, P.J., Veta, M.: Segmentation and classification of melanoma and nevus in whole slide images. In: *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*. pp. 263–266. IEEE (2020)
25. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
26. Wang, F., Wang, M., Zakhor, A., McCalmont, T.: Transformers for microscopy slide image segmentation of invasive melanoma. *Electronic Imaging* **36**, 1–5 (2024)
27. Wu, J., Fu, R., Fang, H., Liu, Y., Wang, Z., Xu, Y., Jin, Y., Arbel, T.: Medical sam adapter: Adapting segment anything model for medical image segmentation. arXiv preprint arXiv:2304.12620 (2023)
28. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in neural information processing systems* **34**, 12077–12090 (2021)
29. Xiong, Y., Varadarajan, B., Wu, L., Xiang, X., Xiao, F., Zhu, C., Dai, X., Wang, D., Sun, F., Iandola, F., et al.: EfficientSAM: Leveraged masked image pretraining for efficient segment anything. arXiv preprint arXiv:2312.00863 (2023)
30. Yuan, Y., Chen, X., Wang, J.: Object-contextual representations for semantic segmentation. In: *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI* 16. pp. 173–190. Springer (2020)

31. Zhang, Y., Zhou, T., Wang, S., Liang, P., Zhang, Y., Chen, D.Z.: Input augmentation with sam: Boosting medical image segmentation with segmentation foundation model. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 129–139. Springer (2023)
32. Zhou, T., Zhang, Y., Zhou, Y., Wu, Y., Gong, C.: Can sam segment polyps? arXiv preprint arXiv:2304.07583 (2023)