Plug-and-Play PPO: An Adaptive Point Prompt Optimizer Making SAM Greater

Supplementary Material

In the supplementary materials, we first evaluate the runtime efficiency of PPO across different GPU devices. Next, we provide a qualitative analysis of the optimization effects on initial segmentation generated by CS. We then explore the optimization of different initial prompts generated by FM using various reference images. Finally, we investigate the necessity of negative prompts for SAM.

1. Runtime Efficiency Analysis

To further demonstrate the plug-and-play capability of PPO in optimizing initial prompts, we tested the entire workflow—including FM-based initial prompt generation, PPO optimization, and segmentation—on different devices. The results, shown in Table 1, reveal that PPO achieves very short optimization times for 224×224 images across five different types of GPU servers, with the fastest averaging just 0.1956 seconds per image.

Additionally, when using FM as the initial prompt generation paradigm, the entire segmentation workflow achieves target segmentation without additional training, relying solely on a one-shot reference image. This is accomplished through DRL, which effectively coordinates two foundational vision models. The process is highly efficient, approaching real-time segmentation, ensuring task transferability, and offering significant potential for practical applications.

Table 1. Runtime for different GPUs.

GPU	Runtime (Unit: Second)		
	Generate prompts	PPO	Segment
GeForce RTX 3080	0.2199	0.4698	0.5280
GeForce RTX 4090	0.0914	0.1956	0.2321
RTX A4500	0.3029	0.2800	0.6875
Tesla A40	0.3944	0.3106	1.0399
Tesla V100	0.2744	0.3206	0.6100

2. Further Analysis of CS-based Prompts

In Section 4.5, we quantitatively demonstrated that PPO enhances SAM's segmentation performance by optimizing initial prompts generated through various methods. In this section, we further illustrate the advantages of PPO through qualitative analysis. As shown in Figure 1, Figure 1(b) represents the results of coarse segmentation. Due to limited training data, the coarse segmentation results exhibit significant false positives and false negatives, particularly in challenging regions.

In Figure 1(c), prompt points sampled from the coarse segmentation are fed into SAM, partially mitigating the issue of false negatives. However, this approach also introduces new challenges. Since point sampling may occur in regions of false positives, these incorrect prompts can mislead SAM, resulting in erroneous segmentation outputs that deviate from the ground truth.

Figure 1(d) shows the results after optimizing the sampled points using PPO. By leveraging the heterogeneous graph constructed from the sampled prompts, PPO effectively identifies and removes incorrect prompts, significantly reducing their misleading effects on SAM. The optimized prompts enable SAM to produce segmentation results that align more closely with the ground truth, addressing both false positives and false negatives.

This qualitative analysis further highlights the robustness of PPO in optimizing various types of point prompt schemes, demonstrating its capability to enhance segmentation performance across diverse scenarios.

3. Different Initial Prompts Optimization

CS-based prompt generation primarily depends on the quality of the pseudo-masks produced by upstream segmentation, which limits the variability in the distribution of sampled prompts within the pseudo-mask. In contrast, FM-based prompt generation is influenced by the choice of reference images, which can significantly affect the quality of the generated prompts. To further validate the effectiveness of PPO, we generated different initial prompt schemes using FM with various reference images. The quantitative results are presented in Table 2, where each sample represents a prompt scheme generated using a different reference image for each dataset.

Additionally, Figures 2 and 3 present the qualitative results on two datasets, with the first column showing the selected reference images. These results demonstrate that different reference images produce varying prompt schemes, and the quality of the prompts is influenced by the similarity between the reference and target images. The initial prompts generated from different reference images lead to unstable segmentation results from SAM. However, after applying PPO, the segmentation performance improves significantly and becomes more consistent.

PPO effectively mitigates errors in the prompts, such as incorrect prompts and excessively dense prompts, by refining the prompt distribution. This optimization reduces the impact of reference-target differences and enhances the effectiveness of the prompts, ultimately improving SAM's

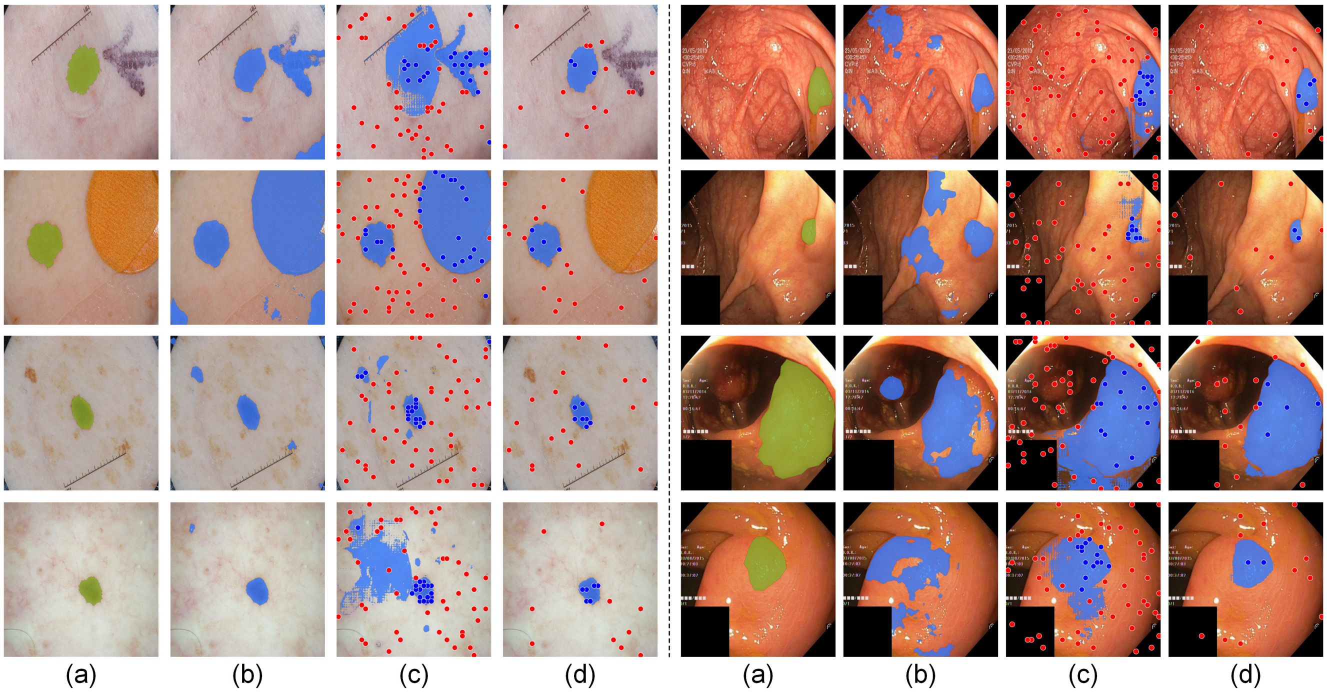

Figure 1. Qualitative segmentation results of PPO on the coarse segmentation sampling-based initial prompt scheme are as follows: (a) ground truth, (b) coarse segmentation results, (c) segmentation results with initial prompts, and (d) segmentation results with optimized prompts. PPO significantly improves SAM's segmentation accuracy by optimizing randomly sampled prompts in regions with coarse segmentation errors and reducing redundant prompts within the same area.

Table 2. Optimization performance of initial prompts generated from different reference images.

Reference	ISIC				Kvasir			
	Before PPO		After PPO		Before PPO		After PPO	
	mDSC (%↑)	mHD (↓)	mDSC (%↑)	mHD (↓)	mDSC (%↑)	mHD (↓)	mDSC (%↑)	mHD (↓)
Sample 1	70.9	8.3	78.0	8.0	30.4	15.7	66.5	11.5
Sample 2	68.4	9.1	80.6	7.9	48.6	14.2	72.0	10.5
Sample 3	72.0	8.5	77.6	8.4	55.6	12.6	68.1	10.2
Sample 4	73.3	8.1	75.4	8.1	43.1	14.7	75.4	9.9
Sample 5	75.4	8.2	78.7	8.0	34.1	15.0	77.6	9.3
Average	72.0 ± 2.6	8.4 ± 0.4	78.0 ± 1.8	8.1 ± 0.2	42.3 ± 10.3	14.4 ± 1.1	71.9 ± 4.7	10.3 ± 0.8

segmentation performance.

4. The Necessary for Negative Prompt Points

PPO optimizes prompts by leveraging the interplay between positive prompts (targets) and negative prompts (backgrounds). In this section, we further illustrate the necessity of introducing negative prompts, with qualitative results shown in Table 3. Specifically, when optimizing prompt schemes without negative prompts, we first add negative prompts for PPO optimization and then remove them during segmentation for experimentation. For FM-generated initial prompts, the results indicate that the absence of negative prompts leads to a significant drop in segmentation performance. Although PPO can mitigate some errors by removing incorrect points, the performance remains inferior to using both optimized positive and negative prompts.

Quantitative results in Figure 4 demonstrate that, without negative prompts, background regions are frequently missegmented due to a lack of constraints. While adding unoptimized negative prompts reduces false positives, segmentation performance is still suboptimal. After PPO optimization, the accuracy of positive prompts improves, enhancing segmentation results. Combining optimized positive and negative prompts further enhances SAM's segmen-

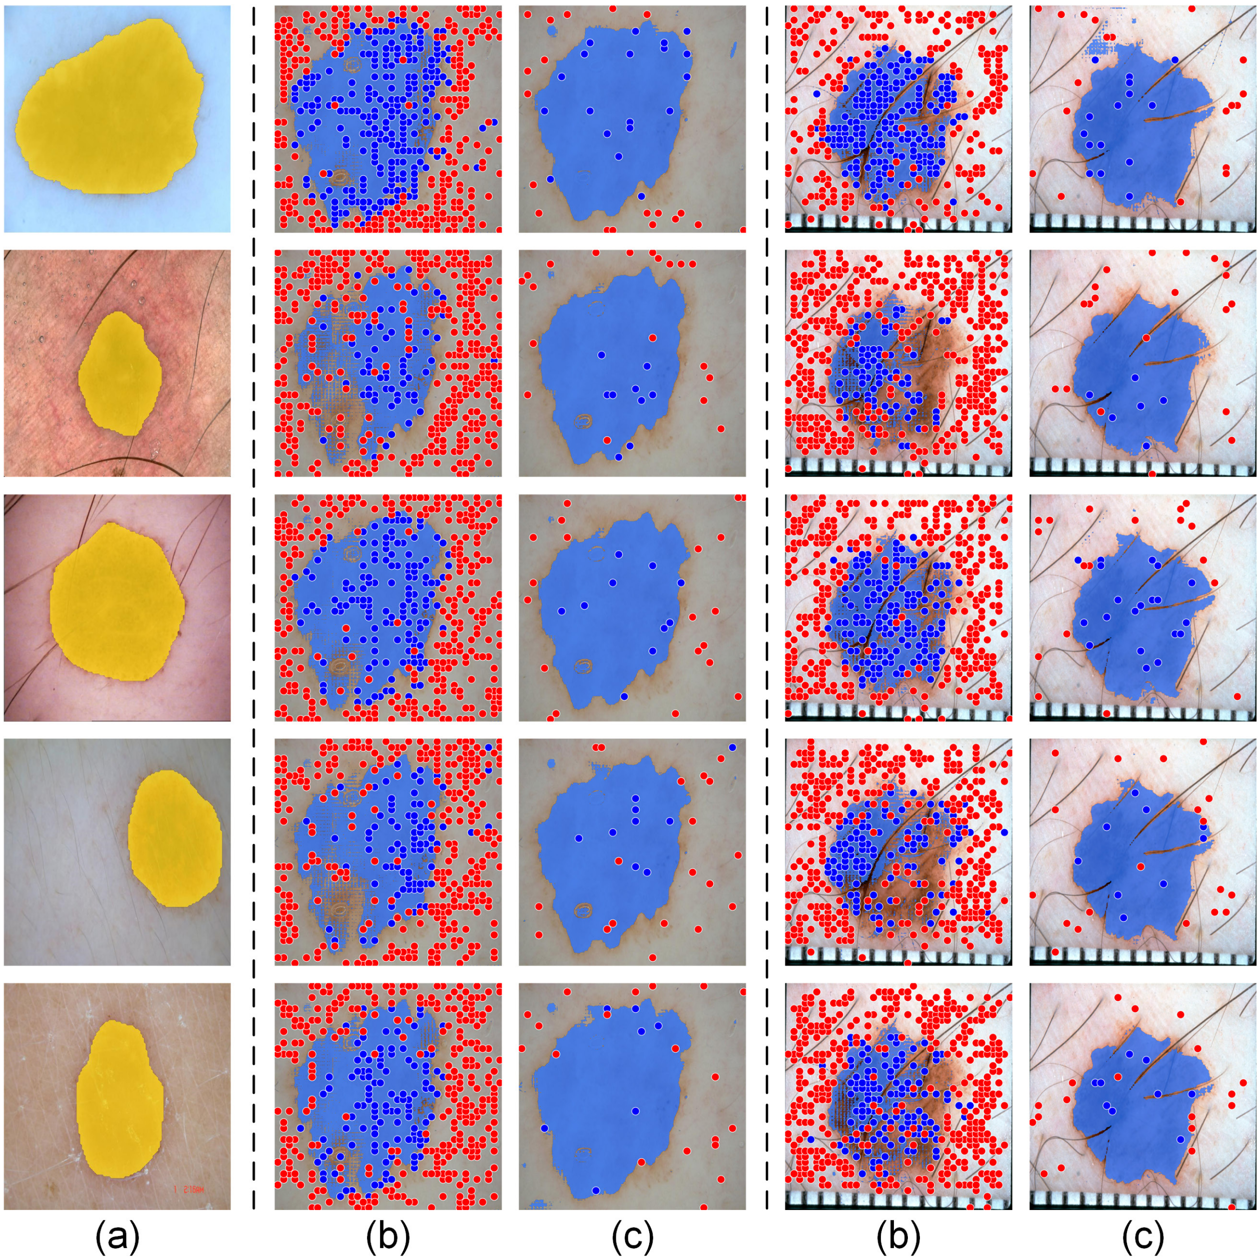

Figure 2. Qualitative results of optimizing initial prompts generated from different reference images in the ISIC dataset. (a) reference image, (b) initial prompts generated by FM, and (c) prompts optimized by PPO.

tation performance, achieving the best results.

Additionally, Matcher [1], a pioneering FM-based paradigm, generates segmentation results by optimizing only positive prompts. Matcher predicts all possible combinations of the positive prompt set to select the optimal mask as the final output. To validate the necessity of optimizing negative prompts, we incorporated negative prompts into Matcher for experimental analysis. However, applying the same optimization strategy to both positive and negative prompts would result in an unmanageable computational burden due to the excessive combinations of prompts. Therefore, we adopted a random selection strategy for negative prompts in Matcher. As shown in Table 3, adding negative prompts improves Matcher's segmentation performance, while incorporating PPO-optimized negative prompts further boosts results. Similar trends are observed in the quantitative results depicted in Figure 5. Optimizing only positive prompts often leads to substantial false

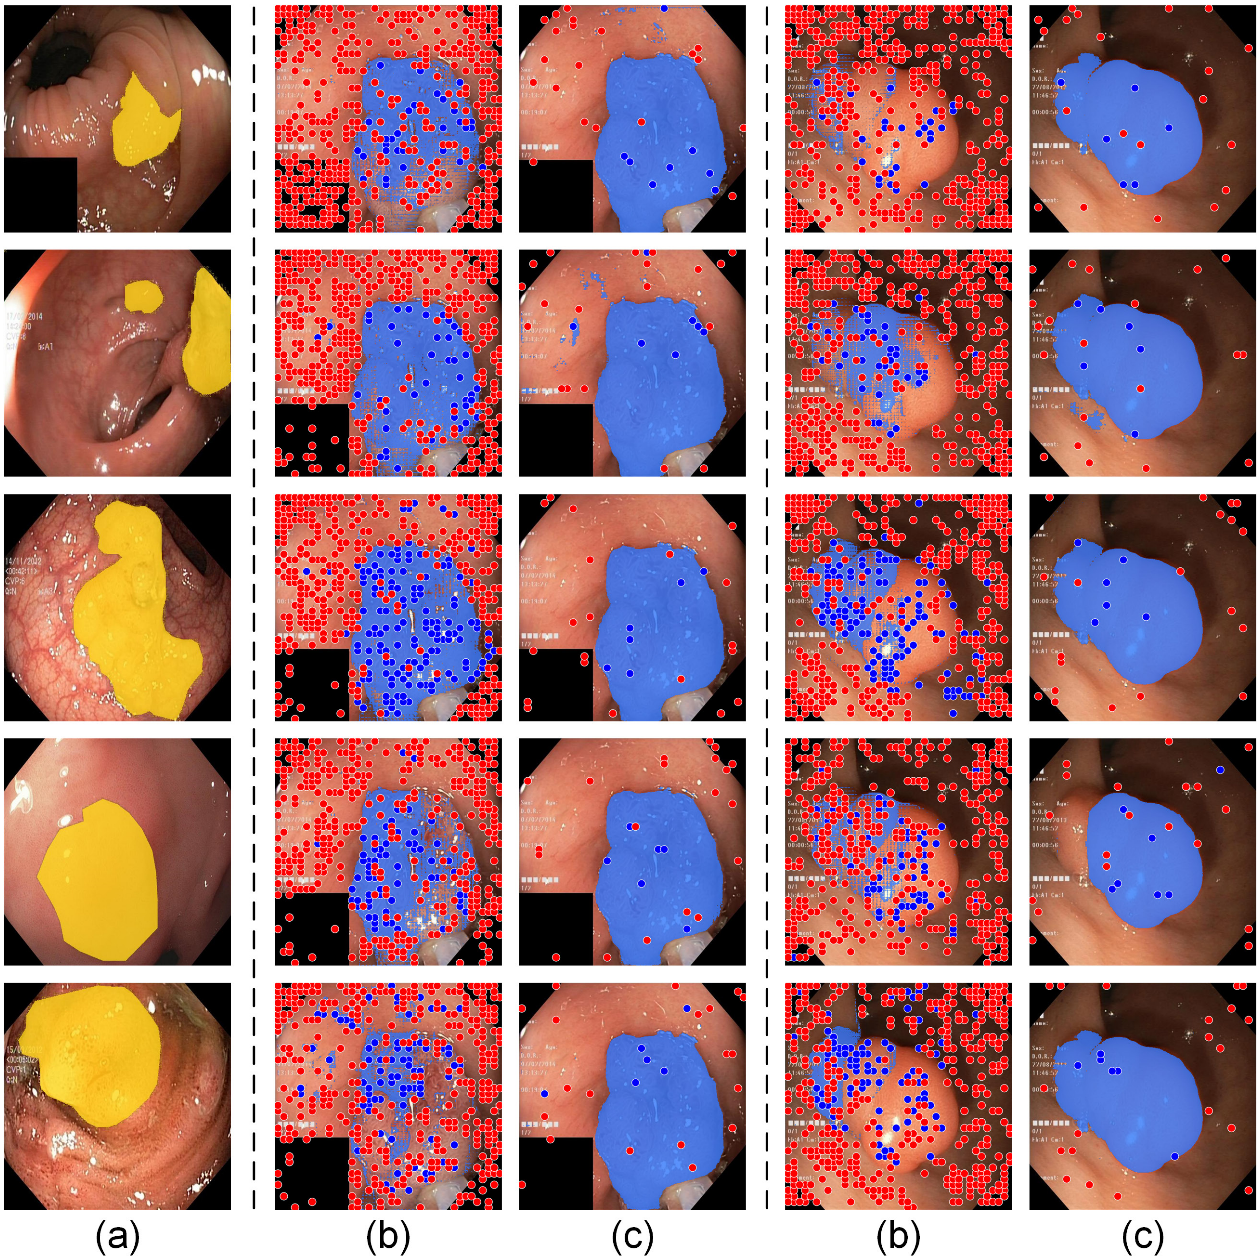

Figure 3. Qualitative results of optimizing initial prompts generated from different reference images in the Kvasir dataset. (a) reference image, (b) initial prompts generated by FM, and (c) prompts optimized by PPO.

positives. While PPO-optimized positive prompts partially mitigate this issue, adding negative prompts significantly reduces false positives. Further optimization of these negative prompts via PPO leads to even greater performance gains.

Our experiments demonstrate the effectiveness of negative prompts and validate the necessity of PPO for optimizing prompts. Furthermore, they highlight PPO's plug-and-play capability to improve segmentation performance across different paradigms through prompt optimization.

References

[1] Yang Liu, Muzhi Zhu, Hengtao Li, Hao Chen, Xinlong Wang, and Chunhua Shen. Matcher: Segment anything with one shot using all-purpose feature matching. *arXiv preprint arXiv:2305.13310*, 2023. 3, 5

Table 3. Ablation study on negative prompts and PPO optimization.

Methods	Negative prompts	PPO	ISIC mDSC (%↑)	ISIC mHD (↓)	Kvasir mDSC (%↑)	Kvasir mHD (↓)
FM	✗	✗	41.6	19.9	34.3	17.4
	✓	✗	54.4	16.6	49.1	14.6
	✗	✓	55.6	11.5	36.3	11.7
	✓	✓	75.4	8.1	70.3	9.0
Matcher [1]	✗	✗	73.7	9.7	46.2	16.0
	✓	✗	78.2	9.0	68.3	11.1
	✗	✓	79.1	8.2	71.6	9.6
	✓	✓	85.9	9.0	77.6	9.3

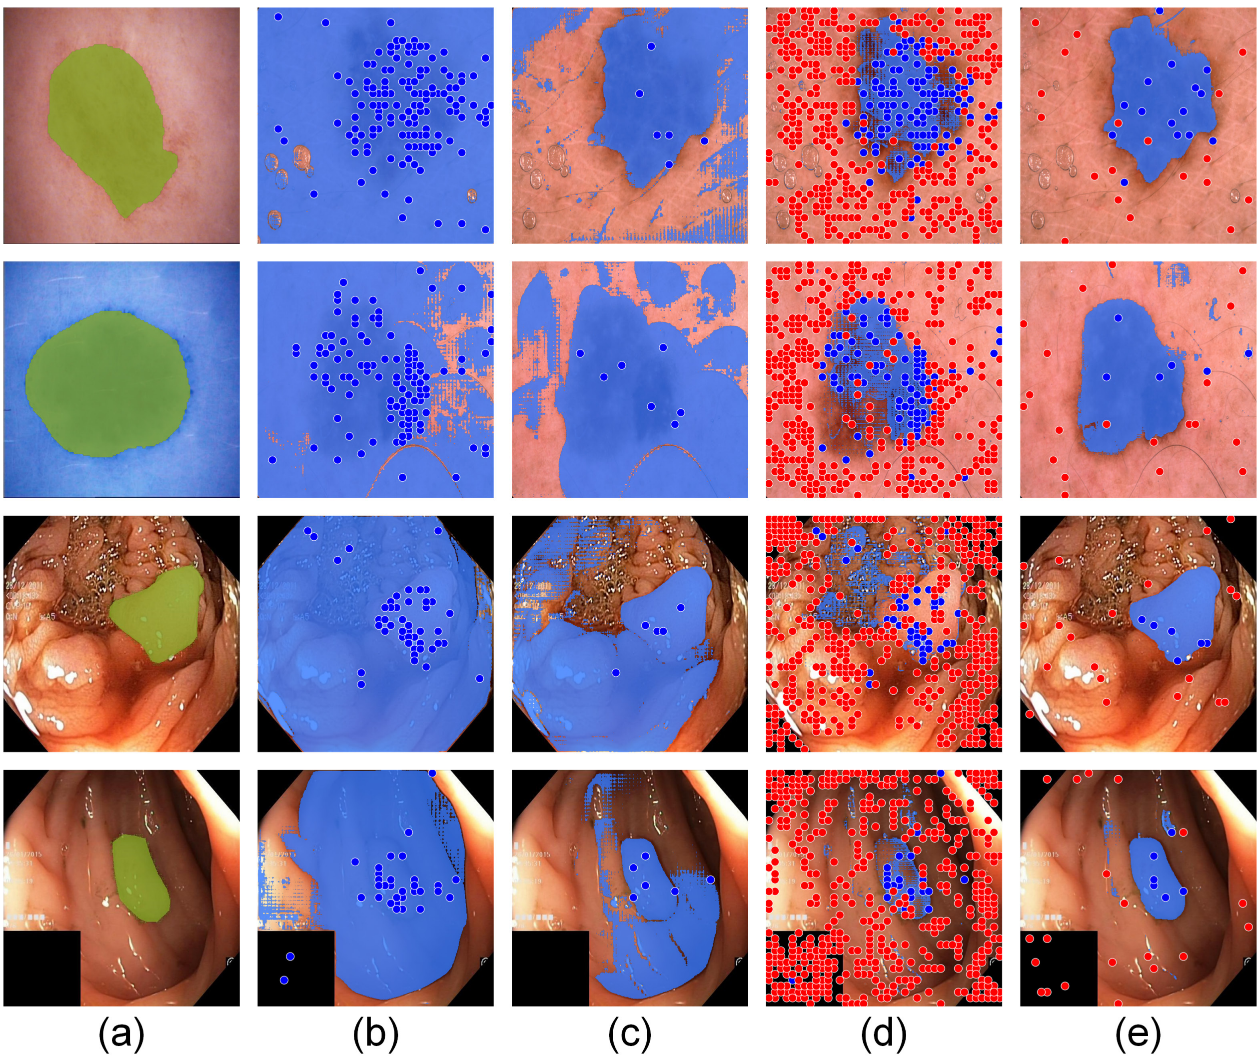

Figure 4. Qualitative results of the ablation study on negative prompts and PPO optimization in FM generation method. (a) ground truth, (b) without negative prompts and PPO, (c) without negative prompts but with PPO, (d) with negative prompts but without PPO, and (e) with both negative prompts and PPO.

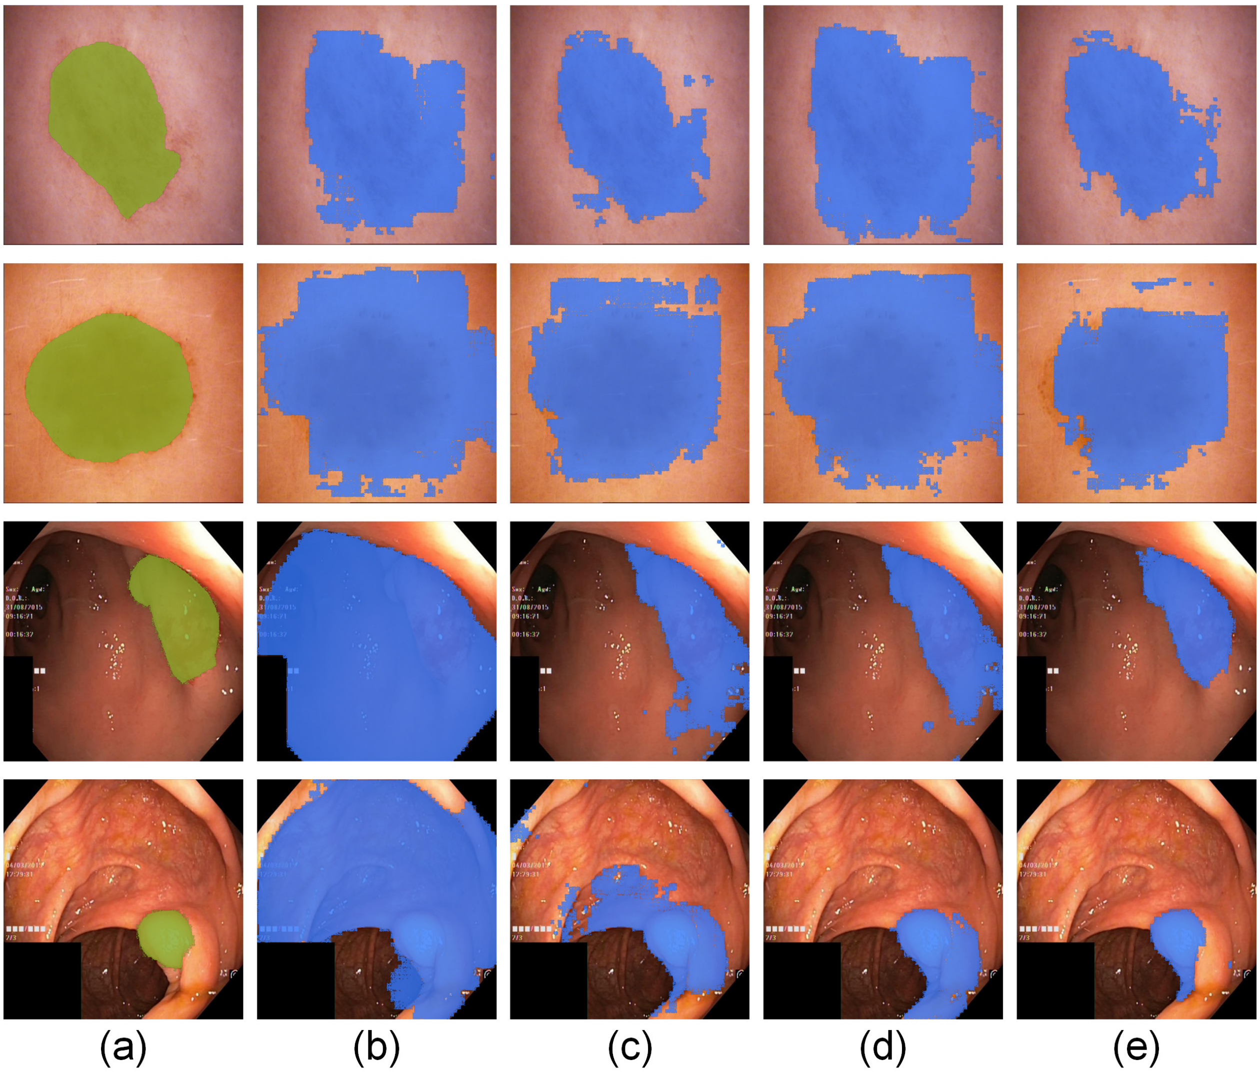

Figure 5. Qualitative results of the ablation study on negative prompts and PPO optimization with Matcher. Each scenario is represented as follows: (a) ground truth, (b) without negative prompts and PPO, (c) without negative prompts but with PPO, (d) with negative prompts but without PPO, and (e) with both negative prompts and PPO.