

Optimizing Efficiency and Effectiveness in Sequential Prompt Strategy for SAM using Reinforcement Learning

Yifei Huang¹[0009–0004–9309–3337], Chuyun Shen²[0009–0001–3622–1193], Wenhao Li³[0000–0003–2985–1098], Xiangfeng Wang², Bo Jin⁴, and Haibin Cai¹

¹ School of Software, East China Normal University, Shanghai, China
51255902080@stu.ecnu.edu.cn, hbcai@sei.ecnu.edu.cn

² School of Computer Science, East China Normal University, Shanghai, China

³ School of Data Science, The Chinese University of Hong Kong, Shenzhen, China

⁴ School of Software, TONGJI University, Shanghai, China

Abstract. In the rapidly advancing field of medical image analysis, Interactive Medical Image Segmentation (IMIS) plays a crucial role in augmenting diagnostic precision. Within the realm of IMIS, the Segment Anything Model (SAM), trained on natural images, demonstrates zero-shot capabilities when applied to medical images as the foundation model. Nevertheless, SAM has been observed to display considerable sensitivity to variations in interaction forms within interactive sequences, introducing substantial uncertainty into the interaction segmentation process. Consequently, the identification of optimal temporal prompt forms is essential for guiding clinicians in their utilization of SAM. Furthermore, determining the appropriate moment to terminate an interaction represents a delicate balance between efficiency and effectiveness. To provide sequential optimal prompt forms and best stopping time, we introduce an Adaptive Interaction and Early Stopping mechanism, named **AIES**. This mechanism models the IMIS process as a Markov Decision Process (MDP) and employs a Deep Q-network (DQN) with an adaptive penalty mechanism to optimize interaction forms and ascertain the optimal cessation point when implementing SAM. Upon evaluation using three public datasets, AIES identified an efficient and effective prompt strategy that significantly reduced interaction costs while achieving better segmentation accuracy than the rule-based method.

Keywords: Interactive Medical Image Segmentation · Reinforcement Learning · Early-Stopping · Reward Shaping.

1 Introduction

Deep learning-based segmentation has revolutionized various domains, including scene understanding, medical image analysis, and augmented reality [15]. Numerous algorithms have been proposed and evaluated [19, 6]. However, their success often hinges on the availability of large, well-annotated datasets. In the




Fig. 1: The architecture of the proposed AIES mechanism: AIES consists of two primary components: the general interactive segmentation model SAM and an RL-based agent. The RL-based agent processes input states, including original medical image slices, previous logits, and prior prompts, subsequently generating recommendations for subsequent interactions within clinical settings. If not terminal, clinicians give SAM new prompts. The medical image slices undergo encoding through SAM’s image encoder, while SAM’s prompt encoder encodes the prompts provided by clinicians. Subsequently, SAM’s lightweight mask decoder generates segmentation logits for the next interaction.

realm of medical imaging, the scarcity of data and ambiguous boundaries between regions present significant obstacles for these algorithms [25].

Interactive Medical Image Segmentation (IMIS) is a promising approach that enhances model accuracy through human feedback [13]. This feedback, in the form of user corrections or refinements, informs model iterations and updates predictions. User guidance may occur during training or application via clicks, scribbles, or other interactions. The recently introduced Segment Anything Model (SAM) represents a significant leap in IMIS due to its zero-shot generalization capability [9]. The integration of human feedback with SAM enhances prompt-driven segmentation, opening new avenues in the field [26]. SAM has been evaluated across various anatomical structures, imaging modalities, and ten public medical image segmentation datasets [16, 12]. However, despite its promise, SAM often performs suboptimally on medical datasets, with a reported mean Dice score of 58.52. This limitation has prompted the development of MedSAM, which aims to enhance diagnostic tools and personalize treatment plans by tailoring SAM to specific medical tasks. Such fine-tuning has led to significant improvements, with Dice score increases of 4.39% and 6.68% for ViT-B and ViT-H, respectively [5]. These findings highlight the effectiveness of SAM in medical image segmentation and the benefits of fine-tuning for IMIS.

However, SAM’s zero-shot capability shows significant variability in medical versus natural image segmentation, mainly due to its sensitivity to various prompt forms such as points or bounding boxes [5]. This issue is especially pronounced in the IMIS context [21]. Two main factors drive this issue: First, the interdependence of segmentation stages means prompt choices in one stage affect the next. Second, variability and randomness in human feedback can overlook how prompt choices impact performance and inter-prompt relationships. In clinical settings, determining the optimal moment to conclude an interaction remains challenging, often resulting in wasted efforts and reduced trust if outcomes are suboptimal [17]. Conversely, an ineffective interaction strategy may fail to meet desired outcomes. There’s a need for interactive strategies to use SAM more **efficiently** and **effectively**. Here, **efficiently** refers to minimizing the number of interactions required to attain satisfactory segmentation results, reducing the time and effort expended by users. Simultaneously, **effectively** signifies the capacity to generate accurate and reliable segmentation outcomes that fulfill users’ expectations and clinical requirements.

The Temporally-Extended Prompts Optimization (TEPO) framework was introduced to adaptively offer suitable prompt forms for SAM’s varying interaction sensitivity [21]. However, TEPO’s focus on fixed-length interactions risks efficiency by overlooking the potential benefits of adaptive interaction rounds and termination strategies. In contrast, we propose an **Adaptive Interaction and Early Stopping (AIES)** mechanism, which enhances both the efficiency and effectiveness of prompt strategies by considering not just the form of interaction, but also identifying the optimal termination timestep. Our approach includes several early termination strategies that optimize the segmentation process with SAM, ensuring the interaction ceases at the most beneficial point to prevent performance degradation and minimize computational waste. Utilizing reinforcement learning (RL), a proven method for sequential decision tasks across multiple domains [18, 10, 11, 7], we frame the issue as a Markov Decision Process (MDP). Through RL, the AIES mechanism assists users in choosing the optimal prompt form and deciding when to end interactions, effectively boosting SAM’s performance and segmentation accuracy. To align improvements in dice score with reduced interaction costs, the reward function integrates two components: a positive reinforcement signal that enhances dice scores, and a penalty that curtails interaction costs. This design enables the AIES mechanism to identify the most suitable prompt forms and determine the appropriate termination timing.

The contributions of this paper are threefold: **(1) Sequential prompt strategy optimization:** We propose a novel approach that leverages reinforcement learning to optimize interaction choices in the context of IMIS with SAM. By formulating the interaction selection process as an MDP, our method effectively determines the optimal prompt forms and termination timesteps, thereby significantly enhancing SAM’s zero-shot performance. **(2) Adaptive penalty for early termination:** Adaptive penalty helps users flexibly choose the desired number of interaction steps and obtain better segmentation results under the same interaction cost. **(3) Validation of generalization:** Integrating SAM

with reinforcement learning and early stopping, we prove our effectiveness on various medical image datasets, showcasing their wide applicability and superior performance in segmentation.

2 Method

To enhance the efficiency and effectiveness of SAM for clinical practitioners, this paper introduces the AIES mechanism. AIES uses adaptive feedback from human experts to choose the best prompt forms and termination timing. It treats guiding clinicians as a MDP. AIES boosts efficiency and effectiveness through a reward system. This system values changes in the Dice coefficient and interaction costs. Reinforcement learning is used for training. A comprehensive illustration of the AIES mechanism can be found in Fig. 1.

2.1 MDP Modeling of AIES

Considering the ability of reinforcement learning to model dynamic uncertainty on sequential decision-making tasks[20], our introduced approach AIES formalizes IMIS as an MDP. In this MDP framework, an agent (responsible for guiding the user in selecting prompt forms and deciding on termination) interacts with the environment (user’s interactive medical image segmentation with SAM) in discrete time steps. At each time step t , the agent observes a state s_t , performs an action a_t , and receives a reward r_t . The environment then determines the next state s_{t+1} based on the current state and the action taken. The MDP components include

State Space (S): At each discrete time step t , the state S_t is a composite entity consisting of the current medical image slice I , the segmentation logits P_{t-1} from the previous time step, and the set of interaction prompts U_{t-1} .

Action Space (A): The action space encompasses various prompt forms and an early termination action. Specifically, the action space A is defined as $A = \{1, 2, 3\}$, where Action 1 means plotting a bounding box to identify the target object. Action 2 means clicking the center of the error region in the previous segmentation prediction. Action 3 signifies stopping the training process.

Reward Function (R): The reward at each timestep, R_t , contains two parts: the change in dice coefficient and penalty. the change in dice coefficient ΔDice_t is computed as the change in the Dice coefficient between SAM’s prediction logit and the ground truth, across timesteps:

$$\Delta\text{Dice}_t = \text{Dice}(P_t, G) - \text{Dice}(P_{t-1}, G), \quad (1)$$

where G denotes the ground truth labels, $\text{Dice}(P_t, G)$ measures the similarity between the current prediction P_t and the ground truth, and $\text{Dice}(P_{t-1}, G)$ evaluates the similarity for the previous prediction.

Considering that after some interactions, the results of SAM do not significantly enhance, which decreases the efficiency of interaction with SAM, AIES

introduces a penalty for too many interactions. The reward at timestep t as:

$$R_t = \Delta\text{Dice}_t - h(P_t, P_t - 1, G), \quad (2)$$

where $h(P_t, P_t - 1, G)$ is a penalty function adjusting the penalty based on the current prediction, previous prediction, ground truth, and we will introduce this penalty in detail in the following subsection. This mechanism encourages the timely cessation of interaction, thereby ensuring high accuracy while minimizing unnecessary uncertainty. The adaptive mechanism optimizes the balance between efficiency and effectiveness.

2.2 Adaptive Penalty

In some situations, the prompt may not contribute to the incremental improvement of interactive segmentation performance, rendering such interventions ineffective and leading to the misallocation of valuable clinical resources. To mitigate this inefficiency, we propose the implementation of a penalty mechanism specifically designed to reduce the number of interaction steps needed. Specifically, AIES introduces an adaptive penalty with the objective of expediting the termination of the training process. This mechanism is represented by the following equation:

$$h(P_{t-1}, P_t, G) = \begin{cases} \lambda, & \text{if } \text{Dice}(P_t, G) < \epsilon, \\ \lambda + \max(0, \Delta\text{Dice}_t - \alpha), & \text{otherwise,} \end{cases} \quad (3)$$

where $\Delta\text{Dice}_t = \text{Dice}(P_t, G) - \text{Dice}(P_{t-1}, G)$, and α serves a lower threshold for additional penalties. According to the equation, a constant penalty, λ , is applied when the Dice coefficient is below the extra penalty bound ϵ , and the dice coefficient increase is lower than α . Nevertheless, an additional penalty is added. This adaptive penalty aims to optimize efficiency by considering the current segmentation results.

2.3 Optimization via Deep Q-Networks

Within the AIES framework, the reinforcement learning paradigm is employed to refine the strategy π , which maps states to actions, maximizing the cumulative reward over time. One popular approach to achieve this is by approximating the Q-function, $Q(s, a; \theta)$, which estimates the expected cumulative reward for taking action (a) in the state (s) and following the strategy π afterward. Deep Q-Networks (DQNs) are a type of RL algorithm that uses deep neural networks to approximate the Q-function. The loss function is formulated as follows:

$$L(\theta) = \mathbb{E}_{(s,a) \sim \pi} \left[\left(Q(s, a; \theta) - \left(r_t + \gamma \max_{a'} Q(s', a'; \theta') \right) \right)^2 \right], \quad (4)$$

where r_t denotes the immediate reward after the execution of action a in the state s , and γ represents the discount factor, emphasizing future rewards' importance. The term θ' refers to the target network's parameters, updated periodically for learning stability.

Table 1: Dice coefficient score across strategies for *Spleen* [1], *Polyp* [8, 4, 23, 22, 24], and *brats2020* [14]. AIES(x) denotes the AIES strategy with average actual interaction steps x, and AIES(6) is trained without action ‘stop’.

Spleen			Polyp-			Brats2020		
Length	Policy	Dice	Length	Policy	Dice	Length	Policy	Dice
1	Box_center	0.5548	1	Box_center	0.7888	1	Box_center	0.6901
	Center	0.5578		Center	0.6444		Center	0.4657
2	Box_center	0.6614	2	Box_center	0.8158	2	Box_center	0.6547
	Center	0.8197		Center	0.7354		Center	0.6213
	Random	0.4232		Random	0.5243		Random	0.4335
	AIES(2.3298)	0.8350		AIES(1.549)	0.8312		AIES(1.7207)	0.7144
3	Box_center	0.6462	3	Box_center	0.8564	3	Box_center	0.6890
	Center	0.8705		Center	0.8174		Center	0.7195
	Random	0.5951		Random	0.6661		Random	0.5800
	AIES(2.8404)	0.8576		AIES(2.3478)	0.8589		AIES(2.7041)	0.7489
4	Box_center	0.7740	4	Box_center	0.8774	4	Box_center	0.7478
	Center	0.9012		Center	0.8514		Center	0.7710
	Random	0.6806		Random	0.7345		Random	0.6327
	AIES(3.5532)	0.8825		AIES(3.7719)	0.8781		AIES(3.8366)	0.7649
5	Box_center	0.7740	5	Box_center	0.8774	5	Box_center	0.7791
	Center	0.9084		Center	0.8744		Center	0.7988
	Random	0.7188		Random	0.7794		Random	0.6913
	AIES(4.2553)	0.8873		AIES(4.3747)	0.8790		AIES(4.7651)	0.7889
6	Box_center	0.8484	6	Box_center	0.9012	6	Box_center	0.8062
	Center	0.9100		Center	0.8847		Center	0.8176
	Random	0.7890		Random	0.8427		Random	0.7914
	AIES(6)	0.8922		AIES(6)	0.9046		AIES(6)	0.8203

3 Experiments

This section provides the experimental results of our method and other rule-based strategies on different datasets. The datasets and details of the experiments and evaluation metrics are introduced as follows. The experiments aim to explore: 1) Whether SAM is sensitive to prompts forms in the context of IMIS? 2) The effectiveness of AIES across various medical imaging datasets. 3) Is the adaptive penalty more effective than the constant penalty? 4) Can AIES reduce misunderstanding cases?

- **Datasets:** For brain tumor segmentation, we used the *Brats2020* dataset from the BraTS Challenge [14, 2, 3]. It includes training, validation, and test sets with four MRI modalities, each sized $240 \times 240 \times 155$, applying data augmentation and preprocessing to enhance images. For colorectal cancer detection, we utilized *Polyp* datasets from Kvasir-SEG [8] and CVC-ClinicDB [4], with additional testing on five databases [8, 4, 23, 22, 24] for comprehensive evaluation. The *Spleen* dataset [1] provided CT scans for segmentation algorithm development. AIES were tested on diverse medical image datasets. The BraTS dataset, known for its complexity, was the primary focus due to its significance in medical imaging.

- **Implementation Details & Evaluation Metrics:** Images are standardized to 200×150 during preprocessing. Training images are augmented with flips, rotations, noise, and transformations for better generalization, while test images are cropped and scaled for accurate evaluation. Total training epochs is 100 epochs, and 2000 steps in each epoch, updating the Q-network every 100 steps. We utilized the Adam optimizer with a $1e - 5$ learning rate, γ of 0.9, and a 64 batch size. For penalty, we set the extra penalty bound ϵ as 0.8, and α is set at 0.2, and penalty values were optimized via grid search, we will introduce it in detail in the following subsection. Performance was evaluated using the Dice coefficient to compare segmentation with ground truth where the higher Dice coefficient indicates better segmentation, and step lengths to assess the efficiency of early stopping.

- **Comparison of results across various datasets:** To validate the effectiveness of AIES across various medical imaging datasets, we engaged three particularly challenging datasets: *Spleen* [1], *Polyp* [8, 4, 23, 22, 24], and *brats2020* [14] among various strategies (**Box-center**, **Center**, and **Random Policy**). The **Box-center** strategy involves selecting a bounding box initially, followed by choosing the center of the false area. The **Center** strategy solely selects the center of the false area, while the **Random** strategy randomly selects a bounding box, center point, or terminal, each with varying probabilities. We introduced perturbations, such as random offsets of 10 pixels, for both center and box prompts to evaluate the robustness of these strategies. In our work, users are only required to click the center point of the largest error area. A comprehensive comparison of the results can be found in Table 1.

Different prompt forms influence SAM results, highlighting its sensitivity to prompt formats. AIES notably outperformed in Dice scores, achieving 0.9046 on *Polyp* and 0.8203 on *Brats2020*, although it was slightly behind on the *Spleen* dataset compared to the Center strategy but required fewer interactions.

- **Adaptive penalty results:** To minimize futile interactions, the RL agent can employ a simple strategy like AIES-Constant, which uses a consistent penalty per step. However, this approach makes controlling the average step length challenging, as illustrated in Figure 2 where the relationship between penalty and average step length is nonlinear, particularly unstable around a penalty of 0.02. In contrast, AIES-Adaptive adopts a more dynamic penalty mechanism that controls average interaction lengths more predictively. This adaptability not only allows users to better balance performance and efficiency but also improves performance consistency across the same number of steps. In essence, the adaptive penalty offers easier customization and typically performs better, catering to diverse user needs.

- **Misunderstanding statistical results:** In the context of IMIS, the interactive algorithm may erroneously interpret the significance of the interaction, leading to inaccurate adjustments of the interaction results, a phenomenon referred to as misinterpretation[20]. In this study, instances where the Dice coefficients exhibit a reduction of greater than 0.1 following interaction are considered misinterpretations. This experimental setup is illustrated in Fig. 3. The histogram




Fig. 2: The top plot shows Average length vs. Dice coefficient (LOWESS smoothed). The bottom plot shows Penalty vs. average length with constant-penalty in green and adaptive-penalty in blue, with fitted lines. The horizontal axis represents average effective step length.




Fig. 3: Misunderstandings and the corresponding performance under different numbers of interactions: The histogram represents the percentage of misunderstanding under different strategies, and the points represent the corresponding dice.

presents the percentage of interactive misinterpretations of total interactions at the 2nd, 4th, and 6th steps. Meanwhile, the scatter plot shows average Dice scores for images without interaction stops at the 2nd, 4th, and 6th steps. The misinterpretation rate of AIES is lower than that of the Center and Box-center strategies and higher than the random strategy. However, it's important to note that AIES outperforms the random strategy in overall effectiveness. The random strategy, which fails to learn effective interaction forms, exhibits fewer misinterpretations due to its conservative nature. This indicates that AIES effectively reduces misunderstandings while achieving notable performance.

- **Qualitative experimental analysis:** The Fig. 4 visualize a comparison of strategies including Center, Box-center, Random, and AIES. AIES outperforms others by finding optimal stopping points for higher Dice coefficients. Unlike the Random strategy terminates too early, AIES stops judiciously, showcasing its efficiency and effectiveness.

4 Conclusion

This study introduces the AIES mechanism to aid users unfamiliar with SAM in IMIS. AIES optimizes interaction timing and temporal prompt forms with SAM, enhancing efficiency and effectiveness. Experiments demonstrate that SAM is sensitive to different sequential prompt forms. Our AIES method effectively finds an interaction strategy across various datasets, not only reducing interaction costs but also improving interactive segmentation results with fewer misunderstandings. In addition, the adaptive penalty in the reward function not only adjusts interaction costs more flexibly but also enhances performance, paving the way for further method optimizations.




Fig. 4: Visualization of strategies involves Star symbols for center points, green rectangles for bounding box prompts, and red rectangles for actual interaction steps. Subsequent to the red box, counterfactual interactions occur, where interactions continue even though they were actually stopped. In these counterfactuals, the continued interaction specifically uses the "center" form.

Acknowledgments. This study was supported in part by the Science and Technology Commission of Shanghai Municipality under grant number 22511106000 and 22511106004, and National Natural Science Foundation of China under grant number 62231019.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Antonelli, M., Reinke, A., Bakas, S., Farahani, K., Kopp-Schneider, A., Landman, B.A., Litjens, G., Menze, B., Ronneberger, O., Summers, R.M., et al.: The medical segmentation decathlon. *Nature communications* **13**(1), 4128 (2022)
2. Bakas, S., Akbari, H., Sotiras, A., Bilello, M., Rozycki, M., Kirby, J.S., Freymann, J.B., Farahani, K., Davatzikos, C.: Advancing the cancer genome atlas glioma mri collections with expert segmentation labels and radiomic features. *Scientific Data* **4**(1), 1–13 (2017)
3. Bakas, S., Reyes, M., Jakab, A., Bauer, S., Rempfler, M., Crimi, A., Shinohara, R., Berger, C., Ha, S., Rozycki, M., et al.: Identifying the best machine learning algorithms for brain tumor segmentation. Progression Assessment, and Overall Survival Prediction in the BRATS Challenge **10** (2018)
4. Bernal, J., Sánchez, F.J., Fernández-Esparrach, G., Gil, D., Rodríguez, C., Vilar-iño, F.: Wm-dova maps for accurate polyp highlighting in colonoscopy: Validation vs. saliency maps from physicians. *Computerized Medical Imaging and Graphics* **43**, 99–111 (2015)
5. Cheng, D., Qin, Z., Jiang, Z., Zhang, S., Lao, Q., Li, K.: Sam on medical images: A comprehensive study on three prompt modes. *arXiv preprint arXiv:2305.00035* (2023)
6. 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)
7. Gu, Z., Cheng, J., Fu, H., Zhou, K., Hao, H., Zhao, Y., Zhang, T., Gao, S., Liu, J.: Ce-net: Context encoder network for 2d medical image segmentation. *IEEE Transactions on Medical Imaging* **38**(10), 2281–2292 (2019)
8. Jha, D., Smedsrud, P.H., Riegler, M.A., Halvorsen, P., de Lange, T., Johansen, D., Johansen, H.D.: Kvasir-seg: A segmented polyp dataset. In: *The 26th International Conference on MultiMedia Modeling*. pp. 451–462. Springer (2020)
9. 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. *arXiv preprint arXiv:2304.02643* (2023)
10. Liao, X., Li, W., Xu, Q., Wang, X., Jin, B., Zhang, X., Wang, Y., Zhang, Y.: Iteratively-refined interactive 3d medical image segmentation with multi-agent reinforcement learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9394–9402 (2020)
11. Ma, C., Xu, Q., Wang, X., Jin, B., Zhang, X., Wang, Y., Zhang, Y.: Boundary-aware supervoxel-level iteratively refined interactive 3d image segmentation with multi-agent reinforcement learning. *IEEE Transactions on Medical Imaging* **40**(10), 2563–2574 (2020)
12. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. *Nature Communications* **15**(1), 654 (2024)
13. Marinov, Z., Jäger, P.F., Egger, J., Kleesiek, J., Stiefelhagen, R.: Deep interactive segmentation of medical images: A systematic review and taxonomy. *arXiv preprint arXiv:2311.13964* (2023)
14. Menze, B.H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., et al.: The multimodal brain tumor image segmentation benchmark (brats). *IEEE Transactions on Medical Imaging* **34**(10), 1993–2024 (2014)

15. Minaee, S., Boykov, Y., Porikli, F., Plaza, A., Kehtarnavaz, N., Terzopoulos, D.: Image segmentation using deep learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(7), 3523–3542 (2021)
16. Mohapatra, S., Gosai, A., Schlaug, G.: Brain extraction comparing segment anything model (sam) and fsl brain extraction tool. *arXiv preprint arXiv:2304.04738* (2023)
17. Nourani, M., King, J., Ragan, E.: The role of domain expertise in user trust and the impact of first impressions with intelligent systems. In: *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*. vol. 8, pp. 112–121 (2020)
18. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al.: Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* **35**, 27730–27744 (2022)
19. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *The 18th International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 234–241. Springer (2015)
20. Shen, C., Li, W., Xu, Q., Hu, B., Jin, B., Cai, H., Zhu, F., Li, Y., Wang, X.: Interactive medical image segmentation with self-adaptive confidence calibration. *Frontiers of Information Technology & Electronic Engineering* **24**(9), 1332–1348 (2023)
21. Shen, C., Li, W., Zhang, Y., Wang, Y., Wang, X.: Temporally-extended prompts optimization for sam in interactive medical image segmentation. In: *2023 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. pp. 3550–3557. IEEE (2023)
22. Silva, J., Histace, A., Romain, O., Dray, X., Granado, B.: Toward embedded detection of polyps in wce images for early diagnosis of colorectal cancer. *International Journal of Computer Assisted Radiology and Surgery* **9**, 283–293 (2014)
23. Tajbakhsh, N., Gurudu, S.R., Liang, J.: Automated polyp detection in colonoscopy videos using shape and context information. *IEEE Transactions on Medical Imaging* **35**(2), 630–644 (2015)
24. Vázquez, D., Bernal, J., Sánchez, F.J., Fernández-Esparrach, G., López, A.M., Romero, A., Drozdal, M., Courville, A., et al.: A benchmark for endoluminal scene segmentation of colonoscopy images. *Journal of Healthcare Engineering* **2017** (2017)
25. Wang, G., Li, W., Zuluaga, M.A., Pratt, R., Patel, P.A., Aertsen, M., Doel, T., David, A.L., Deprest, J., Ourselin, S., et al.: Interactive medical image segmentation using deep learning with image-specific fine tuning. *IEEE Transactions on Medical Imaging* **37**(7), 1562–1573 (2018)
26. Zhang, Y., Shen, Z., Jiao, R.: Segment anything model for medical image segmentation: Current applications and future directions. *Computers in Biology and Medicine* p. 108238 (2024)