

Interactive Medical Image Segmentation with Self-Adaptive Confidence Calibration

Wenhao Li[‡] Qisen Xu[‡] Chuyun Shen[‡] Bin Hu[†] Fengping Zhu[†] Yuxin Li[†]
 Bo Jin[‡] Xiangfeng Wang[‡]

Abstract

Medical image segmentation is one of the fundamental problems for artificial intelligence-based clinical decision systems. Current automatic medical image segmentation methods are often failed to meet clinical requirements. As such, a series of interactive segmentation algorithms are proposed to utilize expert correction information. However, existing methods suffer from some segmentation refining failure problems after long-term interactions and some cost problems from expert annotation, which hinder clinical applications. This paper proposes an interactive segmentation framework, called interactive **M**edical segmentation with self-adaptive **C**onfidence **C**alibration (**MECCA**), by introducing the corrective action evaluation, which combines the action-based confidence learning and multi-agent reinforcement learning (MARL). The evaluation is established through a novel action-based confidence network, and the corrective actions are obtained from MARL. Based on the confidential information, a self-adaptive reward function is designed to provide more detailed feedback, and a simulated label generation mechanism is proposed on unsupervised data to reduce over-reliance on labeled data. Experimental results on various medical image datasets have shown the significant performance of the proposed algorithm.

1 Introduction

Medical image segmentation is one of the essential tasks for computer-aided medical diagnosis. However, due to the pathological variability, dark lesion areas, as well as the uneven quality of the training data (lacking the consistency between imaging scanners, operators, and annotators), the accuracy of traditional convolutional neural networks (CNNs) type segmentation algorithms usually fail to meet clinical demands Wang et al. (2018a,b); Liao et al. (2020). To further refine the relatively inaccurate segmentation results, interactive image segmentation algorithms that take advantage of interactive correction information (*e.g.*, clicks, scribbles, or bounding boxes) are introduced Rajchl et al. (2016); Xu et al. (2016); Lin et al. (2016); Wang et al. (2018b,a); Bredell et al. (2018); Liao et al. (2020); Ma et al. (2020). The general interactive segmentation process is depicted in Figure 1, which contains two modules, *i.e.*, interactive module and utilization module. In the interactive module, the users (or experts) provide some interaction correction information (*e.g.*, hint information like clicks and scribbles). In the utilization module, the segmentation algorithm takes advantage of the interaction correction information to refine the previous model efficiently. The interactive segmentation algorithms can achieve better performance than traditional segmentation algorithms by utilizing the additional interactive correction information. The interaction process can be considered

[‡]School of Computer Science and Technology, East China Normal University, Shanghai, China. Email addresses: {52194501026@stu, 51184501067@stu, 52205901007@stu, bjin@cs, xfwang@cs}.ecnu.edu.cn. The first four authors contribute equally.

[†]Huashan Hospital, Fudan University, Shanghai, China. Email: {08301010188, zhufengping, liyuxin}@fudan.edu.cn

as the continuous cooperation between the model and the human expert. As a result, an excellent interactive segmentation model should understand experts’ interactive information and update itself to collaborate better.

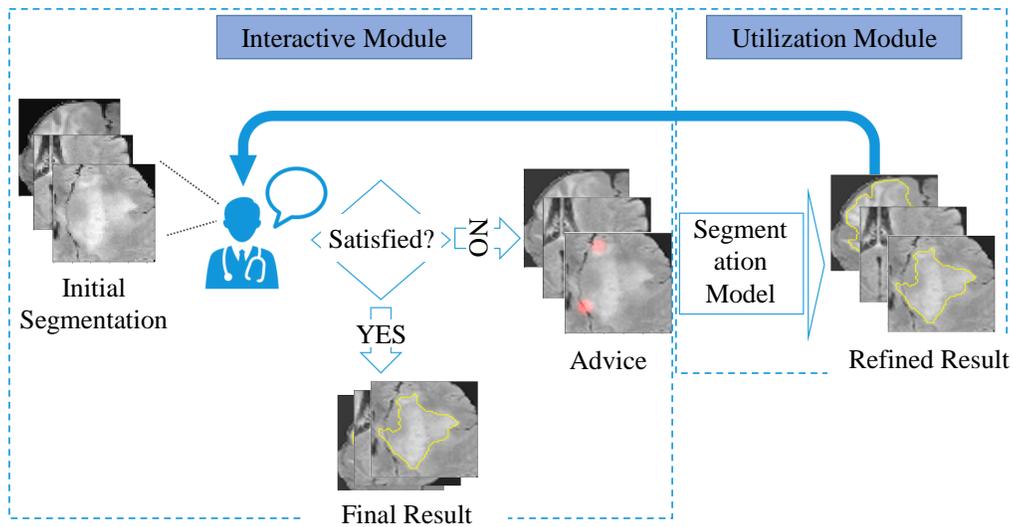


Figure 1: The framework of the interactive segmentation process. Interactive module: the expert observes current (or initial) segmentation and provides further correction information (red hints); Utilization module: new segmentation is refined based on the correction information.

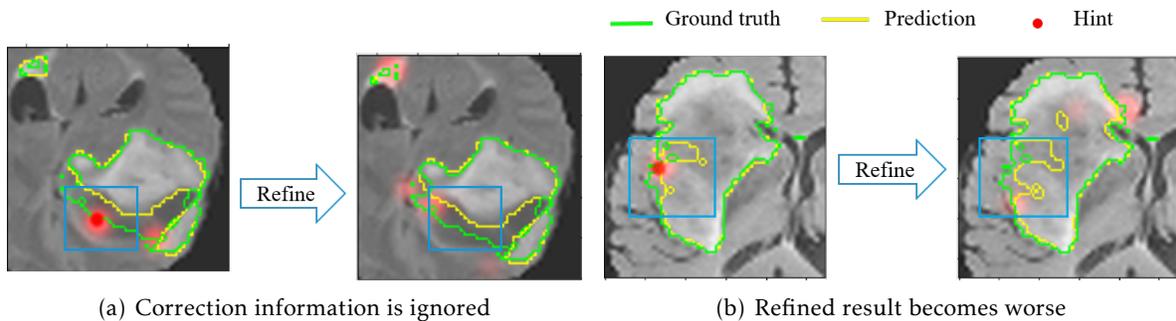


Figure 2: Segmentation refining failure after long-term interactions on BraTS2015 images: a) The segmentation model can not fully understand the hint information or just ignores it. b) The interactive segmentation model could misunderstand the expert’s interaction correction, which results in the worse result;

Methods such as BIFSeg Wang et al. (2018a) and DeepIGeos Wang et al. (2018b) model user interactions as hard constraints through conditional random fields (CRFs) end-to-end training. However, these model more focus on the *one-step* interaction and the refined segmentation after the first step cannot be efficiently used in these model Liao et al. (2020); Ma et al. (2020). Thus, these model can not utilize the long-term interactive information but the short-term interactive information efficiently. The follow-up InterCNN Bredell et al. (2018) and other methods Liao et al. (2020); Ma et al. (2020) are improved on DeepIGeos and BIFSeg, modeling the problem as an *iterative* interaction problem, which are more focus on the *multi-step* interaction. Therefore, these method can cover the data distribution corresponding to different interaction levels in the training phase to effectively use long-term

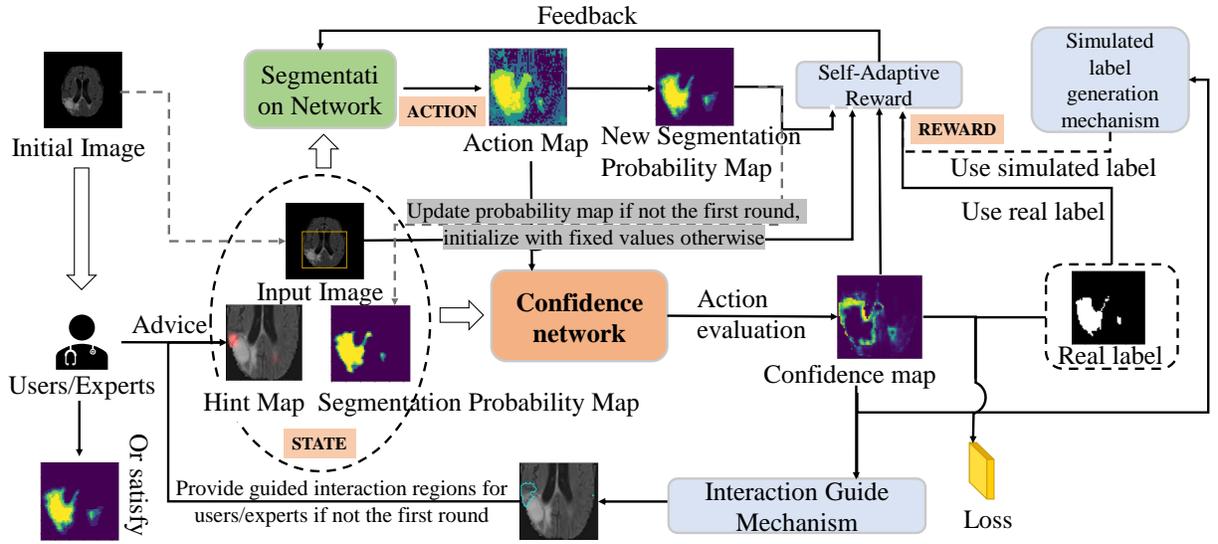
interactions, while ignoring the stochasticity or uncertainty of the model makes it difficult for them to effectively use short-term interactive information. Existing interactive methods cannot effectively utilize both short-term and long-term interactive information at the same time, which leads to the *interactive misunderstanding* phenomenon[‡]. Figure 2 presents this phenomenon when implementing the popular interactive segmentation algorithm InterCNN Bredell et al. (2018) on the BraTS2015 dataset. The algorithm ignores the expert’s correction information as in Figure 2(a), and even be adversely affected by correction information as in Figure 2(b). These inconsistencies between hint information and the refined results indicate that existing algorithms still face the critical challenge of inefficiency utilization of interactive correction information. The reason for segmentation refining failure is that, at the end of the interactive procedure, the total loss of the model will guarantee the main area’s priority and ignore some small but challenging areas, such as edges. The regions that are insensitive to interactive correction information could be considered hard-to-segment regions Nie et al. (2019); Shrivastava et al. (2016), which leads to mediocre segmentation refinement. If these regions are not paid more attention during the training stage, the segmentation model will inevitably dominate those easy-to-segment regions. This problem is more serious for medical images, where the hard-to-segment regions usually are tumor boundaries and are very important for clinical diagnosis and surgery. Therefore, it becomes urgent to improve the utilization of correction information, especially for the hard-to-segment regions.

Besides, the requirement of large amounts of expert-annotated images is another crucial issue. On the one hand, accurate annotations of medical images require plenty of time of experts, so it is expensive to acquire sufficient data with high-quality annotations. On the other hand, more unlabeled images can be obtained with much less cost but existing popular interactive segmentation methods usually ignore utilizing these low-cost images. Thus it becomes more and more critical to reduce the need for expert-annotated images while simultaneously taking advantage of those unlabeled images.

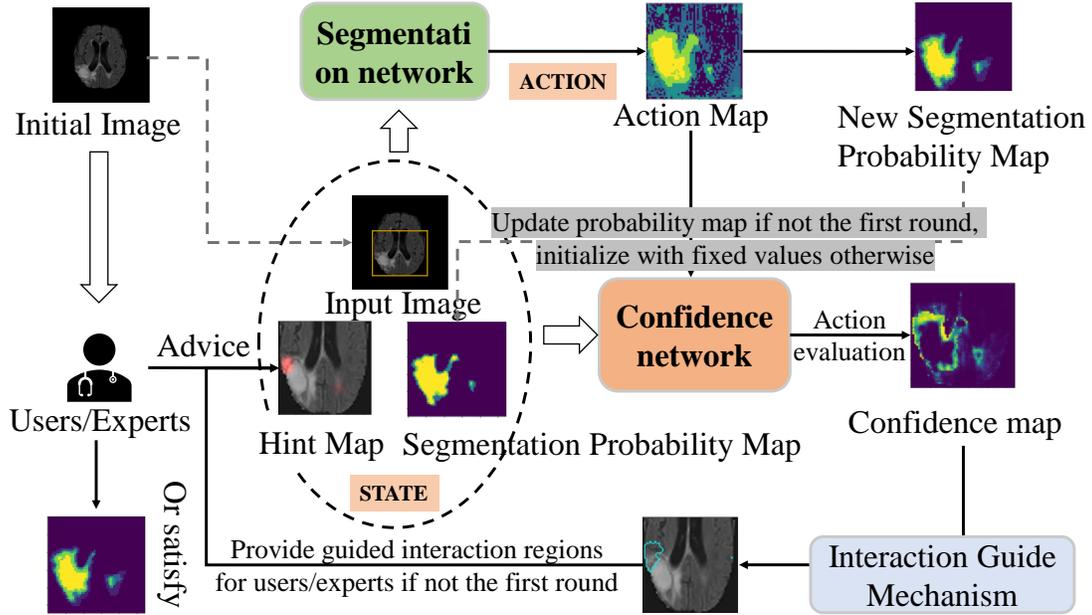
In this paper, we propose a novel interactive segmentation algorithm for 3D medical images called interactive **ME**dicinal segmentation with self-adaptive **CA**libration (**MECCA**). MECCA combines an action-based confidence learning network with the multi-agent reinforcement learning framework. The action-based confidence network could evaluate the corrective actions’ quality by directly calculating the refinement actions’ confidence. Unlike learning the confidence of the overall segmentation result, the confidence of actions could better evaluate whether the segmentation model has correctly utilized the experts’ interactive correction information. We formulate the iterative interactions process as a Markov decision process(MDP) to model the dynamic process and further introduce the reinforcement learning technique Song et al. (2018). Further, instead of setting each image or each patient (a series of images) as the agent, we consider each voxel[†] as an agent while each agent aims to learn the segmentation policy and makes its decision Furuta et al. (2020). Specifically, after receiving the interactive correction information, each agent will modify its label by changing (increasing or decreasing) the category probability. The novel action-based confidence network will directly evaluate each agent’s corrective action obtained from the corrective policy of MARL. By utilizing this action-based confidence network, two additional techniques are further proposed to improve the utilization efficiency of the interactive correction information: 1) compared with former manual-designed rewards, a self-adaptive reward function of each action is constructed, which could provide more meticulous feedback under a flexible framework; 2) a simulated label generation mechanism is established by utilizing the interactive correction information as the weakly supervisory signal. By combining the confidence network, the simulated label generation mechanism can approximately generate the labels for unsupervised images and reduce over-reliance on labeled data. The overall algorithm framework of our proposed algorithm is shown in Figure 3(a). The

[‡]This phenomenon will be discussed in detail in the Section 3.3.

[†]The smallest unit in three-dimensional space.



(a) The architecture of MECCA.



(b) The testing stage of MECCA.

Figure 3: **(a)** The architecture of MECCA. The segmentation module outputs actions to change the segmentation probability of each voxel(agent) at each interaction step. Meanwhile, the confidence network will estimate the confidence of actions, which will generate the self-adaptive reward and simulated label. The confidence map can provide the advice regions of the next interaction step to experts. **(b)** The testing stage of MECCA. In the first instance, all voxels in the segmentation probability map are initialized to a fixed value (set to 0.5 in the experiment). Users or experts randomly mark hint points according to the initialized segmentation probability map.

main contributions of this work are summarized as follows: 1) A novel framework is proposed for interactive medical image segmentation by combining action-based confidence learning network with multi-agent reinforcement learning; 2) A self-adaptive feedback mechanism (self-adaptive

reward) is constructed with the action-based confidence network to alleviate the effect of interactive misunderstanding phenomenon during the interaction process; 3) Simulated supervisory signals can be generated based on the confidence learning network and actions; hence much less labeled data (ground truth data) are needed to achieve the same performance.

The following is the roadmap of this paper. Section 2 provides a brief but complete introduction to related works in image segmentation and interactive image segmentation, while the background material on confidence learning is also provided. Section 3 describes the proposed interactive image segmentation algorithm, including the segmentation policy, the evaluation of corrective actions, the self-adaptive rewards, and the label generation scheme. Extensive experiments are presented in Section 4. Conclusions and future work discussions are proposed in Section 5.

2 Related Works

Image segmentation is the fundamental problem of computer vision or image processing that has been widely and long-termly studied. Deep learning (DL) has further promoted the development of automatic segmentation algorithms, the same as other applications. CNN-type methods are typical DL algorithms for image segmentation, e.g., fully convolutional networks (FCN) Long et al. (2015) and DeepLab Chen et al. (2018). The U-Net Ronneberger et al. (2015), which is considered an evolutionary variant of FCN, becomes one of the SOTA methods and performs better for medical image segmentation. Medical image segmentation is the key to modern auxiliary diagnosis and treatment response evaluation. A series of related works were proposed with the progressive performance for medical image segmentation Milletari et al. (2016); Kamnitsas et al. (2017b); Li et al. (2017). In the following, we will review the development of interactive image segmentation methods and discuss confidence learning for image segmentation relevant to our proposed algorithm.

Traditional interactive image segmentation methods. The classical random Walk Grady (2006) can create a weight map with pixels as vertices and segment the image based on user interactions. GrabCut Rother et al. (2004) and GraphCut Boykov & Jolly (2001) could associate image segmentation with the maximum flow and minimum cut algorithms on graphs, respectively, while Geos Criminisi et al. (2008) was proposed to measure the similarity between pixels geodesic distance. These traditional interactive image segmentation methods aim to utilize additional expert interaction information to modify the segmentation performance further.

DL-based interactive image segmentation methods. Xu et al. (2016) segments images based on CNN interactively. DeepCut Rajchl et al. (2016) and ScribbleSup Lin et al. (2016) both employed weakly supervised expert hints to establish interactive image segmentation methods. DeepIGeoS Wang et al. (2018b) employed geodesic distance metric to construct a hint map. The interactive segmentation process can be considered as a sequential iterative process. It becomes natural to introduce the RL framework to model the interactive segmentation process. Polygon-RNN Castrejon et al. (2017) fundamentally segmented each target as a polygon and iteratively chose the polygon vertexes through a recurrent neural network(RNN). Polygon-RNN+ Acuna et al. (2018) employed almost the same idea of Polygon-RNN but learned to choose vertexes by RL. SeedNet Song et al. (2018) trained an expert interaction generation RL model, which obtains new simulated interaction information at each interaction step. IteR-MRL Liao et al. (2020) and BS-IRIS Ma et al. (2020) both modeled the dynamic interaction process as an MDP and employed multi-agent RL(MARL) models to segment images. Some researches also aim to reduce the annotation cost of interactive image segmentation. IFSL Feng et al. (2021) introduces interactive learning into the few-shot learning strategy and addresses the annotation burden of medical image segmentation models. IOG Zhang et al. (2020) proposes

a practical Inside-Outside Guidance approach for minimizing the labeling cost. These interactive methods are difficult to effectively utilize experts' short-term and long-term interaction information simultaneously, thus making error correction operations.

Uncertainty estimation for image segmentation Uncertainty estimates are helpful in the context of deployed machine learning systems as they are capable of detecting when a neural network is likely to make an incorrect prediction or when the input may be out-of-distribution. Traditionally, much of the work are inspired by Bayesian statistics, or Bayesian Neural Network (BNN) [MacKay \(1992\)](#); [Neal \(2012\)](#). Unfortunately, Bayesian inference is computationally intractable in practice, so much effort has been put into developing approximations of Bayesian neural networks that are easier to train. Recent efforts to approximate the BNNs in this area include Monte-Carlo Dropout [Gal & Ghahramani \(2016\)](#), Multiplicative Normalizing Flows [Louizos & Welling \(2017\)](#), and Stochastic Batch Normalization [Atanov et al. \(2018\)](#). These methods are capable of producing uncertainty estimates, although with varying degrees of success. The main disadvantage with these BNN approximations is that they require sampling to generate the output distributions. As such, uncertainty estimates are often time-consuming or resource-intensive to produce, requiring 10 to 100 forward passes through a neural network to produce useful uncertainty estimates at inference time. An alternative to BNNs is ensembling methods [Dietterich \(2000\)](#); [Kamnitsas et al. \(2017a\)](#); [Mehrtash et al. \(2018\)](#); [Lakshminarayanan et al. \(2017\)](#); [Mehrtash et al. \(2020\)](#), which propose a frequentist approach to the problem of uncertainty estimation by training many models and observing the variance in their predictions. However, this technique is still quite resourcing intensive, as it requires inference from multiple models to produce the uncertainty estimate. A promising alternative to sampling-based methods is to instead have the neural network learn what its uncertainty should be for any give input, i.e., *learning-based uncertainty estimation* or *confidence learning*, as demonstrated in [Kendall & Gal \(2017\)](#); [DeVries & Taylor \(2018a\)](#); [Robinson et al. \(2018\)](#); [DeVries & Taylor \(2018b\)](#); [Jungo & Reyes \(2019\)](#); [Moeskops et al. \(2017\)](#); [Hung et al. \(2018\)](#); [Nie et al. \(2019\)](#). These methods commonly consist of a segmentation and confidence network and are more computationally efficient than other techniques. Thus they are better suited when computational resources are limited or when real-time inference is required, such as the interactive segmentation scenario considered in this paper. Specifically, [Kendall & Gal \(2017\)](#) introduces a confidence network to predict the *aleatoric* and *epistemic* uncertainties by imitating classic Bayesian tools. The segmentation network of [DeVries & Taylor \(2018a,b\)](#) produces two separate outputs: prediction probabilities and a confidence estimate. Confidence estimates are motivated by interpolating between the predicted probability distribution and the target distribution during training, where the degree of interpolation is proportional to the confidence estimate. A series of work [Nie et al. \(2019\)](#); [Hung et al. \(2018\)](#); [Moeskops et al. \(2017\)](#) focus on incorporating the uncertainty estimation into the adversarial learning process, where the segmentation network corresponds to the generator, and the confidence network is the discriminator accordingly. [Moeskops et al. \(2017\)](#) firstly employed GANs to improve the CNN-based brain MRI segmentation method. The semi-supervised learning technique is used in [Hung et al. \(2018\)](#) to predict trustworthy regions in unlabeled images. [Nie et al. \(2019\)](#) proposed a difficulty-aware attention mechanism to handle those hard samples or challenging regions. Different from the previous work to learning uncertainty through imitation, joint training, or adversarial learning, a simple but powerful alternative is to introduce an auxiliary task, such as to predict the overlap between a proposed segmentation and its ground truth [Robinson et al. \(2018\)](#), or to predict the voxel-wise false positive and false negative [Jungo & Reyes \(2019\)](#).

Remark. In our proposed algorithm, the confidence network should evaluate the confidence of calibrating actions instead of the segmentation result, which is the most significant difference between

our interactive method and previous learning-based uncertainty estimation methods. Therefore, we design a novel action-oriented auxiliary task to predict whether the *direction* of voxel-wise action is consistent with the ground truth.

3 The Proposed Algorithm

This section will introduce our proposed interactive segmentation algorithm MECCA, which can iteratively evaluate the refinement actions and feedback to the segmentation model. The algorithm framework follows the multi-agent reinforcement learning structure (Section 3.1 and Section 3.2) and an action-based confidence learning module (Section 3.3) is introduced to evaluate the confidence of corrective actions. This action-based confidence learning module can be used to establish the self-adaptive reward scheme and simulated label generation mechanism to utilize the interactive correction information efficiently. The architecture overview of MECCA has been depicted in Figure 3(a). The model’s state information includes the original 3D image, the previous segmentation probability, and the hint map generated from interaction and confidence maps. Based on the current state information, the segmentation module gives suggested actions to refine previous segmentation results by adjusting the segmentation probability of each voxel (agent). Further, the state information will be utilized to evaluate the confidence of obtained actions through the confidence network, with a confidence map as the output. The self-adaptive reward is designed through a self-adaptive weighting scheme based on the action confidence evaluation (Section 3.4). The self-adaptive reward map can be considered a value map with the same size as the original input image (each agent has its self-adaptive reward), which can reflect the performance of the corresponding agent’s action. Besides, MECCA will suggest some low confidence regions for the user to interact next (Section 3.5). Besides, the confidence map can also be used to generate the simulated label by comparing it with actions, which will be described in Section 3.6. The newly obtained hint map, the adjusted segmentation probability result, and the original 3D image form a new state. The process described above is repeated until the segmentation result meets requirements. To emphasize, during the testing stage in Figure 3(b), there is no need to calculate the self-adaptive rewards, and at the same time, we only need the obtained actions and suggested interaction areas.

3.1 MARL-driven Interactive Segmentation Framework

This study employs the multi-agent reinforcement learning structure to formulate the interactive segmentation process and continuously give error-corrective actions at each interaction step. Let $\mathbf{x} = (x_1, \dots, x_N)$ denotes the input image and x_i denotes the i -th voxel of the image. In the setting of MARL, every voxel x_i is treated as an agent with its own refinement policy $\pi_i(a_i^{(t)}, s_i^{(t)})$. At time step t , agent x_i gets action $a_i^{(t)}$ from the segmentation network according to its current state $s_i^{(t)}$. After taking the action, agent will receive a reward $r_i^{(t)}$ according to the segmentation result.

The state $s_i^{(t)}$ for agent x_i is concentrated by its voxel value b_i , its current segmentation probability $p_i^{(t)}$ and the value $h_i^{(t)}$ on the hint map. In particular, the segmentation probabilities of all agents are initialized to 0.5 and range from 0 to 1. The hint map $h^{(t)}$ is transformed from the user’s hints which are in the form of edge points. At each step, users click on some edges, which are not correctly predicted, as hints. In order to let the model receive hint information, we generates a 3D Gaussian (with 8-voxels kernel size) centered on each of the edge points as the hint map input to the segmentation network.

The action $a_i^{(t)}$ for agent x_i is sampled from its policy and used to adjust its previous segmentation probability:

$$a_i^{(t)} \sim \pi_\theta(a_i^{(t)} | s_i^{(t)}), \quad (1)$$

$$p_i^{(t+1)} = \text{clip}(p_i^{(t)} + a_i^{(t)}, 0, 1), \quad (2)$$

where $a_i^{(t)} \in \mathcal{A}$ and the *clip* operation modifies the probability to the interval $[0,1]$. The action set \mathcal{A} contains actions of different scales, allowing the agent to select the proper action. In our setting, the $\mathcal{A} = \{\pm 0.1, \pm 0.2, \pm 0.4\}$. 0 is not used as one of the actions is mainly due to the following three reasons: 1) in the early stages of the experiment, we found that using 0 as action will make the algorithm converge slow, and it requires more than 2 times the number of interactions to achieve the performance of baselines. Therefore, we remove 0 from the action space and encourage the agent to explore in the early stage of training. Not only does the algorithm converge faster, but it also requires fewer interactions; 2) in order to prevent the algorithm’s output from exceeding the valid range, we have clipped the output action (this is also a common technique in reinforcement learning in order to ensure that the output of the policy is in the valid range), as shown in the (2). 3) due to the existence of the self-adaptive confidence calibration mechanism, the action output by the algorithm will generally change in the direction of the correct label. Combined with the clipping technique used in the second point, the algorithm can guarantee effectual output even if there is no 0-action.

The reward $r_i^{(t)}$ is the feedback (positive or negative) of the action and used to update the refinement policy. The design of reward is a significant part of our algorithm, and we will introduce it in details in Section 3.4.

3.2 Segmentation Network

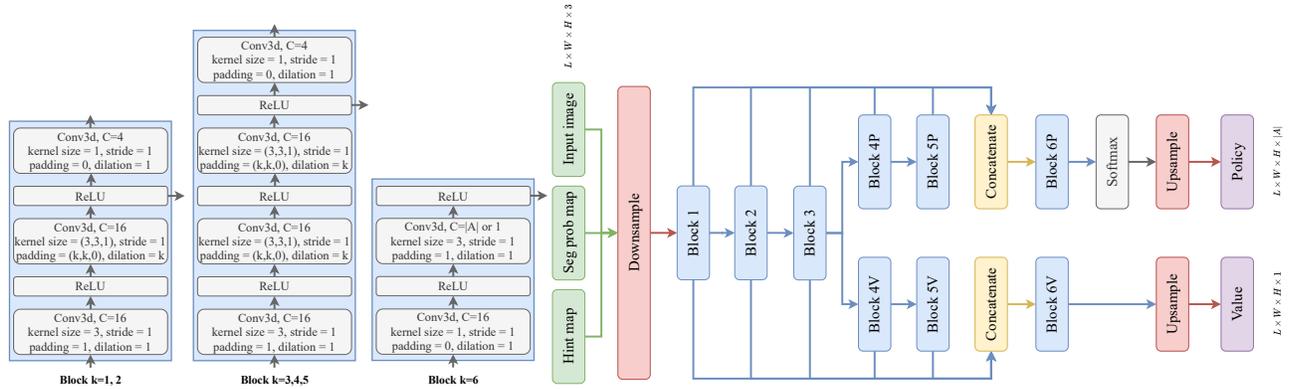


Figure 4: The detailed architecture of the segmentation network. The confidence network architecture is the same as the value branch of the segmentation network, but the parameters are not shared. $|A|$ represents the size of the action space.

The segmentation network adopts the P-Net, proposed in Wang et al. (2018b), as the backbone. The segmentation network S has two heads: policy head and value head. Two heads share the first three 3D convolutional blocks to extract low-level features. Each of the blocks has two convolution layers and the size of the convolution kernel is fixed as 3×3 in all these convolution layers. All the convolution kernels are dilated convolution, which can reduce the loss of the resolution. Both of the heads have another two 3D convolutional blocks to extract specific high-level features. The detailed architecture is shown in Figure 4. Specifically, the policy head is to output the policy $\pi_i(a_i^{(t)}, s_i^{(t)})$, which is the distribution of action probabilities under current state. By taking actions on different scales, the probabilities will be dynamically adjusted. The value head is to estimate the value of the current state, which evaluates how good the current state is and estimates the expected return:

$$V_i^{(t)} = \mathbb{E} \left[R_i^{(t)} | s_i = s_i^{(t)} \right] = \mathbb{E} \left[\sum_{k=t}^T \gamma^{t-1} r_i^{(k)} \right], \quad (3)$$

where T is the terminal time step in the interaction process, and γ is the discount factor. θ_v denotes the parameters of the value head, and the gradient with respect to θ_v is:

$$d_{\theta_v} = \nabla_{\theta_v} A^2(s^{(t)}, a^{(t)}), \quad (4)$$

$$A(s^{(t)}, a^{(t)}) = \sum_{k=t}^T \gamma^{k-t} \bar{r}^{(k)} - V(s^{(t)}), \quad (5)$$

where $\bar{r}^{(k)}$ is the mean reward of all voxels at timestep k . $A(s^{(t)}, a^{(t)})$ represents the advantage function Mnih et al. (2016).

The goal of the policy head is to maximize the expected return by selecting proper actions in state $s^{(t)}$. This study uses θ_p to denote the parameters of the policy head, and the gradient for θ_p is denoted as

$$d_{\theta_p} = -\nabla_{\theta_p} \pi(a^{(t)}|s^{(t)}) A(a^{(t)}, s^{(t)}). \quad (6)$$

Usually, the policy head is updated after the value head.

3.3 Action Confidence Learning

As we mentioned in section 1, there will be some situations where the segmentation model misunderstands or ignores the hint information. To some extent, these samples (or regions) with the phenomenon of interactive misunderstanding are hard samples (or regions). Although these samples may account for a small percentage of the dataset, they are critical for improving generalization and robustness. The most important thing is finding a *professional easy-or-hard representer* Nie et al. (2019) to identify them when interacting. Focal loss Lin et al. (2017) evaluates the hard-or-easy samples through predicted probability. Nie et al. (2019) applies adversarial learning to train the easy-or-hard representer. Both of them have their advantages, but all evaluate hard-or-easy samples based on the final segmentation result. As such, these methods cannot be directly applied to interactive segmentation. For example, if the model predicts the category probability of a voxel to be 0.8 ($p_i^{(t)} = 0.8$), and then takes an action $a_i^{(t)} = -0.1$. The next prediction will be 0.7 after interaction ($p_i^{(t+1)} = 0.7$). If $y_i = 1$, before and after results are all correctly predicted because $p_i^{(t)}$ and $p_i^{(t+1)}$ are all greater than 0.5. While for interactive segmentation, the probability is changing in the wrong direction. This change is what we called the *interactive misunderstanding* phenomenon and the formally definition is shown in following.

Definition 1 (*Interactive Misunderstanding*). For a binary classification problem, the sign of the foreground label $y = 1$ is denoted as positive, and the sign of the background label $y = 0$ is denoted as negative accordingly. In an interactive medical image segmentation task (i.e. a voxel-wise binary classification problem), for any voxel i , if the sign of the change of segmentation probability output by algorithm for two consecutive interaction steps $\text{sign}(\Delta(p^{(i)}))$ is not equal to $\text{sign}(y_i)$, then this phenomenon is defined as *interactive misunderstanding*.

Our proposed confidence network learns the confidence of given actions to avoid misunderstanding hint information and take accurate actions. We argue that confidential information can be used to regularized action choices and suggest more efficient interaction. The confidence network structure also uses P-Net Wang et al. (2018b) as the backbone. The confidence network contains six 3D convolutional blocks. Each of the blocks has two convolution layers, and the size of the convolution kernel is fixed as 3×3 in all these convolution layers. The detailed architecture is shown in Figure 4.

The confidence network is trained using the previous state and action as input and a confidence map as output. The confidence network is optimized by minimizing the summation of binary cross-entropy loss over actions (shown in (7)) at each time step t . Here we use C to denote the confidence

network, w_C denotes the parameters of the confidence network, while L_{BCE} denotes the binary cross-entropy loss:

$$\begin{aligned} &L_C(s^{(t)}, a^{(t)}; w_C) \\ &= L_{BCE}(C(s^{(t)}, a^{(t)}), g^{(t)}) + L_{BCE}(C(s^{(t)}, -a^{(t)}), 1 - g^{(t)}), \end{aligned} \quad (7)$$

where

$$g^{(t)} = \begin{cases} 0 & \text{if } a^{(t)} \oplus y^{(t)} == 1, \\ 1 & \text{otherwise,} \end{cases} \quad (8)$$

where $g^{(t)}$ means whether the direction of action is consistent with the label. $a \oplus b$ is defined as that the statement is only true if either $a > 0$ or $b > 0$, but not both.

One potential issue when training the confidence network is the imbalance of samples. Early in RL training, the error rate of actions output by the segmentation network is very high, while most actions are correct when the network gradually converges. Inspired by the training of the discriminator in generative adversarial networks, this study introduces symmetric samples into the (7) to speed up training. An obvious advantage is that it improves sample utilization efficiency because the confidence network can know what a ‘‘bad sample’’ is and get a corresponding ‘‘good sample.’’

3.4 Self-Adaptive Reward

Essentially, no matter applying focal loss or adversarial learning to the training of the segmentation model, they all try to weigh these hard samples to prevent the model dominated by easy samples. However, the interactive segmentation task differs from those fully automatic segmentation tasks because the interactive segmentation model needs to cooperate with the user and understand the user’s hint information. The action to refine the segmentation result shows how the segmentation model understands hint information. Therefore, it is necessary to ensure that the hint information is correctly understood and the correct action is taken.

Specifically, the previously described action-confidence learning can provide the segmentation model with a confidence map to alleviate the interactive misunderstanding phenomenon. By this confidence map, hard-or-easy samples can be better recognized as the confidence values for these ‘hard regions’ are relatively lower than in other regions. This paper formulates this voxel-level action-aware as the self-adaptive reward function, $r^{(t)}$, which is shown in (9), to adapt this mechanism to the training of MARL:

$$r^{(t)} = \sum_{i=1}^N \alpha (2 - c_i)^\beta \text{gain}_i^{(t)}, \quad (9)$$

where c_i is the value on the confidence map. The setting of hyperparameters α and β are described in Section 4.1. In addition, $\text{gain}_i^{(t)}$ denotes the relative gain of cross-entropy:

$$\text{gain}_i^{(t)} = \mathcal{X}_i^{(t-1)} - \mathcal{X}_i^{(t)}, \quad (10)$$

$$\mathcal{X}_i^{(t)} = -y_i \log(p_i^{(t)}) - (1 - y_i) \log(1 - p_i^{(t)}), \quad (11)$$

where $\mathcal{X}_i^{(t)}$ denotes the cross-entropy between current segmentation probability and ground truth. If an agent gets a positive reward, its current action is good, and the refined segmentation result is closer to the ground truth.

With the self-adaptive reward function in (9), the confidence value c_i of these wrong actions is lower, and they will be punished more when training our segmentation model.

3.5 Interaction Guide

Another challenge for interactive image segmentation is that users usually need to decide where to interact in many slices, which is very time-consuming, especially for 3D images. Our framework also provides users with an interaction guide mechanism to save the user’s interaction time. After refinement, our framework will suggest some possible areas for users to interact with next. Specifically, our framework will filter out those areas with low action confidence and provides them for users (see Figure 5). Firstly, the original 3D image will be segmented with super voxels; each super voxel can be regarded as a group of voxels that share common characteristics. We use simple linear iterative clustering (SLIC) [Achanta et al. \(2012\)](#) technique with $spacing = [2, 2, 2]$, $compactness = 0.1$ to generate supervoxels, and the number of initial supervoxels equals to 100 and gradually declines during the refinement iterations for training and testing. Secondly, the proposed algorithm will compute the mean action confidence in each super voxel and rank them in descending order. Finally, the top 5 super voxels will be marked and recommended to users. What users need to do is to select the best interaction positions from these super voxels.

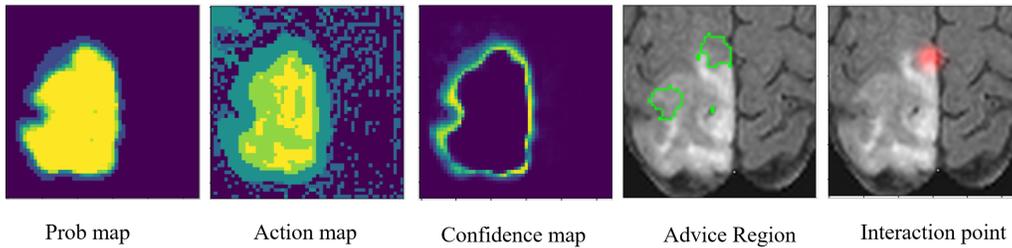


Figure 5: The illustration of interaction guide mechanism. The areas surrounded by the green lines are advice regions provided to users, and the red point is the real hint information selected from advice regions. The brighter the color (closer to yellow), the larger the positive value; on the contrary, the darker the color (closer to black), the smaller the negative value.

3.6 Simulated Label Generation

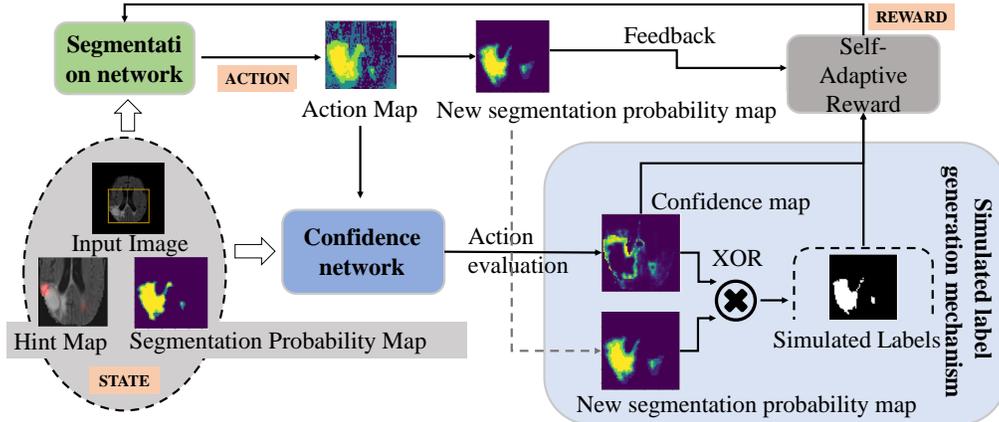


Figure 6: The illustration of simulated label generation. The simulated label generation mechanism utilizes the confidence map and the action map to generate the simulated label. The confidence map is used to calibrate the action, and the direction of the calibrated action is the simulated label of each voxel.

In medical imaging, the unlabeled data is much more than labeled data due to difficulties in labeling medical images. To address this lack of annotations, the proposed algorithm leverages the action confidence not only to improve the utilization efficiency of hint information but also to generate a simulated label, shown in Figure 6, for unlabeled data. We define the simulated label as follow:

$$\hat{y}^{(t)} = \begin{cases} 1 & \text{if } a^{(t)} \oplus c^{(t)} == 1, \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

The $\hat{y}^{(t)}$ is the simulated voxel-level label generated from the confidence map and detailed illustration is shown in Figure 6. Based on this mechanism, this study considers using these unlabeled data to assist in training the segmentation network. Specifically, this study divides the dataset into two parts during the training stage: the first part contains voxel-level annotation information, and the second part contains only hint information during training. The labeled data is trained as above, while the unlabeled data is trained with the simulated label. During the training stage, the backward of gradients will be masked:

$$d\theta_v = M \nabla_{\theta_v} A(s^{(t)}, a^{(t)})^2, \quad (13)$$

$$d\theta_p = M \nabla_{\pi} (a^{(t)} | s^{(t)}) A(a^{(t)}, s^{(t)}), \quad (14)$$

where

$$M = I(\max(c^{(t)}, 1 - c^{(t)}) > \delta). \quad (15)$$

$A(s^{(t)}, a^{(t)})$ is the advantage (which is defined in Mnih et al. (2016)) at time step t of taking $a^{(t)}$ in condition of state $s^{(t)}$, which indicates the actual accumulated reward without being affected by the state and reduces the variance of gradient. A mask, M , is used to constraint the training of unlabeled data. The gradients of unlabeled data backward only when the action confidence exceeds the threshold δ (which is gradually increased during the training process). Unlike traditional pseudo-label training, the supervised signal does not come from the segmentation but the confidence network. These filtered data with hint information are more valuable and provide more accurate supervised signals. Generally, the training process with labeled and simulated labeled data is carried out simultaneously. When using the simulated label generation mechanism, the pseudocode is shown in Algorithm 1.

4 Experiments and results

4.1 Dataset and Implementaion details

To comprehensively evaluate our proposed method, we apply our algorithm to four 3D medical image datasets. All datasets are divided into two parts: D_{train}/D_{test} . The details of these datasets are as follows: 1) **BraTS2015**: Brain Tumor Segmentation Challenge 2015 Menze et al. (2015) contains 274(234/40) multiparametric MRI(Flair, T1, T1C, T2) from brain tumor patients. In our task, we only use the Flair image and segment the whole brain tumor. 2) **BraTS2020**: Brain Tumor Segmentation Challenge 2020 Menze et al. (2015) contains 285(235/50) multiparametric MRI(Flair, T1, T1C, T2) from brain tumor patients. In our task, we only use the Flair image and segment the whole brain tumor. 3) **MM-WHS**: Multi-Modality Whole Heart Segmentation Zhuang & Shen (2016) contains 24(20/4) multi-modality whole heart images covering the whole heart substructures. In our task, this study chooses to segment the left atrium blood cavity. 4) **Medical Segmentation Decathlon**: This is a generalisable 3D semantic segmentation datasets containing different organ segmentation tasks Simpson et al. (2019). This study chooses to use the spleen and liver dataset, which provides 61(41/20) and 106(96/10) CT images respectively.

Algorithm 1 MECCA: Interactive Medical Image Segmentation with Self-Adaptive Confidence Calibration

```
1: Initialize the segmentation network  $S$  with  $\theta$ ;  
2: Initialize the confidence network  $C$  with  $w$ ;  
3: for every sample in labeled datasets do  
4:   Set the segmentation probability of each voxel to 0.5;  
5:    $s^{(0)} \leftarrow (x, p^{(0)}, h^{(0)})$ ;  
6:   for every interaction time step  $t$  do  
7:     Take action  $a^{(t)} \leftarrow S(s^{(t)})$ ;  
8:     Get reward  $r^{(t)}$  and observe the next state  $s^{(t+1)}$ ;  
9:     Compute the gradient of  $S$  via (4) and (6);  
10:    Compute the gradient of  $C$  via (7);  
11:   end for  
12:   Get a sample  $(x', y')$  from unlabeled dataset;  
13:   Initialize the state:  $s'^{(0)} \leftarrow (x', p'^{(0)}, h'^{(0)})$ ;  
14:   for every interaction time step  $t$  do  
15:     Take action  $a'^{(t)} \leftarrow S(s'^{(t)})$ ;  
16:     Generate the simulated label  $\hat{y}^{(t)}$  by (12);  
17:     Observe the reward and next state  $s'^{(t+1)}$ ;  
18:     Compute the gradient of  $S$  via (13) and (14).  
19:   end for  
20: end for
```

We implement our method with PyTorch [Paszke et al. \(2019\)*](#). The segmentation and confidence network are both initialized by Xavier [Glorot & Bengio \(2010\)](#) method, and learning rates are initialized to $1e-4$. Other parameters are set as follow: $T = 5$, $\gamma = 0.95$, $\alpha = 0.8$, $\beta = 1$. The mask M ranges from 0.85 to 0.99 and ascents 0.00025 at every epoch. Adam [Kingma & Ba \(2015\)](#) is adopted as the optimizer. The original image is cropped by the bounding box based on the ground truth with a random extension in the range of 1 to 11 voxels. Each image is then resized and normalized to $55 * 55 * 30$. The data is augmented by random flip and rotation. The proposed algorithm training time with one Nvidia 2080ti GPU varies from 5 to 13 hours for different datasets.

4.2 Interaction Settings and Evaluation Metrics

During the process of interaction, this study adopts the edge points as the hint information for two reasons. On the one hand, click operation saves more time than scribbles or other interactive ways. On the other hand, edge points can provide more information about object edges as the bottleneck of medical image segmentation is usually the inaccurate segmentation of object edges. Specifically, this study provides each method with 45 points[‡] during the whole interaction process. At every step, users will click some edges which are not correctly predicted. This study generates a 3D Gaussian (with 8-voxels kernel size) centered on each of the edge points as the hint map to let networks receive hint information. Then the hint map will be input to the segmentation network as part of the state. This study will use the Dice score and the average symmetric surface distance (ASSD) to evaluate the performance of the segmentation result. According to these evaluation metrics, doctors can judge the patient’s condition:

$$\text{Dice}(S_p, S_g) = (2|S_p \cap S_g|) / (|S_p| + |S_g|),$$

*The demo video of MECCA algorithm is available at <https://bit.ly/mecca-demo-video>.

‡Please refer to Section IV-C for details.

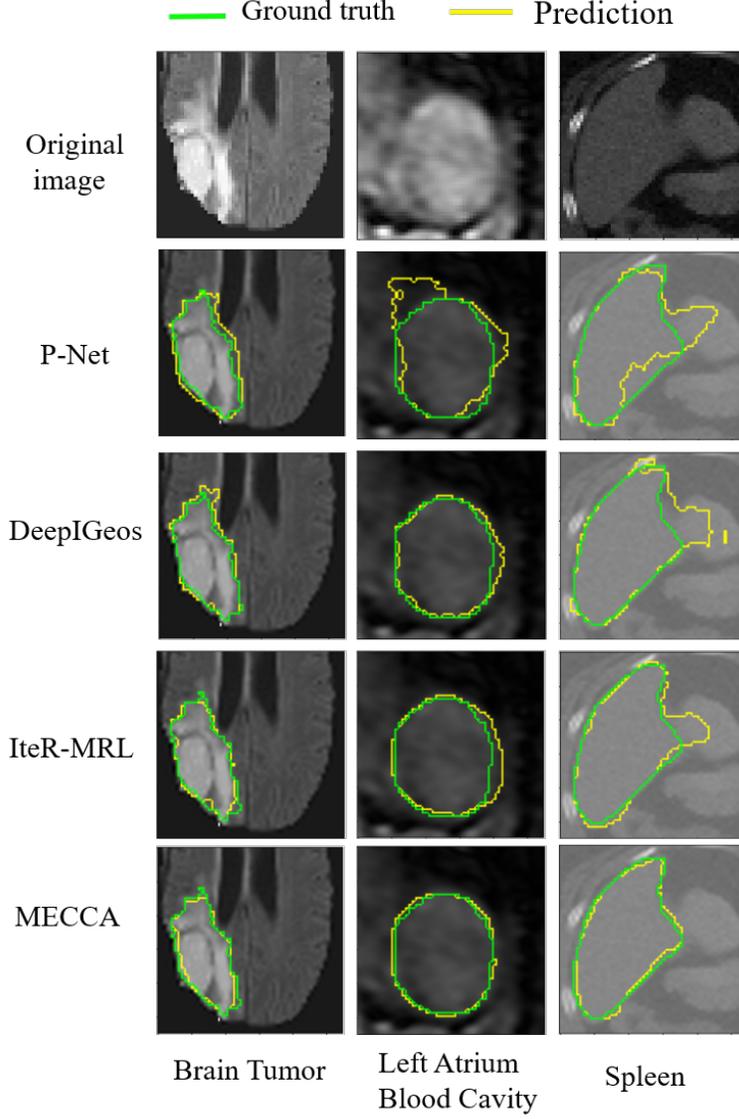


Figure 7: Visualization of segmentation results by different methods. The green lines represent the boundaries of ground truth, and the yellow lines represent the predictive boundaries.

where S_p, S_g denote the prediction of an algorithm and the ground truth respectively.

$$ASSD = \left(\sum_{i \in \mathcal{S}_a} d(i, \mathcal{S}_b) + \sum_{i \in \mathcal{S}_b} d(i, \mathcal{S}_a) \right) / (|\mathcal{S}_a| + |\mathcal{S}_b|),$$

where \mathcal{S}_a and \mathcal{S}_b represent the set of surface points of the segmentation result predicted by the algorithm and the ground truth, respectively. $d(i, \mathcal{S}_b)$ is the shortest Euclidean distance between i and \mathcal{S}_b . The Dice score and ASSD in all tables are the average value of five test results of algorithms.

Table 1: Dice scores at each interaction step by different methods. The value in parentheses represents the improvement relative to the previous step.

Step	1	2	3	4	5
DeepIGeos	87.36	87.53 (+0.17)	87.67 (+0.14)	87.99 (+0.32)	88.32 (+0.33)
InterCNN	87.21	88.59 (+1.38)	88.54 (-0.04)	88.39 (-0.16)	88.26 (-0.12)
IteR-MRL	84.56	85.35 (+0.80)	88.15 (+2.80)	88.11 (-0.08)	88.94 (+0.83)
MECCA	86.49	88.53 (+2.03)	89.56 (+1.03)	89.12 (-0.44)	90.29 (+1.17)

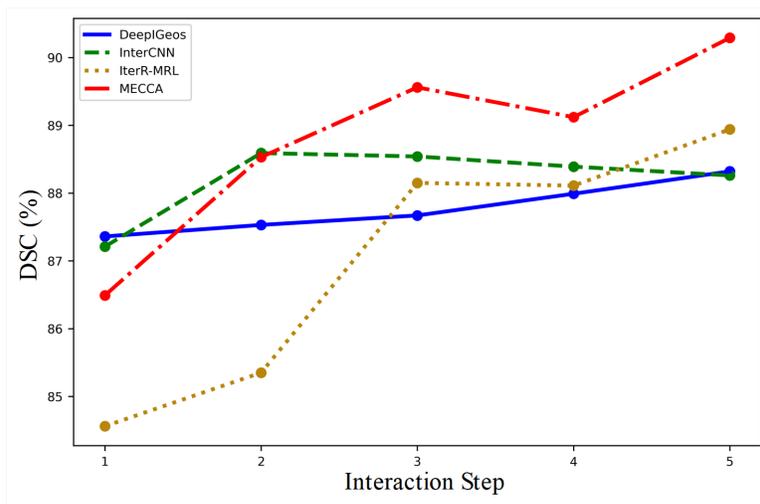


Figure 8: Visualization of the performance improvement of different interactive medical segmentation methods at the different interaction steps. All these testing results are performed on the BraTS2015 dataset.

4.3 Comparisons with State-of-the-art Methods

We compare MECCA with five state-of-the-art interactive segmentation methods: DeepIGeos Wang et al. (2018b)[§], InterCNN Bredell et al. (2018)[¶], IteR-MRL Liao et al. (2020) and BS-IRIS Ma et al. (2020). InterCNN is the multi-step version of DeepIGeos. We also introduce the SOTA method, U-Net Ronneberger et al. (2015)^{||}, of medical image segmentation as a comparable baseline. Table 2 show the quantitative comparison of the six interactive segmentation methods on different datasets. For a fair comparison, all CNN-based methods adopt the same network structure (P-Net), which is proposed in Wang et al. (2018b). We can see that our proposed MECCA performs better than other state-of-the-art methods on all three datasets. This study also visualizes the results in Figure 7, which shows that our method is more accurate in edge segmentation.

To demonstrate that our method can take advantage of hint information more efficiently, this study also compares the relative improvement of different methods at each interaction step. Due to the different interaction ways of different methods, this study sets up a unified interaction process: all

[§]<https://github.com/taigw/GeodisTK>.

[¶]<https://github.com/gbredell/interCNN>.

^{||}https://github.com/liyun-lu/unet_and_vnet.

Table 2: Quantitative comparison of 3D medical image datasets segmentation by different methods. In particular, the P-Net is the method without hint information. Significant improvement (p-value < 0.05) is marked in bold.

Methods	BraTS2020		BraTS2015		MM-WHS		Spleen		Liver	
	Dice(%)	ASSD(pixels)								
P-Net	83.67±8.35	5.78±4.01	84.00±12.01	5.10±3.72	81.40±1.48	3.28±0.45	88.08±2.25	4.25±2.07	35.89±2.61	34.46±23.82
U-Net	84.72±10.42	4.09±3.89	84.66±11.25	6.17±4.69	80.96±1.65	3.72±0.39	87.95±2.87	5.12±1.09	56.00±1.93	22.38±22.42
DeepIGeos	88.54±.97	2.11±1.30	88.32±5.34	2.28±1.24	88.48±0.71	1.53±0.18	91.97±1.51	0.93±0.46	48.57±2.52	10.28±3.45
InterCNN	88.39±6.01	2.01±1.09	88.26±7.07	1.81±2.09	87.85±1.15	0.80±0.15	93.52±0.94	0.54±0.83	59.92±2.20	5.95±2.76
IteR-MRL	89.22±4.65	2.07±0.91	88.94±4.81	1.41±0.22	89.55±0.87	0.90±0.11	91.50±1.34	0.67±0.21	62.29±1.93	0.87±0.59
BS-IRIS	90.47±5.23	1.82±0.33	89.74±3.86	1.61±0.42	89.12±0.98	1.19±0.16	92.35±1.13	0.54±0.19	67.25±2.01	4.34±1.18
MECCA	91.02±5.86	1.15±0.20	90.29±5.07	1.50±0.33	90.39±5.89	0.80±0.01	94.96±1.44	0.30±0.16	71.46±1.41	2.36±0.99

these methods have 5 interaction steps and receive 25 identical points at the first step to generate the initial segmentation result. After that, users will click 5 points at each step in the following 4 interaction steps (total of $25 + 5 \times 4 = 45$ points), and these methods will iteratively refine previous segmentation results. The DeepIGeos method does not model the interaction sequence, and it just combines current and previous hint information to refine the previous segmentation. The experimental results are shown in Table 1 and Figure 8.

The results show that MECCA can perform better under the same amount of hint information and improve more dice scores at most steps. Compared with CNN-based methods, the main advantage of RL-based methods is that they can always keep notable improvement. There are two reasons for this result. The first one is that RL-based methods model the whole interaction process to avoid interaction conflict. The second reason is the relative entropy-based reward which encourages the model to keep refining results. However, we should realize, the RL-based methods still can not guarantee the high confidence of the corrective actions. As we can see in Figure 8, the performance of IteR-MRL is not as good as other methods at the beginning, which is caused by numerous incorrect actions. However, after learning the confidence of actions and applying the self-adaptive reward to update the segmentation model, our proposed MECCA can perform well at each step and keep significantly refining the result.

Figure 13(a) presents the visualization of the segmentation process of our proposed method. Since our proposed method models the whole interaction process, Figure 13(a) shows the current interaction step and its previous and next interaction steps. At each step, the second column is the confidence map. This study finds that the confidence value of object edges is always lower than in other regions, and these regions will receive more ‘punishment’ when rewards are generated. This study observes that MECCA can gradually correct the edges around the user clicks (the red regions).

Table 3: The DICE of our method varies with the number of interactions under different difficult cases.

	Interaction times				
	1	2	3	4	5
The easy case	94.05	95.15	95.38	95.91	96.25
The hard case	41.45	47.38	65.78	77.79	79.37

Table 4: MECCA’s tolerance for inaccurate interaction points. Significant improvement (p-value < 0.05) is marked in bold.

	DICE (%)	ASSD (pixels)
Disturbed Interactions	88.75±9.72	1.11±0.23
Non-Disturbed Interactions	90.29±5.07	1.50±0.33

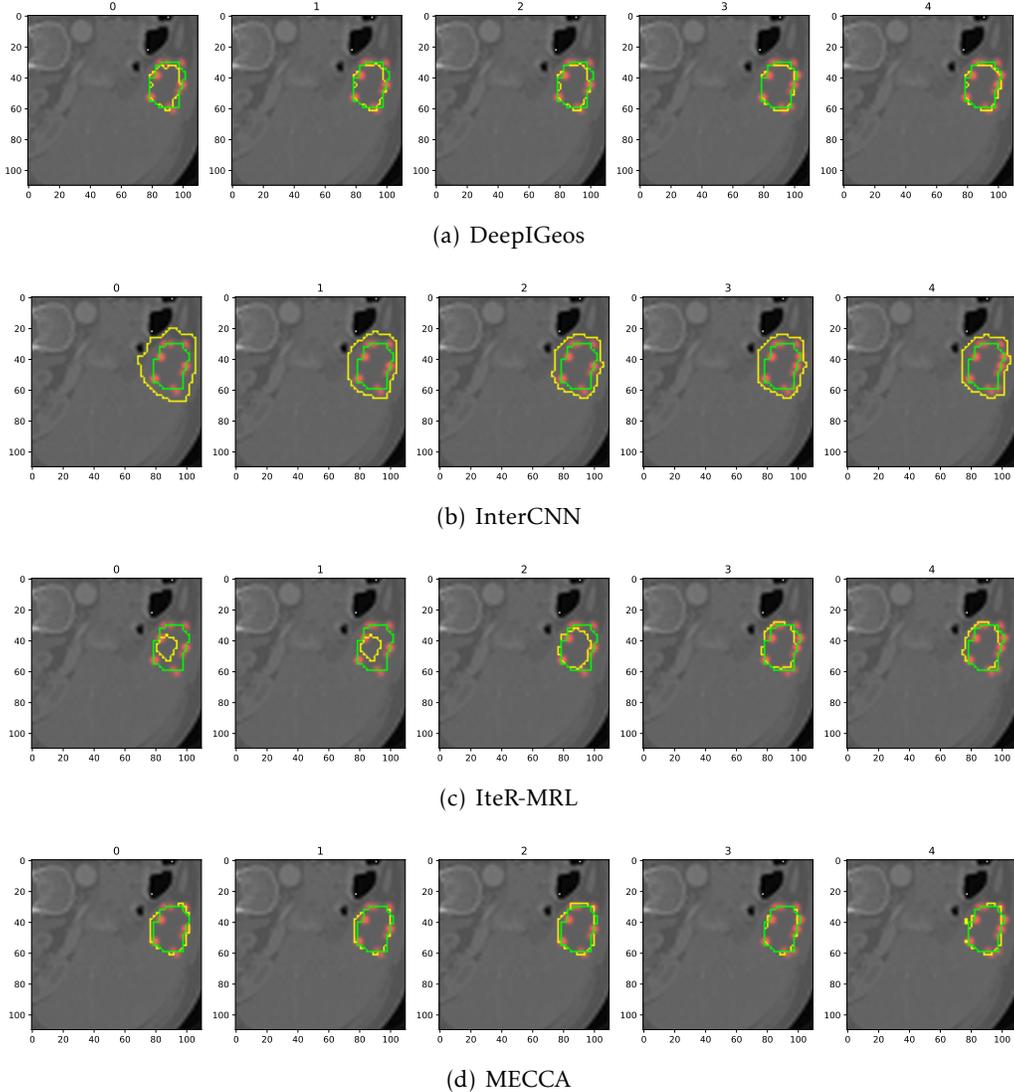


Figure 9: The results of different methods response to the same user interactions according to the same initial segmentation on 4/10 testing instance, 7/30 channel, Liver 2 of MSD dataset.

Besides, we select the best and worst-performing samples from the test set for more analysis. The results are shown in the Table 3. The improvement of easy samples through interaction is not considerable, but the improvement of hard samples through interaction can reach twice as much. However, significantly more interactions are required for hard samples to achieve the same segmentation accuracy of easy samples. Moreover, it is very difficult and time-consuming to accurately mark edge points in a real scene, and is less practical to ask user to click the accurate edge points. To verify MECCA’s tolerance for inaccurate edge points, we conducted a robustness study under the same settings as Table 2. Specifically, during each interaction in the training phase, random noises are added to simulated edge points. The noise range is plus or minus 2 voxels in the three directions of x , y , and z . In this way, the edge points used by the algorithm will be randomly selected from 64 ($4 \times 4 \times 4$) voxels within the real edge point neighborhood. This disturbance radius can cover the edge ambiguity area in most cases. In the testing phase, we also adopted the same disturbance operation, and the final results are shown in Table 4.

From the table, it can be seen that the perturbed edge points will make the DICE value of the

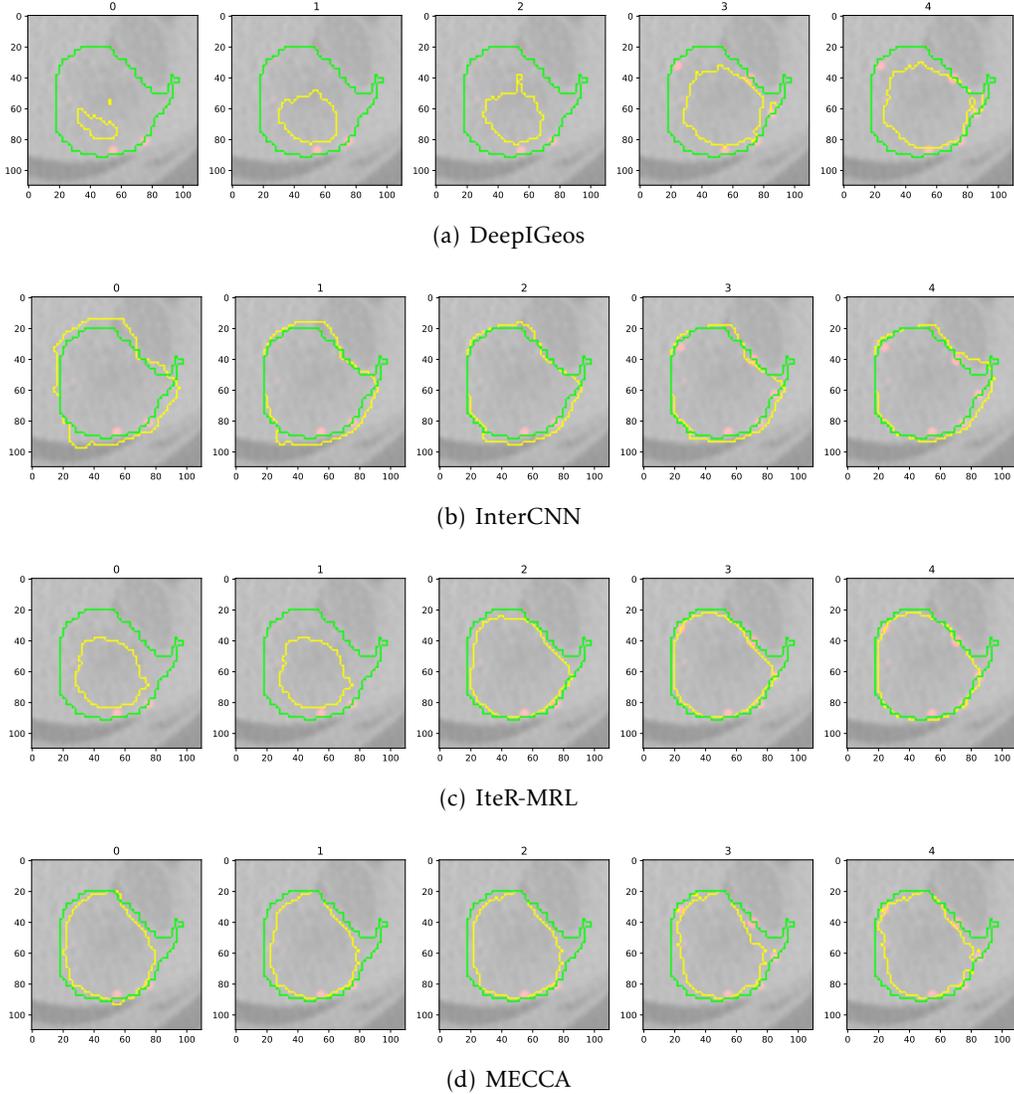


Figure 10: The results of different methods response to the same user interactions according to the same initial segmentation on 6/10 testing instance, 18/30 channel, Liver 2 of MSD dataset.

algorithm drop and have a more considerable variance. However, the maximum value of DICE exceeds the accurate edge points. The perturbed edge points can even exceed the performance of accurate edge points in the ASSD value. We believe that there may be two reasons for the above phenomenon. First, MECCA uses adaptive confidence calibration to improve the information-misunderstanding of the iterative algorithm, but it also makes the algorithm more “conservative”. The perturbed edge point information may exceed the actual object boundary, which can make the segmentation effect of our method better or more “radical”; Second, perturbing the interactive information during the training phase can enable the reinforcement learning algorithm to explore the environment (medical images). Existing SOTA reinforcement learning algorithms generally impose entropy constraints on the policy, enhance the randomness of the policy, and encourage exploration. Moreover, our perturbation of interactive information will indirectly affect the policy of the algorithm. In a word, MECCA has good robustness to inaccurate edge points.

Finally, we show how different methods respond to the same user interactions according to the

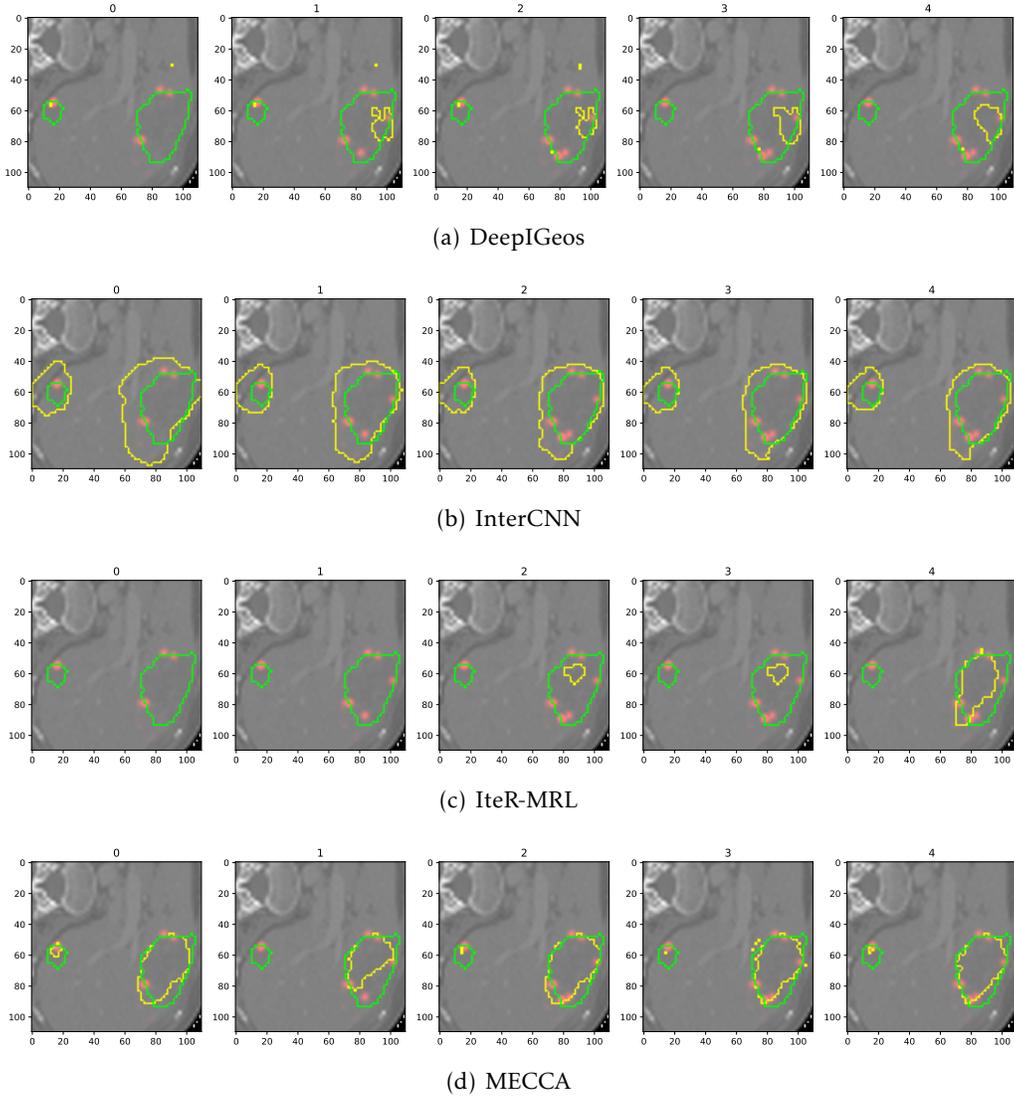


Figure 11: The results of different methods response to the same user interactions according to the same initial segmentation on 7/10 testing instance, 10/30 channel, Liver 2 of MSD dataset.

same initial segmentation, especially for hard cases. We selected some images with poor performance in baselines or MECCA as the research objects. Specifically, the generation mechanism of the same user interactions is described as following. Total 45 hint points are randomly selected from the intersection area of the boundary of the foreground object and the error region of the initialization segmentation. Then, these 45 points are allocated to the 5 interactive steps of all methods according to the combination of 25, 5, 5, 5, and 5. The number of hint points used in each interactive step is the same as that of all methods during training and testing. The results are shown in Figure 9-11. It can be seen from the results that MECCA can use the hint points information stably in all cases. The *interactive misunderstanding* phenomenon arises in other methods. These method either ignore the expert's correction information, or even be adversely affected by correction information as in figures.

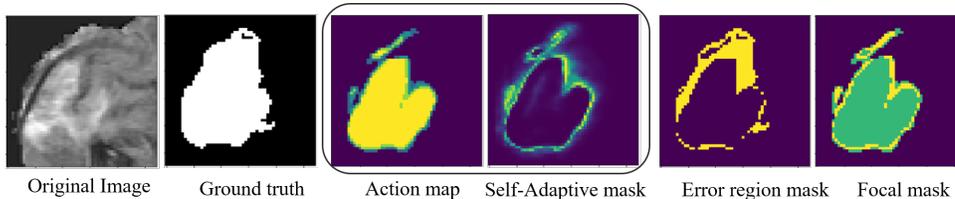


Figure 12: Visualization of the self-adaptive mask, the Error region mask and the Focal mask. All masks are obtained at the first interaction step. The brighter the color (closer to yellow), the larger the positive value; on the contrary, the darker the color (closer to black), the smaller the negative value.

4.4 Comparison of Different Weighted Rewards

One main contribution of our work is that the self-adaptive reward mechanism is further proposed based on the confidence map. This mechanism makes different actions have different feedback levels so that the segmentation network can identify wrong actions as much as possible. To measure the impact of different reward weighting ways on the segmentation result, this study compares MECCA with the original IteR-MRL without weighting reward, the Focal-Reward IteR-MRL (Focal-IteR-MRL), and the one with weighted reward in segmentation error regions (Err-IteR-MRL). The Focal-Reward IteR-MRL adapt the idea of Focal loss [Lin et al. \(2017\)](#):

$$FocalReward_i^{(t)} = \mathcal{X}_i^{(t-1)} - \mathcal{X}_i^{(t)}, \quad (16)$$

where

$$\mathcal{X}_i^{(t)} = \begin{cases} -(1 - p_i)y_i \log(p_i^{(t)}) & \text{if } y_i = 1, \\ -p_i(1 - y_i) \log(1 - p_i^{(t)}) & \text{if } y_i = 0. \end{cases} \quad (17)$$

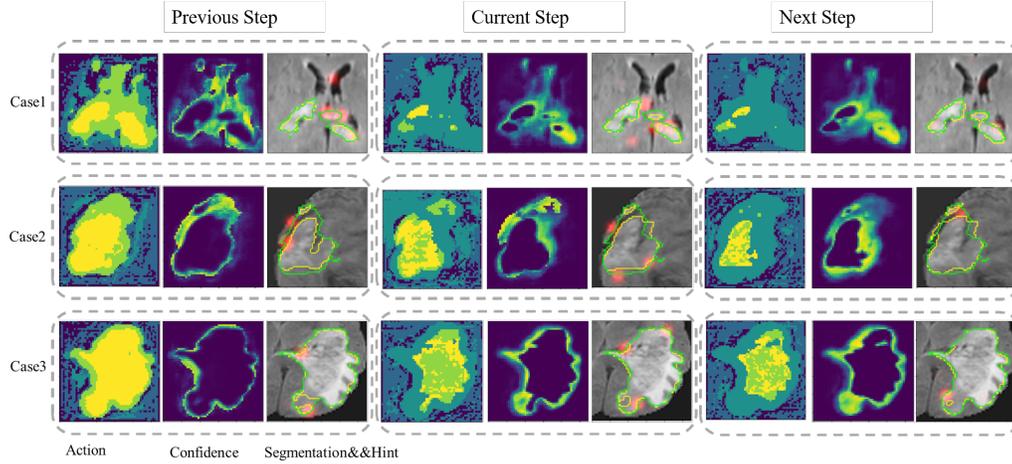
The reward function of Err-IteR-MARL is a linear scaling of reward of IteR-MARL:

$$ErrReward_i^{(t)} = \begin{cases} (1 + \lambda_i)gain_i^{(t)}, & \text{if voxel } i \text{ is misclassified,} \\ gain_i^{(t)}, & \text{otherwise,} \end{cases} \quad (18)$$

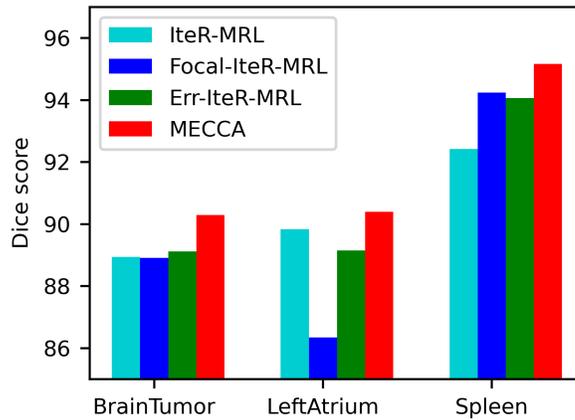
where λ_i is a positive real number hyperparameter. The performance of different weighting rewards is shown in Figure 13(b). Action-based reward weighting methods are more suitable for interactive segmentation since the improvement of our proposed method is more notable than other methods. Also, this study visualizes different masks that weigh the primary reward in Figure 12. As we can see, the focal mask pays more attention to regions with low prediction probabilities, and the errors region mask contains all regions segmented incorrectly in the previous result. It is not enough for the focal mask to obtain more structured information by only considering the probability of prediction. As such, the performance of Focal-IteR-MRL is erratic, even the poorest on the left atrium dataset. The performance of Err-IteR-MRL is more stable than other methods, but the improvement is not notable. It is because that the weighting way of Err-IteR-MRL is based on the segmentation result while the refinement process of the interactive image segmentation is based on actions.

4.5 Weakly-supervised Interactive Segmentation

As mentioned in Section 3.6, MECCA can reduce the dependence on voxel-level annotations of images by using the simulated label generated from the action confidence map. This study validates Algorithm 1 on the BraTS2015 dataset by randomly selecting different proportions of samples as fully labeled data and using the rest of the training images as unlabeled data, which only provided



(a)



(b)

Figure 13: (a) The segmentation process of MECCA. At each step, the first column is the action map, the second column is the confidence map, and the third column is the segmentation result with the user’s hint information. (b) Average Dice scores of methods with different weighting rewards. The reward function of IteR-MRL is not weighted; the reward function of Focal-IteR-MRL is weighted via (16), the reward function of Err-IteR-MRL is weighted in segmentation error regions, and the the reward function of MECCA is weighted through action confidence. The brighter the color (closer to yellow), the larger the positive value; on the contrary, the darker the color (closer to black), the smaller the negative value.

hint information when interaction. This study compares MECCA against the P-Net (without hint), DeepIGeos, and IteR-MRL Liao et al. (2020). However, only MECCA uses unlabeled data, and the remaining three baselines only use a fixed proportion of labeled data. This does not fully demonstrate the performance of MECCA on semi-supervised problems. For this reason, we additionally introduce a state-of-the-art semi-supervised method UA-MT Yu et al. (2019). Table 5 shows the results of different methods. MECCA achieves a Dice score of 87.14% with only 12.5% labeled data, and it performs better than the other methods with 25% labeled data.

Since UA-MT is not an interactive segmentation algorithm, its absolute performance is worse than

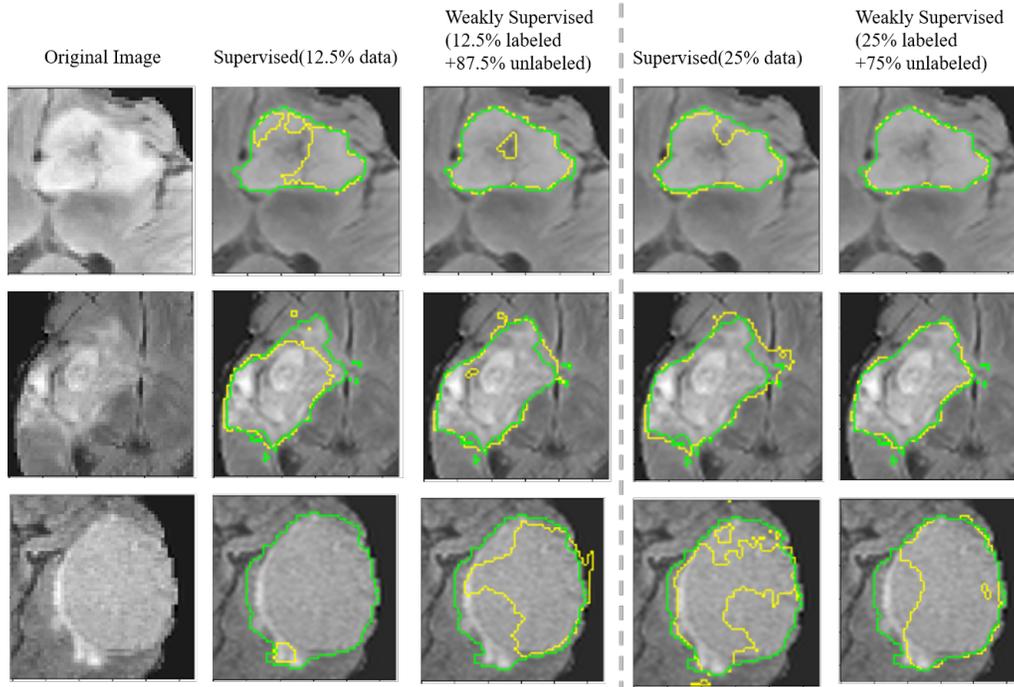


Figure 14: Qualitative segmentation results of MECCA for BraTS2015 validation set. The first column shows the cropped original images. The second to fifth columns respectively show: the result of supervised learning only using 12.5% of labeled data, the results of weakly-supervised learning using 12.5% labeled and other unlabelled images, the result of supervised learning only using 25% of labeled data, the results of weakly-supervised learning using 25% labeled and other unlabelled images. Both supervised learning and weakly supervised learning is based on our proposed method.

interactive segmentation baselines but better than the non-interactive method P-Net. In addition, it can be seen from the last column of results that UA-MT has the most negligible performance loss on different proportions of labeled data. After MECCA introduces the simulated label generation mechanism, the performance loss can also be controlled at a similar magnitude to UA-MT. Interestingly, DeepIGeos can maintain a good performance loss without using a semi-supervised learning mechanism. Combined with the performance loss of IteR-MRL, we can conclude that in the semi-supervised interactive segmentation task, the interactive misunderstanding phenomenon will exacerbate the performance loss caused by the missing data. DeepIGeos and MECCA alleviate the interactive misunderstanding phenomenon through hard constraints and self-adaptive confidence calibration, respectively, so both of them can achieve better results under semi-supervised settings. However, it is worth mentioning that DeepIGeos does not consider multi-step interactions and the relationship between consecutive interactions. It cannot fully utilize long-term interactive information, so it has a large gap with MECCA in absolute performance.

Figure 14 shows the qualitative segmentation results of our method trained with different data sizes. It shows that the model trained with 12.5% labeled data can only capture the main region of the tumor. However, the model is unable to distinguish the infiltration areas of the tumor. For instance, the boundaries of tumors in the figure are hard to distinguish as it is more similar to these healthy regions. In this case, the models trained with little data tend to ignore these boundary regions of the tumor, while the model trained with both labeled data and unlabeled data can get more smooth and accurate boundaries. This phenomenon is mainly due to the distribution of the training set is not consistent with the validation set. The main advantage of the model trained with both labeled and

Table 5: Quantitative comparison between MECCA and other methods on BraTS2015 Dataset of different sizes. Significant improvement (p-value < 0.05) is marked in bold.

Data amount	1/8	1/4	1/3	1/2	1/1	$\Delta(1/8,1/1)$
P-Net	75.86	80.83	80.88	83.02	84.00	10.73%
DeepIGeos	85.50	85.90	87.70	87.60	88.32	3.30%
IteR-MRL	84.36	86.54	87.38	88.66	88.94	5.07%
UA-MT	83.08	84.47	84.62	84.39	84.66	1.90%
MECCA	87.14	88.23	88.31	89.17	90.29	3.60%

unlabeled data is that it can minimize the gap between training and validation data.

Table 6: DICE of the algorithm under different number of interactions and percentage of labeled data.

% Labeled Data	Interaction Steps				
	3	4	5	6	7
25%	84.03	86.18	88.23	87.32	88.24
33%	85.32	87.60	88.31	87.99	88.53
50%	85.39	87.07	89.17	88.93	89.13
100%	89.93	90.40	90.29	89.85	90.79

We additionally test MECCA’s performance under different interaction times and the percentage of labeled data, and the results are shown in Table 6. Based on the experimental results, we can obtain the following conclusions: 1) if there is enough labeled data, MECCA can achieve good results after a few interactions; 2) the number of interactions required by MECCA to achieve the same performance is roughly inversely proportional to the number of labeled data; 3) although MECCA requires more interactions (about 2-3 times the number of interactions) when only part of the labeled data is available, it can approach the algorithm’s performance trained with all labeled data in the end.

4.6 Ablation Study

Table 7: Ablation study of our proposed algorithm on the BraTS2015 dataset with 25% labeled data and 75% unlabeled data . Significant improvement (p-value \leq 0.05) is marked in bold.

Self-adaptive reward	Interaction guide	Simulated labels generation	Dice(%)
			86.54
✓			87.21
	✓		86.56
		✓	88.01
✓	✓		87.24
✓	✓	✓	88.23

Table 7 and 8 shows the impact of different mechanisms to the MECCA. It is clear that the stimulative labels generation mechanism and self-adaptive reward mechanism significantly improve the algorithm’s performance. The stimulative labels generation mechanism can achieve a 1.47% and

Table 8: Ablation study of our proposed algorithm on the Liver 2 dataset of MSD with 25% labeled data and 75% unlabeled data. Significant improvement (p-value < 0.05) is marked in bold.

Self-adaptive reward	Interaction guide	Simulated labels generation	Dice(%)
			66.43
✓			67.45
	✓		66.56
		✓	67.93
✓	✓		67.59
✓	✓	✓	69.99

3.56% gain over the initial algorithm without any mechanisms for the Brats2015 and Liver 2 dataset, respectively. The main reason for this significant improvement is the utilization of these unlabeled data.

Table 9: Quantitative comparison of interactive segmentation methods in computational and interaction time.

Methods	Total inference time per iteration	Interaction time per iteration
DeepIGeos	510ms	226ms
InterCNN	532ms	237ms
IteR-MRL	869ms	235ms
BS-IRIS	990ms	250ms
MECCA	631ms	123ms

Table 10: Quantitative comparison of interactive segmentation methods in interaction time on realistic segmentation platform operated by experts. Demo video of MECCA is available at <https://bit.ly/mecca-demo-video>.

Interaction Step(s)	Interaction Time (<i>Second</i> , Real World)				
	1 (25 points)	2 (5 points)	3 (5 points)	4 (5 points)	5 (5 points)
DeepIGeos	70	15	15	8	8
InterCNN	76	19	18	8	7
IteR-MRL	75	17	17	8	8
BS-IRIS	80	20	20	11	10
MECCA	41	12	11	5	5

Table 9 shows the comparison of interactive segmentation methods in computational and interaction time. The result shows that the MECCA only takes half the time of other methods when performing an interaction. The main contribution to the reduction of interaction time is the interaction guide mechanism mentioned in Section 3.5. Furthermore, to explore the efficiency of MECCA in authentic tasks, we designed an interactive software platform. We asked oncologists to use it for natural interactive segmentation, and the results are in Table 10. It can be seen from the results that the MECCA algorithm also takes only about half of the time of the baselines due to the interaction guidance. In addition, as the segmentation results become more and more accurate, the interaction time required gradually decreases.

5 Conclusion

This paper presents a novel action-based confidence learning method for interactive 3D image segmentation. Specifically, this paper proposes to learn the confidence of actions that continuously refine the segmentation result during the interaction process so that hint information can be used more effectively. Based on this, the self-adaptive reward is proposed for the segmentation module, which can prevent the interactive misunderstanding phenomenon during the interaction and help reduce the time cost of interaction by providing users with the advice regions to interact next. Besides, the confidence map can also replace the ground truth to generate feedback for the unlabeled samples for the segmentation module. These samples without voxel-level annotations can also be used to train our model. By integrating these components, MECCA is demonstrated to medical segmentation efficiently via less voxel-level annotated samples.

References

- Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., and Süsstrunk, S. SLIC Superpixels Compared to State-of-the-Art Superpixel Methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(11):2274–2282, 2012.
- Acuna, D., Ling, H., Kar, A., and Fidler, S. Efficient Interactive Annotation of Segmentation Datasets with Polygon-RNN++. In *CVPR*, 2018.
- Atanov, A., Ashukha, A., Molchanov, D., Neklyudov, K., and Vetrov, D. Uncertainty estimation via stochastic batch normalization. In *ICLR*, 2018.
- Boykov, Y. Y. and Jolly, M.-P. Interactive Graph Cuts for Optimal Boundary & Region Segmentation of Objects in ND Images. In *ICCV*, 2001.
- Bredell, G., Tanner, C., and Konukoglu, E. Iterative Interaction Training for Segmentation Editing Networks. In *MLMI International Workshop*, 2018.
- Castrejon, L., Kundu, K., Urtasun, R., and Fidler, S. Annotating Object Instances with a Polygon-RNN. In *CVPR*, 2017.
- Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., and Yuille, A. DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):834–848, 2018.
- Criminisi, A., Sharp, T., and Blake, A. GeoS: Geodesic Image Segmentation. In *ECCV*, 2008.
- DeVries, T. and Taylor, G. W. Learning confidence for out-of-distribution detection in neural networks. *arXiv preprint arXiv:1802.04865*, 2018a.
- DeVries, T. and Taylor, G. W. Leveraging uncertainty estimates for predicting segmentation quality. *arXiv preprint arXiv:1807.00502*, 2018b.
- Dietterich, T. G. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pp. 1–15. Springer, 2000.
- Feng, R., Zheng, X., Gao, T., Chen, J., Wang, W., Chen, D. Z., and Wu, J. Interactive Few-shot Learning: Limited Supervision, Better Medical Image Segmentation. *IEEE Transactions on Medical Imaging*, 2021. doi: 10.1109/TMI.2021.3060551.

- Furuta, R., Inoue, N., and Yamasaki, T. PixelRL: Fully Convolutional Network With Reinforcement Learning for Image Processing. *IEEE Transactions on Multimedia*, 22(7):1704–1719, 2020.
- Gal, Y. and Ghahramani, Z. Dropout as a Bayesian Approximation: Representing model Uncertainty in Deep Learning. In *ICML*, 2016.
- Glorot, X. and Bengio, Y. Understanding the Difficulty of Training Deep Feedforward Neural Networks. In *AISTATS*, 2010.
- Grady, L. Random Walks for Image Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(11):1768–1783, 2006.
- Hung, W.-C., Tsai, Y.-H., Liou, Y.-T., Lin, Y.-Y., and Yang, M.-H. Adversarial Learning for Semi-supervised Semantic Segmentation. In *BMVC*, 2018.
- Jungo, A. and Reyes, M. Assessing reliability and challenges of uncertainty estimations for medical image segmentation. In *MICCAI*, 2019.
- Kamnitsas, K., Bai, W., Ferrante, E., McDonagh, S., Sinclair, M., Pawlowski, N., Rajchl, M., Lee, M., Kainz, B., Rueckert, D., et al. Ensembles of multiple models and architectures for robust brain tumour segmentation. In *International MICCAI brainlesion workshop*, pp. 450–462, 2017a.
- Kamnitsas, K., Ledig, C., Newcombe, V. F., Simpson, J. P., Kane, A. D., Menon, D. K., Rueckert, D., and Glocker, B. Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *Medical Image Analysis*, 36:61–78, 2017b.
- Kendall, A. and Gal, Y. What uncertainties do we need in bayesian deep learning for computer vision? In *NeurIPS*, 2017.
- Kingma, D. P. and Ba, J. Adam: A Method for Stochastic Optimization. In *ICLR*, 2015.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. In *NIPS*, 2017.
- Li, W., Wang, G., Fidon, L., Ourselin, S., Cardoso, M. J., and Vercauteren, T. On the Compactness, Efficiency, and Representation of 3D Convolutional networks: Brain Parcellation as a Pretext Task. In *IPMI*, 2017.
- Liao, X., Li, W., Xu, Q., Wang, X., Jin, B., Zhang, X., Wang, Y., and Zhang, Y. Iteratively-Refined Interactive 3D Medical Image Segmentation with Multi-Agent Reinforcement Learning. In *CVPR*, 2020.
- Lin, D., Dai, J., Jia, J., He, K., and Sun, J. Scribblesup: Scribble-supervised Convolutional Networks for Semantic Segmentation. In *CVPR*, 2016.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., and Dollar, P. Focal loss for dense object detection. In *ICCV*, 2017.
- Long, J., Shelhamer, E., and Darrell, T. Fully Convolutional Networks for Semantic Segmentation. In *CVPR*, 2015.
- Louizos, C. and Welling, M. Multiplicative normalizing flows for variational bayesian neural networks. In *ICML*, 2017.

- Ma, C., Xu, Q., Wang, X., Jin, B., Zhang, X., Wang, Y., and Zhang, Y. Boundary-aware Supervoxel-level Iteratively Refined Interactive 3D Image Segmentation with Multi-agent Reinforcement Learning. *IEEE Transactions on Medical Imaging*, 2020.
- MacKay, D. J. A Practical Bayesian Framework for Backpropagation Networks. *Neural Computation*, 4(3):448–472, 1992.
- Mehrtash, A., Ghafoorian, M., Pernelle, G., Ziaei, A., Heslinga, F. G., Tuncali, K., Fedorov, A., Kikinis, R., Tempany, C. M., Wells, W. M., et al. Automatic needle segmentation and localization in mri with 3-d convolutional neural networks: application to mri-targeted prostate biopsy. *IEEE Transactions on Medical Imaging*, 38(4):1026–1036, 2018.
- Mehrtash, A., Wells, W. M., Tempany, C. M., Abolmaesumi, P., and Kapur, T. Confidence calibration and predictive uncertainty estimation for deep medical image segmentation. *IEEE Transactions on Medical Imaging*, 39(12):3868–3878, 2020.
- Menze, B., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., Lanczi, L., Gerstner, E., et al. The Multimodal Brain Tumor Image Segmentation Benchmark (BraTS). *IEEE Transactions on Medical Imaging*, 34(10):1993–2024, 2015.
- Milletari, F., Navab, N., and Ahmadi, S.-A. V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation. In *3DV*, 2016.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *ICML*, 2016.
- Moeskops, P., Veta, M., Lafarge, M. W., Eppenhof, K. A., and Pluim, J. P. Adversarial Training and Dilated Convolutions for Brain MRI Segmentation. In *DLMIA*, 2017.
- Neal, R. M. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- Nie, D., Wang, L., Xiang, L., Zhou, S., Adeli, E., and Shen, D. Difficulty-aware Attention Network with Confidence Learning for Medical Image Segmentation. In *AAAI*, 2019.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019.
- Rajchl, M., Lee, M. C., Oktay, O., Kamnitsas, K., Passerat-Palmbach, J., Bai, W., Damodaram, M., Rutherford, M. A., Hajnal, J. V., Kainz, B., et al. Deepcut: Object Segmentation from Bounding Box Annotations using Convolutional Neural Networks. *IEEE Transactions on Medical Imaging*, 36(2):674–683, 2016.
- Robinson, R., Oktay, O., Bai, W., Valindria, V. V., Sanghvi, M. M., Aung, N., Paiva, J. M., Zemrak, F., Fung, K., Lukaschuk, E., et al. Real-time prediction of segmentation quality. In *MICCAI*, 2018.
- Ronneberger, O., Fischer, P., and Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *MICCAI*, 2015.
- Rother, C., Kolmogorov, V., and Blake, A. "GrabCut": interactive foreground extraction using iterated graph cuts. *ACM Transactions on Graphics*, 23(3):309–314, 2004.

- Shrivastava, A., Gupta, A., and Girshick, R. Training Region-Based Object Detectors with Online Hard Example Mining. In *CVPR*, 2016.
- Simpson, A. L., Antonelli, M., Bakas, S., Bilello, M., Farahani, K., Van Ginneken, B., Kopp-Schneider, A., Landman, B. A., Litjens, G., Menze, B., et al. A Large Annotated Medical Image Dataset for the Development and Evaluation of Segmentation Algorithms. *arXiv preprint*, 2019.
- Song, G., Myeong, H., and Mu Lee, K. SeedNet: Automatic Seed Generation with Deep Reinforcement Learning for Robust Interactive Segmentation. In *CVPR*, 2018.
- 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, 2018a.
- Wang, G., Zuluaga, M. A., Li, W., Pratt, R., Patel, P. A., Aertsen, M., Doel, T., David, A. L., Deprest, J., Ourselin, S., et al. DeepIGeoS: A Deep Interactive Geodesic Framework for Medical Image Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(7):1559–1572, 2018b.
- Xu, N., Price, B., Cohen, S., Yang, J., and Huang, T. S. Deep Interactive Object Selection. In *CVPR*, 2016.
- Yu, L., Wang, S., Li, X., Fu, C.-W., and Heng, P.-A. Uncertainty-aware Self-ensembling Model for Semi-supervised 3D Left Atrium Segmentation. In *MICCAI*, 2019.
- Zhang, S., Liew, J. H., Wei, Y., Wei, S., and Zhao, Y. Interactive object segmentation with inside-outside guidance. In *CVPR*, 2020.
- Zhuang, X. and Shen, J. Multi-scale Patch and Multi-modality Atlases for Whole Heart Segmentation of MRI. *Medical Image Analysis*, 31:77–87, 2016.