2-2022

SeqSeg: A sequential method to achieve nasopharyngeal carcinoma segmentation free from background dominance

Guihua TAO

Haojiang LI

Jiabin HUANG

Chu HAN

Jiazhou CHEN

See next page for additional authors

Medical Image Analysis 78 (2022) 102381

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

SeqSeg: A sequential method to achieve nasopharyngeal carcinoma segmentation free from background dominance

Guihua Tao[a,1], Haojiang Li[c,1], Jiabin Huang[a], Chu Han[b], Jiazhou Chen[a], Guangying Ruan[c], Wenjie Huang[c], Yu Hu[a], Tingting Dan[a], Bin Zhang[a], Shengfeng He[a], Lizhi Liu[c,*], Hongmin Cai[a,*]

[a] School of Computer Science and Engineering, South China University of Technology, Guangzhou 510006, China
[b] Guangdong Provincial People's Hospital, Guangdong Academy of Medical Sciences, Guangzhou, 510080, China
[c] Department of Radiology, State Key Laboratory of Oncology in South China, Guangdong Key Laboratory of Nasopharyngeal Carcinoma Diagnosis and Therapy, Collaborative Innovation Center for Cancer Medicine, Sun Yat-sen University Cancer Center, Guangzhou 510060, China

ARTICLE INFO

Article history:
Received 10 February 2021
Revised 18 January 2022
Accepted 31 January 2022
Available online 11 February 2022

MSC:
41A05
41A10
65D05
65D17

Keywords:
Nasopharyngeal carcinoma
Background dominance
NPC Detection and segmentation
Deep Q-learning
Recurrent attention

ABSTRACT

Reliable nasopharyngeal carcinoma (NPC) segmentation plays an important role in radiotherapy planning. However, recent deep learning methods fail to achieve satisfactory NPC segmentation in magnetic resonance (MR) images, since NPC is infiltrative and typically has a small or even tiny volume with indistinguishable border, making it indiscernible from tightly connected surrounding tissues from immense and complex backgrounds. To address such background dominance problems, this paper proposes a sequential method (SeqSeg) to achieve accurate NPC segmentation. Specifically, the proposed SeqSeg is devoted to solving the problem at two scales: the instance level and feature level. At the instance level, SeqSeg is forced to focus attention on the tumor and its surrounding tissue through the deep Q-learning (DQL)-based NPC detection model by prelocating the tumor and reducing the scale of the segmentation background. Next, at the feature level, SeqSeg uses high-level semantic features in deeper layers to guide feature learning in shallower layers, thus directing the channel-wise and region-wise attention to mine tumor-related features to perform accurate segmentation. The performance of our proposed method is evaluated by extensive experiments on the large NPC dataset containing 1101 patients. The experimental results demonstrated that the proposed SeqSeg not only outperforms several state-of-the-art methods but also achieves better performance in multi-device and multi-center datasets.

1. Introduction

Nasopharyngeal carcinoma (NPC) is a malignant tumor originating in the nasopharynx that is prevalent in South China, northern Africa, and Alaska Wei and Sham (2005). As a routine clinical procedure for NPC diagnosis, magnetic resonance imaging (MRI) is preoperatively used to assess the progress of the tumor. Radiotherapy is the treatment mainstay for NPC. Tumor segmentation in MR images is an essential step in radiotherapy planning King et al. (2000); Chan et al. (2017). A reliable automatic segmentation model can quickly detect and segment the tumor lesion and then effectively reduce the workload of radiologists for radiotherapy planning.

Traditional machine learning-based methods employ handcrafted visual features to segment NPCs Huang and Liu (2012). However, the selection of these features highly relies on experience, and it is incapable of capturing high-level semantic knowledge. In recent years, deep neural networks have achieved wide applications for image understanding tasks due to their powerful automatic feature learning ability. Some efforts using deep neural networks attempt to solve NPC segmentation and achieve significant results. Huang et al. (2018) proposed a deep convolutional neural network for automatic gross tumor volume segmentation on positron emission tomography-computed tomography images (PET-CT). Zhao et al. (2019a) proposed a method that uses fully convolutional networks with auxiliary paths to achieve automatic NPC segmentation on PET-CT images. Chen et al. (2020) proposed a novel multimodality MRI fusion network (MMFNet), consisting of several encoders and one decoder. MMFNet is committed to utilizing multiple modalities of MR images to complete accurate segmentation. However, these methods fail to provide satisfactory segmentation results for NPC in MR images.

* Corresponding author.
E-mail addresses: csghtao@mail.scut.edu.cn (G. Tao), liulizh@sysucc.org.cn (L. Liu), hmcai@scut.edu.cn (H. Cai).
[1] The authors contributed equally.

https://doi.org/10.1016/j.media.2022.102381
1361-8415/

Since 1) the NPC typically occupies a relatively small volume in MR images, the immense and complex background dominates the feature representation during the standard encoding-decoding process. Moreover, 2) the background is cohesive to the tumor with a shape varied and indistinguishable border, making the segmentation task more challenging. Therefore, it is difficult for such one-scale methods to deal with the above various challenges simultaneously.

The above various challenges are termed background dominance problems; solving the background dominance problems is the key to nasopharyngeal carcinoma segmentation. An effective strategy is to decompose the problem and solve it pertinently. Thus, we propose decomposing the background dominance problem into the above two sub background dominance problems. Additionally, we appropriately designed a sequential method (SeqSeg) for NPC segmentation in MR images to address the inherent background dominance problems from different scales at the instance level and feature level. At the instance level, the segmentation attention is forced to focus on the tumor and its surrounding tissue. At the feature level, through a recurrent attention network (RANet) with the recurrent attention mechanism Huang et al. (2019), segmentation is guided to recurrently focus on tumor-related features along channel-wise and region-wise.

Specifically, at the instance level, we aim to concentrate segmentation attention on the tumor and its surrounding tissue by prelocating the tumor with the most compact bounding box. In this paper, we do not use existing supervised learning models due to unstable detection results and the lack of error correction capabilities Zhao et al. (2019a); He et al. (2017). Such drawbacks were even amplified when facing the imbalanced size of the foreground and background. In contrast, reinforcement learning Henderson et al. (2018); Sutton and Barto (2011); Kaelbling et al. (1996) can iteratively correct the error detection results with a reward mechanism. Thus, we design a deep Q-learning (DQL) Mnih et al. (2015) based model for tumor detection at the instance level. In DQL, we integrate the priori knowledge to guide the detection agent. An efficient exploration algorithm is designed to reduce the exploration space and speed up the convergence rate. Additionally, a novel task-specific reward function is proposed to reasonably reward the agent.

At the feature level, a series of attentions are applied through a RANet. It aims to break free from the cohesive background and achieve accurate segmentation. Therefore, recurrent attention modules (RAMs) are introduced to perform channel-wise and region-wise attention on the feature maps by using the high-level semantic features to guide learning low-level features. Thus, the network gives more attention to more discriminating features. To further emphasize the tumor border, we propose a dilated border weighted loss to obtain a segmentation result with a precise tumor border. Due to the lack of public NPC segmentation datasets, we trained and evaluated our proposed model by our collected dataset, which includes 9531 slices of MR images from 1101 patients. Extensive experiments were conducted to evaluate the effectiveness of our proposed SeqSeg method. Our results outperform existing methods and achieve state-of-the-art performance. The contributions and novelties of this paper are summarized as follows:

1) At the instance level, an NPC detection model based on deep Q-learning is proposed. An efficient reward function and an exploration strategy are proposed to accurately detect NPC tumors.
2) At the feature level, RANet with RAMs is designed for NPC segmentation. RAMs leverage the high-level semantic features in deeper layers to guide learning low-level texture features in shallower layers and perform attention learning on tumor-related features along channel and region. A dilated border weighted loss function is used to penalize the segmentation result for precise tumor borders.
3) Focusing on the background dominance problems, a sequential method SeqSeg from different scales, by incorporating instance level and feature level, is proposed to solve the dominance problems to segment the relatively small NPC from an immense background effectively. Extensive experiments demonstrate its effectiveness. The Dice score of our method is 80.32%, which is higher than that of the other comparison methods.
4) We have collected a large NPC dataset thus far. Experiments on multi-device and multi-center datasets demonstrate the reliability of the proposed method. Additionally, we statistically analyze potential confounders that influence segmentation.

2. Related works

2.1. Recent works on NPC segmentation

Traditional nasopharyngeal carcinoma segmentation methods are time-consuming and prone to cause large bias because of the hand-engineered. The accuracy of the algorithms also heavily depends on the rationality of feature extraction Huang and Liu (2012); Zhou et al. (2006); Huang et al. (2015); Li et al. (2019). In contrast, the convolutional neural network (CNN) has automatic powerful feature learning and representation ability Liu et al. (2020). Wang et al. used CNN to segment the NPC in MR images Wang et al. (2018). Lin et al. designed a CNN with an encoder-and-decoder architecture to segment NPC Li et al. (2018). Lin et al. (2019) applied the atlas-based method to cut the image to obtain training datasets. Then, a VoxResNet-based Chen et al. (2018) method is proposed to segment the NPC. Tang et al. (2021) proposed a dual attention mechanism for feature refinement for NPC segmentation. However, the results of these NPC segmentation methods are still unsatisfactory. Because these methods do not handle the background dominance problem caused by the imbalanced region size between the region of interest and uninterest. Additionally, the background is cohesive with the tumor. These methods also do not effectively deal with such a background dominance problem because these methods do not make full use of the low-level features and high-level features. Therefore, we propose a sequential method termed SeqSeg that incorporates instance level and feature level to appropriately deal with the above background dominance problems.

2.2. Object detection via DQL

Most of the object detection models build on Faster R-CNN Ren et al. (2015), which employs region proposal network to generate region proposals with shared convolutions. However, NPC is small, infiltrative, and variable in shape, with an indistinguishable border. Thus, the NPC is poor in feature representation and suffers greatly from background dominance problems in detection tasks (also in segmentation tasks). Moreover, the medical image datasets are generally smaller than the natural image datasets. Thus, the detection algorithm with a supervised learning scheme is prone to overfitting. Therefore, such existing detection methods for NPC detection are still not satisfactory. In contrast, strategy learning algorithms such as deep reinforcement learning can be used to solve background dominance problems by interacting with the environment through actions and reward strategies. It can also avoid model overfitting because agent exploration introduces randomness and creates a large search space. Finally, priori knowledge is worth utilizing, but it is often ignored. Han et al. used deep Q-learning (DQL) to detect objects in videos Han et al. (2018). After the DQL algorithm outputs the bounding box, a fully convolutional DenseNet Jégou et al. (2017) is employed to segment the video

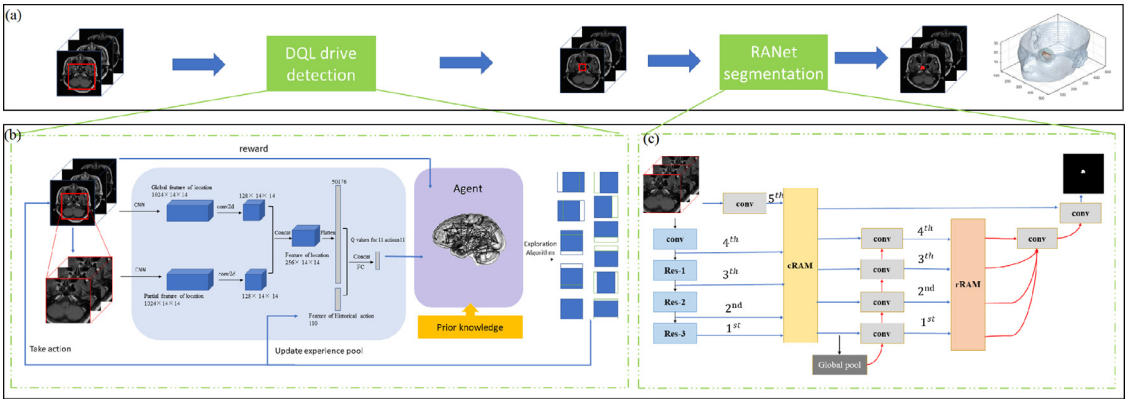

Fig. 1. Schematic illustration of the proposed SeqSeg for NPC segmentation.

frames. Man et al. used DQL to detect the pancreas in CT images. They reward the agent according to the change in interaction over union (IoU) Rezatofighi et al. (2019) between the outputted bounding box and the ground truth Man et al. (2019). However, such an IoU metric is unreasonable in some cases to reward the agent. In this paper, an exploration algorithm integrated prior knowledge is proposed to guide the agent exploration. Additionally, a novel task-specific reward function is proposed to reasonably reward the agent to achieve NPC accurate detection on the MR image.

2.3. Semantic segmentation to extract region of interest

The emergence of fully convolutional networks (FCNs) Shelhamer et al. (2017) is a breakthrough in semantic segmentation. Badrinarayanan et al. proposed SegNet Badrinarayanan et al. (2017) using the encoder-and-decoder architecture. Different from the upsampling by transposed convolution in FCN, SegNet employs the unpooling layer in the decoder for upsampling. Similarly, Ronneberger et al. proposed U-Net Ronneberger et al. (2015) with the encoder-and-decoder architecture. In contrast, after each upsampling in the decoder phase, U-Net clipped the feature maps from the symmetrical position of the encoder to the same size as the upsampling feature maps, and connected it with the upsampling feature maps along the channels. However, in NPC segmentation, most of the features extracted by the networks are rough and weak in expression because of the background dominance problem and the insufficient utilization of features from various levels. The above research did not use a good attention mechanism to emphasize channels and regions with higher relevance to the segmentation target. In this paper, a RANet with RAMs is designed to use the high-level semantic features from the deeper layers to guide learning texture features from the shallower layers, allowing it to electively suppress features with low correlation to segmentation and stimulate features with high correlation to segmentation.

3. Methods

The proposed SeqSeg consists of two major scales to jointly achieve NPC segmentation. Fig. 1 demonstrates the proposed SeqSeg and its architecture. As shown in Fig. 1(a), three modality MR sequences T1-weighted (T1), enhanced contrast T1-weighted (T1C), and T2-weighted (T2) Zhuo et al. (2019) are concatenated and fed into the detection model at the instance level. Upon having the detected bounded area at the instance level, it is then analyzed at the feature level to obtain a final segmentation. The details of the method are illustrated in Fig. 1(b) and 1(c). Section 3.1 illustrates the instance level processing of the proposed SeqSeg, that is, NPC detection. Section 3.2 shows the details of the segmentation at the feature level.

3.1. Instance level: DQL-based NPC detection

At the instance level, DQL-based NPC detection model is proposed. NPC detection is modeled as a decision-making process of the agent aimed at obtaining a bounding box with the highest similarity to the ground truth to reduce the uninterested region for the next scale. At time t, the agent takes action a_t to interact with environment state s_t. The environment transitions to the new state, which is $s_{(t+1)}$, and gives reward r_t for the actions of the agent according to the reward function. During the training phase, the agents take action according to the exploration algorithm. While testing, the agent takes action with the maximum Q value according to the trained Q-network. When the agent takes a terminal action, the detection model outputs the final detection result, i.e., the final bounding box. After each interaction with the environment, a state transition was stored in the experience pool. Fig. 2 illustrates the DQL based NPC detection model.

3.1.1. Agent actions

In this paper, eleven actions, simulated in Fig. 3, are designed for the agent to make decisions. There are six zoom actions designed to decrease the size of the bounding box to be 3/4, four shift actions move the bounding box by 1/4 closer to the ground truth, and a terminal action ends the decision-making process. The agent takes the terminal action once the current action is selected or the maximum search step is reached.

3.1.2. State

The state of the environment contains enough information for the agent to make wise decisions. The location of the bounding box in the MR image and the historical actions are combined to encode the environmental state. When encoding the location, the image features on the whole MR image are extracted as the global features, while the image features on the corresponding region of the bounding box are extracted as the local feature. Since the designed action space is discrete and the number of actions is 11, this paper adopts one-hot encoding to encode the historical actions and obtain the historical action features with dimensions of $h \times 11$, where h denotes the number of historical actions.

To address this problem, we design a new reward function for each type of action. (1) After the agent takes a shift action, if the updated bounding box is closer to the ground truth, the agent then receives a positive reward; otherwise, it receives a negative reward. (2) When the agent takes the zoom action, if the IoU before and after the action is taken are both zero, and the bounding box is still

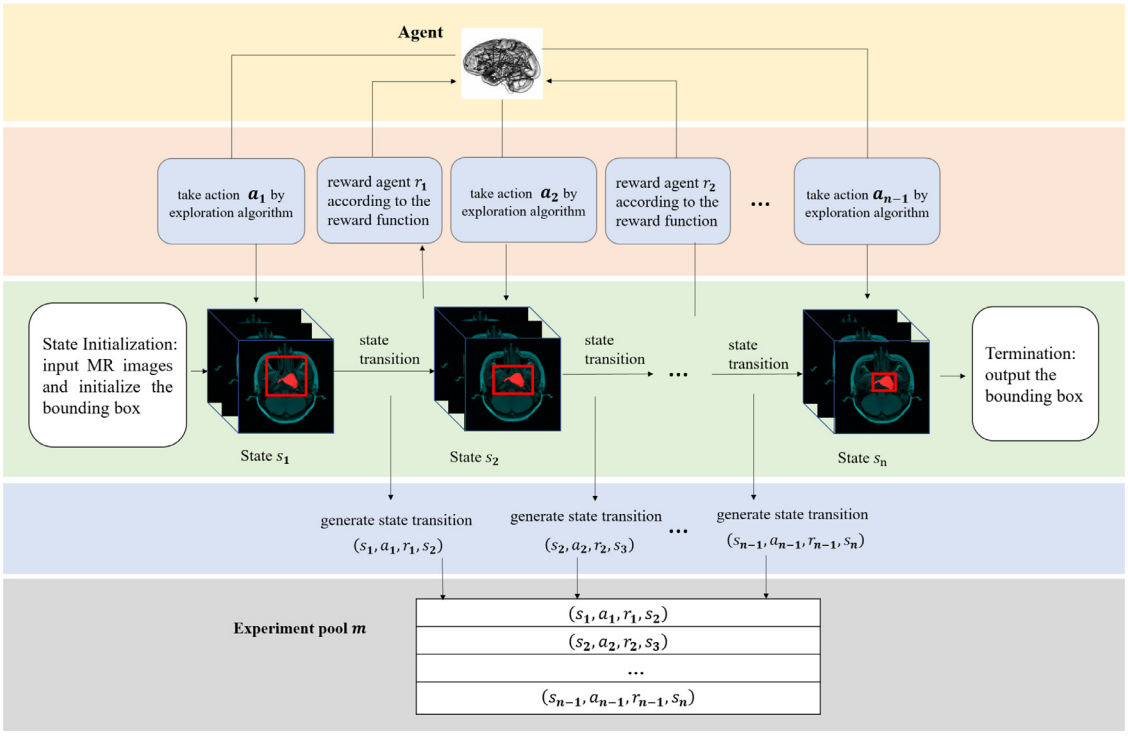

Fig. 2. Illustration of DQL based NPC detection model.

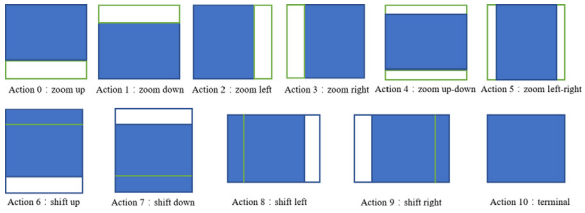

Fig. 3. Illustration of agent actions. The blue region denotes the original bounding box, while the white region denotes the bounding box after the action taken by the agent.

greater than the size of the ground truth (multiplied by a predefined threshold), then the reward is given according to the change between the center of the bounding box before and after the action is taken; otherwise, rewards are given based on the changes in IoU. (3) When the action is the terminal action, a positive reward is given if recall and IoU are both greater than the corresponding predefined threshold. The reward function is illustrated in Algorithm 1, where $\tau_h, \tau_w, \tau_{Recall}$ and τ_{IoU} are the predefined thresholds. b and b' denote the bounding box before and after the action is taken. g denotes the ground truth. The function $\rho(\cdot)$ is the distance between two center points.

3.1.3. Reward function

Most of the existing work on object detection using deep reinforcement learning rewards the agent by the intersection over union (IoU) between the bounding box and the ground truth Wen et al. (2021); Han et al. (2018); Man et al. (2019); Zhao et al. (2019b). However, there are several drawbacks in evaluating the actions of the agent by changing the IoU. When the bounding boxes have no intersections with the ground truth before and after the action is taken, the reward is not positive; even an updated bounding box is closer to the ground truth. Such conflict causes the agent to take exhaustive actions to search through the space, thereby introducing high costs. As shown in Fig. 4, the bounding box becomes closer to the object (red area) by moving the green box to the blue box after the action is taken. However, the IoU between the bounding box and the ground truth remains 0. Therefore, the reward function in Han et al. (2018); Man et al. (2019) prohibits the agent from obtaining a positive reward, resulting in a conflict between the reward and reasonable action.

Algorithm 1 Reward function.

if *action_type* ==′ *shift*′ **then**
  r := sign $\left(\rho\left(b, b^{gt}\right) - \rho\left(b', b^{gt}\right)\right)$
else if *action_type* ==′ *hrink*′ **then**
  if IoU $\left(b', g\right) == 0$ and IoU$(b, g) == 0$ and $b'_h > b^{gt}_h \times \tau_h$ and $b'_w > b^{gt}_w \times \tau_w$ **then**
    r := sign $\left(\rho\left(b, b^{gt}\right) - \rho\left(b', b^{gt}\right)\right)$
  else
    r := sign$(\text{lo}U(b', g) - \text{lo}U(b, g))$
  end if
else
  $r := \begin{cases} +\sigma, & \text{Recall}(b, g) > \tau_{\text{Recall}}, \text{IoU}\left(b, b^{gt}\right) > \tau_{IoU} \\ -\sigma, & \text{Otherwise} \end{cases}$
end if
return r

3.1.4. Exploration algorithm

Most of the work using deep Q-learning adopts the ϵ-greedy algorithm to select actions to explore the environment Man et al. (2019); Liu and Hodgins (2017). The agent needs to make trial-and-error in the large exploration space, resulting in the difficulty of converging the Q-network in the early training phase. In this paper, we utilize priori knowledge according to the location and size of the NPC to guide the agent to reasonably reduce the search space, thereby speeding up the convergence rate of the model and improving the detection performance. The pro-

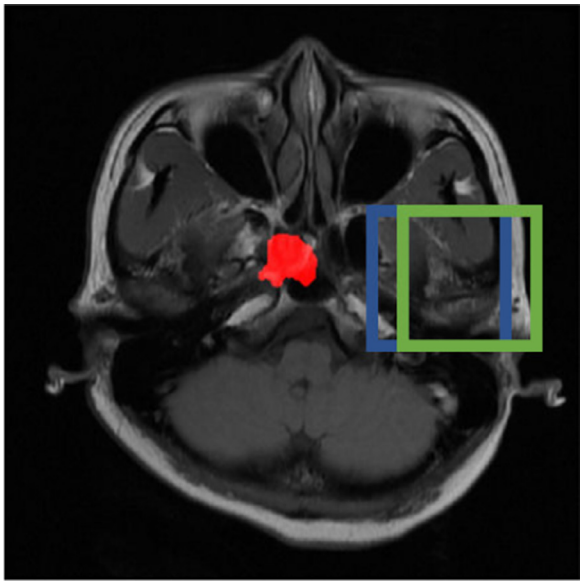

Fig. 4. Illustration of a problem in evaluating the actions of the agent by the change of IoU.

posed scheme explores the relative location and size difference between the bounding box and ground truth to deduce conducive actions. The differences are quantified by a priori knowledge term ϵ_{priori}. (1) If the center distance between the bounding box and the ground truth is greater than a predefined threshold, the shift action is taken. (2) If the center distance is less than a predefined threshold but there is still a gap between the twos in size, the zoom action is taken. (3) In cases other than (1) and (2), the terminal action can be taken. By taking actions based on such priori knowledge, the agent has a higher probability of receiving a positive reward, thus making the Q-network learning samples more balanced. The algorithm for taking the three actions is shown in Algorithm 2. In the early training phase, to reduce the agent search

Algorithm 2 Taking action base on priori knowledge.

if $\rho\left(b, b^{gt}\right) > \tau_{distance}$ **then**
  take shift actions
else if $b'_h > b^{gt}_h \times \tau_h$ or $b'_w > b^{gt}_w \times \tau_w$ **then**
  take zoom actions
else
  take the terminal action
end if

space, a higher probability is used to take actions based on priori knowledge. However, in the later phase of training, the probability of randomly selecting actions ϵ is low, and the Q-network is also more reasonable, so the probability of taking actions based on priori knowledge declines. In summary, the early training phase is driven by priori knowledge, and the process is controlled by a high probability value of ϵ_{priori}. In contrast, as the training progresses, the probability ϵ_{priori} declines linearly until it drops to 0. The proposed exploration algorithm based on priori knowledge is shown in Algorithm 3. ϵ_{priori_decay} is the decay factor of ϵ_{priori}, while ϵ_{decay} is the decay factor of ϵ.

3.1.5. Deep Q-learning

The DQL algorithm utilizes a CNN Q-network to fit the Q-function. To stabilize the training, the DQL algorithm uses a target Q-network with the same structure to calculate the target Q value. The weight of the target Q-network is updated at intervals.

Algorithm 3 Exploration algorithm based on priori knowledge.

for $epoch \in [1, T]$ **do**
  if random()$< \epsilon_{\text{priori}}$ **then**
    taking action using Algorithm 2
  else if random() $< \epsilon$ **then**
    taking action randomly
  else
    taking the action with the highest Q value
  end if
  $\epsilon_{priori} := \epsilon_{priori} - \epsilon_{priori\_decay}, \epsilon := \epsilon - \epsilon_{decay}$
end for

The Q value of the state s_t action a_t is calculated using the Q-network, $q = Q(s_t, a_t; \theta_Q)$, where θ_Q is the Q-network parameter. The Q-value of the target to be fitted is obtained according to the immediate reward of the agent from the environment and the target Q-network. The calculation formula is given by

$$q' = \begin{cases} r_t & a = \text{terminal} \\ r_t + \gamma \max_{a'} Q\left(s_{t+1}, a'; \theta_{\text{target}}\right), & a \neq \text{terminal} \end{cases} \quad (1)$$

where θ_{target} is the target Q-network parameter.

The proposed Q-network is illustrated in the light-blue area in Fig. 1(b). After the location feature and historical action feature are extracted, we use a fully connected layer to output the Q-values of each action.

3.2. Feature level: NPC segmentation through recurrent attention

At this level, we take the output of the instance level as the input of the segmentation model. Then, we design a RANet with RAMs modules to obtain the final segmentation result.

3.2.1. Recurrent attention network (RANet)

RANet uses the encoder-decoder architecture, as shown in Fig. 1(c).

The black and red lines represent the downsampling and upsampling operations, respectively, and the blue lines represent the data flow without changing the feature map resolution. The network takes T1, T1C, and T2 MR images as input and outputs a corresponding probability map as the segmentation result. In the encoder phase, it utilizes ResNet He et al. (2016) trained in advance on ImageNet Krizhevsky et al. (2012) to extract multi-level semantic features. The extracted feature maps are iterated into the channel-wise recuttent attention module (cRAM) in order from the deepest layer to the shallowest layer, and the attention feature map is generated. Then, the global average pool is applied to the deepest attention feature map in the deepest layer to obtain the global contextual feature map. In the decoder phase, the upsampling results are combined with the attention feature map from the cRAM, which is input into the rRAM in order from the shallowest layer to the deepest layer. Finally, the attention feature maps of rRAM are convoluted together with the global contextual feature map of cRAM to obtain the final segmentation result. Illustrations and details of cRAM and rRAM are shown in Suppl. 1.

3.2.2. Channel-wise recurrent attention (cRAM)

Similar to squeeze-and-excited (SE) networks Hu et al. (2018), cRAM is dedicated to stimulating channels with discriminate information, which is beneficial to NPC segmentation. The low-level features in the shallower layers are recalibrated by feedback from long short-term memory (LSTM) Hochreiter and Schmidhuber (1997), which memorizes the high-level contextual information from deeper layers. Given n level feature maps from n layers, there are n level feature maps $X = [X_1, X_2, \ldots, X_n]$, $X_i \in \mathbb{R}^{C_i \times H_i \times W_i}$, where

C, H and W denote the channel, height and width, respectively. Each feature map passes through a global average pooling layer and a convolution layer and then is reshaped to a vector v_i with designative length. At the $(i-1)^{th}$ iteration, LSTM takes v_i as input, takes h_{i-1} as the hidden state and outputs $v_i' = F_{LSTM}(v_i; h_{i-1})$. Then, we obtain attention weight W_i by a convolution operation on v_i'. W_i is then activated by the sigmoid activation function to output the recurrent channel-wise attention map $A_i \in R^{C\times 1\times 1}$. Finally, we obtain the attentional feature map U_i by

$$U_i = A_i \odot X_i, i \in \{1, 2, \ldots, n\} \tag{2}$$

where $\odot$ is the element-wise multiplication operator using array broadcasting.

3.2.3. Region-wise recurrent attention (rRAM) to strengthen features of interest

The rRAM is devoted to enhancing feature learning by focusing on the region that is highly related to targeted segmentation. It exploits the feedback from the LSTM semantic guidance to refine the feature map and learn the attention map across regions rather than channels. At the $(i-1)^{th}$ iteration, the rRAM takes feature maps $X = [X_1, X_2, \ldots, X_n]$ as input, and h_{i-1} as the hidden state, and outputs the corresponding attention maps $A^r = [A_1^r, A_2^r, \ldots, A_n^r]$. Specifically, each of the attention maps can be learned by

$$A_i^r = \sigma\left(f(X_i, h_{i-1})\right), i \in \{1, 2, \ldots, n\} \tag{3}$$

where f represents a series of operations, including adaptive pooling, convolution and LSTM. First, the feature maps X_i are fed into the adaptive pooling layer and transformed to regional statistics Z_i with size C×S×S, where S is a predefined size for Z_i and $S < min(H, W)$. Convolution is then performed on Z_i and outputs a vector v_i. Then, LSTM takes v_i as input and outputs an attentional vector v_i'. Furthermore, the attentional vector is transformed into an attention map A_i^r by successive operations in the order of reshaping, convolution, upsampling, and sigmoid function activation. Finally, element-wise multiplication is performed on A_i^r and X_i to generate the attentional feature map, which stimulates the more relevant regions to enhance segmentation performance.

3.2.4. Dilated border weighted loss function (DBW)

NPC segmentation with a precise tumor border is essential for subsequent tumor evaluation. In this paper, we applied a loss function to force RANet at the feature level to give more attention to the tumor border and its nearby pixels. We first applied the Canny edge detector on the ground truth to obtain the tumor border binary map. Then, the binary map is amplified by using a morphological dilation operation on the border. To retain the morphological structure of the tumor border, a cruciform structural element with a size equal to (2k+1)×(2k+1) is applied on the binary map. Let $w_{i,j}$ represent the weight on the pixel (i, j). If the pixel is at the dilated border, $w_{i,j} = w_b$; otherwise, $w_{i,j} = 1$. The loss values on each pixel are normalized by dividing the summation of all pixel weights. Finally, the dilated border weighted loss function between the ground truth $y_{i,j}$ and the prediction $p_{i,j}$ is formulated by

$$\begin{aligned} L = & \frac{1}{\sum_{i=0}^{H}\sum_{j=0}^{W} w_{i,j}} \quad \sum_{i=0}^{H}\sum_{j=0}^{W} L_{i,j} \\ = & \frac{1}{\sum_{i=0}^{H}\sum_{j=0}^{W} w_{i,j}} \quad \sum_{i=0}^{H}\sum_{j=0}^{W} -w_{i,j} \\ & \left[y_{i,j}\log\left(p_{i,j}\right) + \left(1-y_{i,j}\right)\log\left(1-p_{i,j}\right)\right] \end{aligned} \tag{4}$$

where H and W are the height and width of the input image, respectively.

The proposed dilated border weighted loss function guides the neural network to mine the tumor edge during training to yield precise NPC segmentation.

4. Experiments

Our proposed SeqSeg includes NPC detection model at the instance level and NPC segmentation model at the feature level. Thus, we first conduct verification experiments to evaluate the detection performance at the instance level. Next, experiments are carried out to evaluate the effectiveness of RAMs modules and dilated border weighted loss function applied at the feature level. Then, we evaluate the segmentation performance of the proposed SeqSeg method on a large NPC dataset. Additionally, we evaluate the reliability of the SeqSeg on multi-device and multi-center datasets. Finally, we statistically analyze potential confounders that influence segmentation. We are devoted to increasing the segmentation performance by alleviating the corruption caused by background dominance. Thus, we do not consider the slice with no NPC. For this reason, only the slices that include NPC are used for experiments. The proposed model is a semiautomatic method.

4.1. Dataset

We performed experiments on MR images of 596 NPC patients collected at Sun Yat-sen University Cancer Center from January 2010 to January 2013. The MR images included T1-weighted (T1), T2-weighted (T2), and contrast-enhanced T1-weighted (T1c) images.

The imaging parameters were as follows: axial T1-w imaging (FSE, TR = 540 ms, and TE = 11.8 ms), axial T2-w imaging (FSE, TR = 4,000 ms, and TE = 99 ms), and axial CET1-w imaging (FSE, TR = 540 ms, and TE = 11.8 ms). The slice thickness between adjacent slices was 5 mm on average, and the spacing between slices was approximately 6 mm, in which imaging had a resolution of 0.74 mm × 0.74 mm on average. The tumors occupied an average area of 1.11% in a slice. We divided the training, validation, and test sets into 301, 129, and 166 patients, respectively. Different patients have different numbers of MR images. There were 185 patients with 36 slices, 397 patients with 32 slices, and 14 patients with 16 slices. We selected the slices that had tumors from each patient. The collected slices come from both early and advanced patients. Then, there were 5412 images, including three modalities, each with 1804 images. The corresponding ground truth was delineated and was defined by the consensus of four experienced radiologists. Data augmentation was applied, including rotation and horizontal flipping. Moreover, for multicenter evaluation, we collected additional datasets of 300 patients (2,409 images) and 205 patients (1,710 images) from the Cancer Center of Sun Yat-sen University and the First People's Hospital of Foshan, respectively.

We evaluated the NPC segmentation performance by some standard evaluation metrics, including the Dice similarity coefficient (Dice), coefficient (CC), precision, recall, and Hausdorff Distance (HD). In addition, intersection over union (IoU) is also used as an evaluation metric for object detection.

4.2. Implementation details

4.2.1. Image preprocessing

In this paper, the MR images of T1, T1c, and T2 are concatenated together as the input with three channels. For each channel, by calculating the maximum pixel value of the channel in the whole dataset, the pixels in the MR image are rescaled to [0, 1]. All images were resized to 384×384.

4.2.2. Deep Q-learning settings

We used ResNet-50 as the CNN backbone to extract image features of the environment state, and the image was rescaled to 224×224 before being fed to ResNet-50. The Q-network used the L_1 loss function, while the regulation term with L_2 was adopted

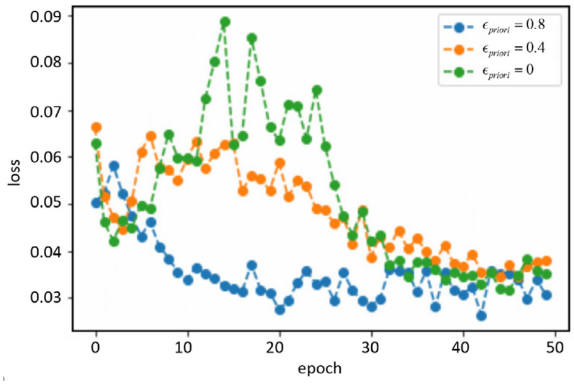

Fig. 5. Convergence curve of our proposed DQL based NPC detection model with different initial ϵ_{priori}.

with a coefficient of 0.0004. Stochastic gradient descent (SGD) with momentum was set as 0.9. The learning rate was 0.0005, and the batch size was 128. The maximum number of steps for the agent was 15. During the training, the hyperparameters in the ϵ-greedy algorithm decayed from 1.0 to 0.1 with a decrease rate $\epsilon_{decay} = 0.1$ in the first 10 epochs. The discount factor γ in learning was set to 0.6, and the size of the experience pool was 16,384. The target Q-network updated its parameter from the Q-network every 1000 steps. A maximum of 50 epochs was trained in each experiment. In addition, the decay factor ϵ_{priori_decay} in the proposed exploration algorithm was 0.2. $\tau_h, \tau_w, \tau_{Recall}$ and τ_{IoU} were set to 0.9, 0.9, 0.95, and 0.5, respectively. $\tau_{distance}$ was set to 1/8 of the diagonal length of the target bounding box. The bounding box was initially set at the center of the image, and its height and width were 42% of the image size.

4.2.3. Sequential NPC segmentation

During the training process of the detection model, a bounding box was obtained once every action was taken. The MR image was cropped according to the bounding box output by the agent. To guarantee enough contextual information, we enlarge the bounding box by 1.5 times the input image of RANet. In the joint training of the two scales, the initial learning rate was 0.001 and declined to 0.0001 after 25 epochs.

4.3. Detection results

4.3.1. Convergence

Using the different initial probabilities of taking actions based on priori knowledge, we analyzed whether the priori knowledge used in the exploration algorithm would have an impact on the model convergence. There were three different initial values ϵ_{priori}, which were 0, 0.4, and 0.8. ϵ_{priori} decayed with $\epsilon_{priori\_decay} = 0.2$ at the end of each epoch until it decayed to 0. Fig. 5 shows that with the value of ϵ_{priori} being zero, that is, no priori knowledge was used, there were several rush jumps in the loss value. The training converged slowly to a large error of approximately 0.04. Similar behavior was also observed with the value of ϵ_{priori} being 0.4. In comparison, when we used a large value of ϵ_{priori} of 0.8, the training rapidly converged to a small error of approximately 0.03. The resulting loss value smoothly decayed until convergence. The comparison demonstrates that the convergence of the Q-network can be accelerated by using priori knowledge.

4.3.2. Ablation study of the exploration algorithm

This subsection analyzes the behaviors of the proposed exploration algorithm by comparing different initial probabilities ϵ_{priori} in exploiting priori knowledge to guide action taking. The results are summarized in Tab. 1. Without priori knowledge, i.e., $\epsilon_{priori} = 0$, the proposed exploration algorithm is reduced to the ϵ-greedy algorithm. The larger the value ϵ_{priori} is, the more obvious the effect of priori knowledge in the initial training phase. Tab. 1 shows the experimental results of the exploration algorithm. When the agent took actions according to priori knowledge with a high probability $\epsilon_{priori} = 0.8$, the detection model achieved the highest score on all metrics except for the max IoU. Therefore, it can be concluded that in the early training phase, the agent selected actions based on priori knowledge with high probability, which reasonably reduced the exploration space of the agent, thus speeding up the convergence speed and improving the detection performance.

Table 1
Experimental results for exploration algorithm.

ϵ_{priori}	Recall			IoU		
	min	max	mean	min	max	mean
0.8	49.00%	100.00%	95.13%	7.53%	81.73%	49.41%
0.4	34.66%	100.00%	87.47%	6.19%	79.60%	48.95%
ϵ-greedy	29.91%	100.00%	86.74%	5.83%	82.90%	43.05%

Table 2
Experimental results for comparison with Faster R-CNN.

Model	Recall			IoU		
	min	max	mean	min	max	mean
Ours	49.00%	100%	95.13%	7.53%	81.73%	49.41%
Faster R-CNN	0	100%	93.54%	0%	98.57%	69.66%

4.3.3. Comparison with Faster R-CNN

We compared the proposed detection model with the benchmark method of Faster R-CNN Ren et al. (2015). For a fair comparison, we used feature pyramid network (FPN) Lin et al. (2017) as the RPN for Faster R-CNN to exploit its merits in detecting small targets. Tab. 2 summarizes the experimental results, where the Faster R-CNN yields a higher mean IoU and maximum IoU than the proposed detection model. However, the proposed detection model achieved a higher mean recall value of 95.13% than the Faster R-CNN of 93.54%. Such a high recall value implies that the proposed detection model detects stable bounding boxes containing suspicious tumors more frequently than Faster R-CNN. Moreover, the minimum recall by Faster R-CNN was 0, while the rate by the proposed model achieved 49.00%. Such superior results demonstrate that the proposed model performs stably in that the agency takes the correct action in moving the bounding box.

An illustrative example is shown in Suppl. 2 to demonstrate the detection process by the proposed DQL. Initially, we start with an initial bounding box in a solid white line. Then, during the detection process, the agent moves the bounding box with various actions to cover the tumor area with a minimal area.

4.4. Segmentation results

In this section, we aim to validate the performance of tumor segmentation at the feature level. For a fair comparison, the whole MR slice other than the partial image detected bounding box was tested by various models.

4.4.1. Test on recurrent attention modules

The RAMs were incorporated into a one-scale U-Net to measure its capability. The experiments were summarized in Tab. 3.

One can observe that the proposed RAMs remarkedly improved the performance of U-Net by increasing the Dice scores by 5.6%. Similarly, U-Net equipped with cRAM is comprehensively superior to U-Net equipped with an SE block in every metric. For visual

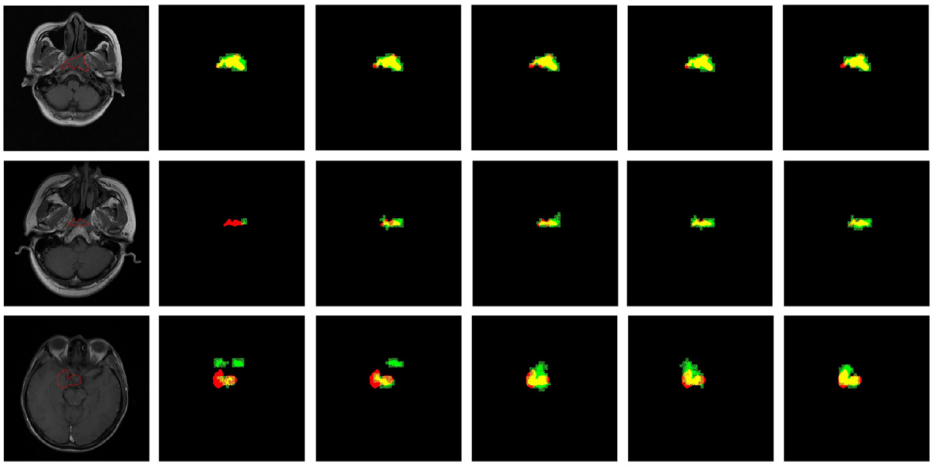

Fig. 6. Results comparison for NPC segmentation in 3 slices. The red region denotes the ground truth, the green region denotes the segmentation results of the models, and the yellow region denotes the interaction between them.

Table 3
Experiments to demonstrate the capability of the proposed recurrent attention modules.

Model	Dice	Precision	Recall	CC
U-Net	66.71%	70.89%	62.99%	0.18%
U-Net+SE	69.76%	72.27%	67.41%	13.29%
U-Net+cRAM	71.87%	**73.20%**	70.58%	21.71%
U-Net+rRAM	71.10%	70.27%	71.94%	18.70%
U-Net+cRAM+rRAM	**72.31%**	70.27%	**75.38%**	**25.02%**

Table 4
Experiments to compare the performance of RANet by different loss functions.

Loss function	Dice	Precision	Recall	CC	HD
Cross entropy	66.71%	**70.89%**	62.99%	0.18%	19.01
Category weighted	61.47%	47.25%	**94.31%**	-25.36%	22.18
Dice loss	66.96%	66.89%	67.05%	1.34%	18.98
Dilated border weighted	**71.65%**	69.17%	77.99%	**20.87%**	**16.75**

Table 5
Experiments to compare RANet (weighted loss) with existing methods.

Model	Dice	Precision	Recall	CC	HD
U-Net	66.71%	70.89%	62.90%	0.18%	19.01
DenseNet	60.63%	53.32%	70.25%	-29.90%	22.74
DDNN	70.98%	69.85%	72.14%	18.22%	18.08
Att Unet	73.43%	70.39%	75.61%	25.68%	16.85
RANet	77.66%	75.54%	79.91%	42.48%	14.72
RANet (weighted loss)	**79.63%**	**78.24%**	**81.07%**	**48.84%**	**13.11**

comparison, the segmentation results on three typical MRI slices after different models are shown in Fig. 6. The tumor labeled by the radiologists is highlighted in red, while the predicted results are in green. Their overlapped parts are in yellow for emphasis. In the first example (first row in Fig. 6), the tumor has a nice shape with a large area and is easy to segment. In the second example, the tumor is tiny; therefore, the segmentation is largely affected by the background dominance problem. The third example is more challenging than the first two; the tumor is infiltrative, and thus, its appearance is irregular. The five models were found to accurately segment the tumor in the first example. Thus, the obtained results are largely covered in yellow. In the second example, U-Net mistakenly segmented the tumor. The error was partially corrected by using the SE block or cRAM. However, there remained mistakes. By using both cRAM and rRAM, U-Net perfectly segmented the suspicious tumor. In the third example, both U-Net with/without SE block mistakenly identified more tumor areas than the labeled ones. In comparison, the segmentation result obtained by jointly using cRAM and rRAM was very similar to the labeled results with irregular appearances.

The experiments demonstrate that the proposed RAMs improve the segmentation performance of U-Net by exploiting the semantics in the deeper layers to guide the feature learning in the shallower layer with the help of both channel and regional information.

4.4.2. Test on dilated border weighted loss function

We compared the dilated border weighted loss function with the typical cross-entropy loss function, the category weighted loss function, and the Dice loss function. The mentioned compared loss functions can be found in Suppl. 3.

For a fair comparison, all of these experiments use a one-scale U-Net as the baseline model. Tab. 4 summarized the results. When using the popular cross-entropy loss function, U-Net achieves the highest precision but a low recall value. When using the category weighted cross-entropy loss function, the recall grows with the increase in the positive category weighting. However, the precision decreases sharply, implying that the segmentation result has a large number of false positives. The Dice loss function achieved a balance in all metrics and outperformed the first two loss functions in key metrics Dice scores, but its performance was still unsatisfactory. In contrast, when using the proposed dilated border weighted loss function, both the Dice score and CC increased steadily with increasing dilation coefficient k and edge pixel weight w_b. The two values reached their peaks when $k = 10$, $w_b = 20$.

4.4.3. Comparative test of segmentation capability by RANet

We compared the RANet performance with that of the designed RAMs under DBW loss with three state-of-the-art methods on NPC segmentation. The comparative methods included U-Net Ronneberger et al. (2015), fully convolutional DenseNet Jégou et al. (2017), deep deconvolutional neural network (DDNN)Men et al. (2017) and the attention U-Net (Att Unet)Oktay et al. (2018). Tab. 5 summarizes the experimental results. It is clear that our method achieves better performance than other methods on all the metrics, where the Dice and CC values of our proposed method are 79.63% and 48.84%, respectively, which are dramatically larger than those of U-Net (66.71%, 0.18%), DenseNet (70.98%, -29.90%), DDNN (70.98%, 18.22%) and Att Unet

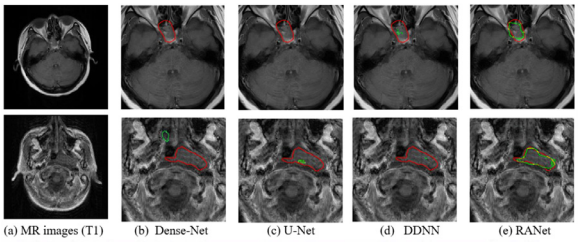

Fig. 7. Comparative results on NPC segmentation after different models. The red contour, green contour and yellow contour represent the ground truth, prediction result and their overlap, respectively. There are two samples of the T1-weighted MR image shown in the first column. The second to fourth columns show the results of two samples obtained by different methods. For the best visualization, we zoomed the segmentation results.

Table 6
Quantitative measurements on five segmentation methods.

Model	Dice	Precision	Recall	CC	HD
RANet-S	79.63%	**78.24%**	81.07%	48.84%	13.11
Mask R-CNN	76.91%	73.84%	83.00%	39.96%	14.34
Lin et al. (2019)	78.39%	73.00%	86.75%	44.87%	13.78
Tang et al. (2021)	77.98%	73.45%	85.00%	43.52%	13.16
SeqSeg (Ours)	**80.32%**	75.53%	**87.57%**	**51.00%**	**12.25**

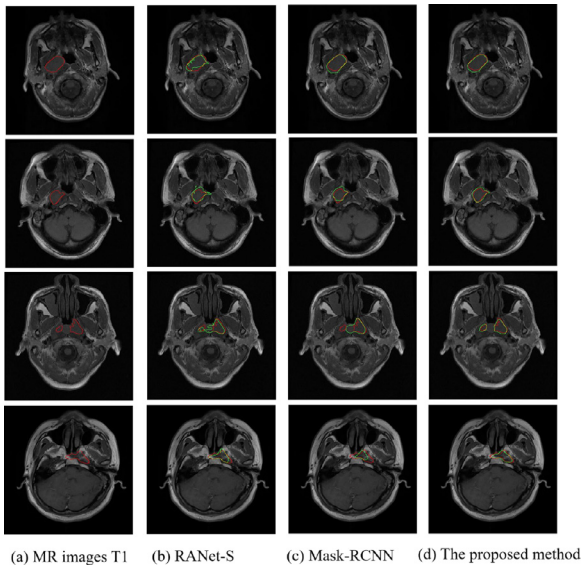

Fig. 8. Comparison of the segmented NPC by different models. The red, green, and yellow contours represent the ground truth, the prediction, and their overlap, respectively.

(72.91%, 25.68%). The proposed method obtained a nearly 10% increasing ratio with respect to the Dice value.

The segmentation results on the two difficult cases after different methods are visualized in Fig. 7. In the first example, the tumor border was unclear, and the volume was small; in the second image, the tumor contrast was weak. RANet also achieves the best visualization results, and its predicted contours fit well with the ground truth. In comparison, the other methods missegment the tumor, and thus the obtained segmentation fails to seriously overlap with the labeled truth.

5. Results of the proposed SeqSeg method

We tested the proposed SeqSeg method on real datasets by comparing it with Mask R-CNN He et al. (2017), Tang et al. (2021) DA-DSUnet, and a network proposed by Lin et al. (2019). We selected the benchmark Mask R-CNN because it is a representative method that inherits from Faster R-CNN and has been successfully tested on various instance detection tasks. Additionally, it is a sequential framework similar to our method for detection and segmentation. To demonstrate the benefit of detection at the instance level to overcome the background dominance problem, RANet is also tested individually, in which the whole MR image rather than the detected region is used. Lin et al. (2019) is an effective NPC segmentation method based on VoxResNet Chen et al. (2018). It first crops patches covering the tumor using an atlas-based method Sims et al. (2009). Then, a VoxResNet-based method is designed to refine the delineation of the tumor. Similar to our method, DA-DSUnet designed an attention mechanism for feature refinement. DA-DSUnet is a dual attention method, including position and channel attention, for NPC segmentation.

The experimental results are summarized in Tab. 6. One will note that the proposed SeqSeg achieves the highest performance with respect to all metrics except precision. Dice and recall are pivotal metrics in the NPC segmentation task. The proposed SeqSeg achieves a higher Dice value of 80.32% than RANet (79.63%) and a higher recall value of 87.57% than RANet (81.07%). Additionally, the proposed SeqSeg achieves the highest score in all metrics compared with Mask R-CNN, Lin Network Lin et al. (2019), and DA-DSUnet. For visual comparison, the segmentation results of each method are shown in Fig. 8. Red contours, green contours, and yellow contours represent the ground truth, segmentation result, and their overlapping part, respectively. In the first and second examples, the tumor has a nice shape with a distinguishable border and is easy to segment. The third example and the fourth example are more challenging than the first two. In the third example, the tumor is surrounded by low contrast tissues. In the fourth example, the tumor is flocculent and unshapely. Additionally, both of them have an irregular shape and an indistinguishable border. This causes the background to largely affect the segmentation by the background dominance problem. The three models were found to accurately segment the tumor in the first example and the second example. Thus, the obtained contours fit well with the ground truth. In the third and fourth examples, the Mask R-CNN results were dissimilar to the original shape. The results of RANet missed the border details. In comparison, the segmentation results by the proposed SeqSeg fit well with the labeled contours without the above problems. The experiments demonstrate that the proposed SeqSeg achieves the best visualization results compared with other methods.

We analyze the influence of different factors on segmentation accuracy. Categorical variables for the testing cohorts were compared by using the x^2 test or Fisher exact test, and numeric variables were compared by using the Mann Whitney U test. The Mann Whitney U test was used to compare Dice and HD between different subgroups. Statistical significance was set at two-tailed P-value$<$0.01. In the subgroup analyses, the proposed method achieved comparable Dice and HD between the post-induction chemotherapy and chemotherapy-naive subgroups (Dice: P-value=0.29; HD: P-value$<$0.001). In HD, the performance of the proposed method was highly related to the chemotherapy condition. Its performance in chemotherapy-naive patients was better than that in post-induction patients. Tab. 7 shows the details of the performance under the different chemotherapy conditions of the proposed method. For the different T-categories, the proposed method achieved a significantly smaller HD in the early T-category tumors than in the advanced T-category tumors (Dice: P-value=0.97; HD: P-value$<$0.001). Tab. 8 shows the details of the performance in the different T-categories of the proposed method. However, we did not observe a difference in Dice between the dif-

Table 7
The performance of our method in different chemotherapy condition (Chemo). Dice: *P*-Value=0.29; HD: *P*-Value<0.001.

Chemo	Dice	Precision	Recall	CC	HD	Sample
0	80.65%	76.09%	87.50%	52.01%	**11.42**	235
1	79.73%	74.77%	87.57%	49.15%	**13.38**	169

Note. 0 = Chemotherapy Naive, 1 = Post induction chemotherapy.

Table 8
The performance of our method in different T-category. Dice: *P*-Value=0.97; HD: *P*-Value<0.001.

T-category	Dice	Precision	Recall	CC	HD	Sample
T1+T2	79.84%	75.01%	86.63%	49.50%	**9.459**	51
T3+T4	80.33%	75.61%	87.66%	51.03%	**12.65**	353

Note. T staging used the American Joint Committee on Cancer staging system.

ferent chemotherapy conditions and different T-categories. In addition, other confounders, including gender, age, MR sequences, image quality, and N-category showed no impact on segmentation performance. The corresponding performances are shown in Suppl. Tab. 1, Tab. 2, Tab. 3, Tab. 4, and Tab. 5, respectively. Importantly, experimental results show that our method maintains stable and accurate segmentation performance on the newly added datasets from multi-devices and multi-centers. The details are shown in Suppl. Tab. 6 and Suppl. Tab. 7.

6. Conclusion

In this paper, we identified that the key to accurate segmentation of NPC is to solve the problem of background dominance. Then, we presented the SeqSeg to deal with the background dominance problem for accurate NPC segmentation. The proposed SeqSeg pertinently solves the background dominance problem at different scales, including instance level and feature level. At the instance level, it aims to capture the relatively small NPC from the giant and complex background. A stable detection model based on DQL was proposed to gradually capture the NPC through the interaction between agents and the environment, to force segmentation to focus attention on the tumor and its surrounding tissue. Additionally, an exploration algorithm was proposed to reduce the exploration space of agents in the detection model; a reward function was proposed to reasonably reward the agent, thus speeding up the convergence and improving the detection performance. At the feature level, the proposed SeqSeg aims to deal with the background dominance problem in which the background is cohesive to the tumor with a shape varied and indistinguishable border. We designed RAMs to guide the segmentation to give attention to tumor-related features by using high-level semantic features to guide learning low-level texture features and perform channel-wise attention and region-wise attention. Additionally, we proposed DBW loss to enlarge the attention on the border and to mine border features.

Experimental results show that the proposed SeqSeg achieves state-of-the-art NPC segmentation performance and demonstrate that solving the background dominance problem is the key to achieving excellent segmentation. Additionally, pertinently solving the background dominance problem at different scales, including the instance level and feature level is an effective strategy. Finally, we conduct a comprehensive experiment and analysis about the potential confounders of segmentation performance on large NPC datasets. Multi-center verification and multi-device verification were also conducted and further verify the reliability of our proposed SeqSeg method.

Declaration of Competing Interest

We declare that we do not have any commercial or associative interest that represents a conflict of interest in connection with the work submitted.

CRediT authorship contribution statement

Guihua Tao: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Haojiang Li:** Investigation, Resources, Data curation. **Jiabin Huang:** Software, Writing – original draft, Data curation, Visualization. **Chu Han:** Writing – review & editing. **Jiazhou Chen:** Writing – review & editing. **Guangying Ruan:** Investigation, Resources, Data curation. **Wenjie Huang:** Investigation, Resources, Data curation. **Yu Hu:** Writing – review & editing. **Tingting Dan:** Data curation. **Bin Zhang:** Data curation. **Shengfeng He:** Writing – review & editing. **Lizhi Liu:** Investigation, Resources, Data curation, Funding acquisition. **Hongmin Cai:** Writing – review & editing, Supervision, Project administration, Funding acquisition.

Acknowledgments

This work was supported in part by the National Key Research and Development Program of China under Grant 2019YFB2102102, the National Natural Science Foundation of China (61771007, 62102153, 62172112, U21A20520, 82171906), Science and Technology Planning Project of Guangzhou City, China (201907010043), the Fundamental Research Fund for the Central Universities (x2jsD2200720), the Key-Area Research and Development of Guangdong Province under Grant 2020B1111190001, and the China Postdoctoral Science Foundation under Grant 2021M691062.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.media.2022.102381.

References

Badrinarayanan, V., Kendall, A., Cipolla, R., 2017. Segnet: a deep convolutional encoder-decoder architecture for image segmentation. IEEE Trans Pattern Anal Mach Intell 39 (12), 2481–2495.
Chan, J.W., Parvathaneni, U., Yom, S.S., 2017. Reducing radiation-related morbidity in the treatment of nasopharyngeal carcinoma. Future Oncol. 13 (5), 425–431.
Chen, H., Dou, Q., Yu, L., Qin, J., Heng, P.-A., 2018. Voxresnet: deep voxelwise residual networks for brain segmentation from 3d mr images. Neuroimage 170, 446–455.
Chen, H., Qi, Y., Yin, Y., Li, T., Liu, X., Li, X., Gong, G., Wang, L., 2020. Mmfnet: a multi-modality mri fusion network for segmentation of nasopharyngeal carcinoma. Neurocomputing.
Han, J., Yang, L., Zhang, D., Chang, X., Liang, X., 2018. Reinforcement cutting-agent learning for video object segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 9080–9089.
He, K., Gkioxari, G., Dollár, P., Girshick, R., 2017. Mask r-cnn. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2961–2969.
He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778.
Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., Meger, D., 2018. Deep reinforcement learning that matters. In: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 32.
Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. Neural Comput 9 (8), 1735–1780.
Hu, J., Shen, L., Sun, G., 2018. Squeeze-and-excitation networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 7132–7141.
Huang, B., Chen, Z., Wu, P.-M., Ye, Y., Feng, S.-T., Wong, C.-Y.O., Zheng, L., Liu, Y., Wang, T., Li, Q., et al., 2018. Fully automated delineation of gross tumor volume for head and neck cancer on pet-ct using deep learning: a dual-center study. Contrast Media & Molecular Imaging 2018.
Huang, J.-B., Zhuo, E., Li, H., Liu, L., Cai, H., Ou, Y., 2019. Achieving accurate segmentation of nasopharyngeal carcinoma in mr images through recurrent attention. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 494–502.

Huang, K.-W., Zhao, Z.-Y., Gong, Q., Zha, J., Chen, L., Yang, R., 2015. Nasopharyngeal carcinoma segmentation via hmrf-em with maximum entropy. In: 2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). IEEE, pp. 2968–2972.
Huang, W.-C., Liu, C.-L., 2012. A hybrid supervised learning nasal tumor discrimination system for dmri. Journal of the Chinese Institute of Engineers 35 (6), 723–733.
Jégou, S., Drozdzal, M., Vazquez, D., Romero, A., Bengio, Y., 2017. The one hundred layers tiramisu: Fully convolutional densenets for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pp. 11–19.
Kaelbling, L.P., Littman, M.L., Moore, A.W., 1996. Reinforcement learning: a survey. Journal of Artificial Intelligence Research 4 (1), 237–285.
King, A.D., Ahuja, A.T., Leung, S.-f., Lam, W.W., Teo, P., Chan, Y.-l., Metreweli, C., 2000. Neck node metastases from nasopharyngeal carcinoma: mr imaging of patterns of disease. Head & Neck: Journal for the Sciences and Specialties of the Head and Neck 22 (3), 275–281.
Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. Imagenet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems, pp. 1097–1105.
Li, Q., Xu, Y., Chen, Z., Liu, D., Feng, S.-T., Law, M., Ye, Y., Huang, B., 2018. Tumor segmentation in contrast-enhanced magnetic resonance imaging for nasopharyngeal carcinoma: deep learning with convolutional neural network. Biomed Res Int 2018.
Li, S., Xiao, J., He, L., Peng, X., Yuan, X., 2019. The tumor target segmentation of nasopharyngeal cancer in ct images based on deep learning methods. Technology in Cancer Research & Treatment 18. 1533033819884561
Lin, L., Dou, Q., Jin, Y.-M., Zhou, G.-Q., Tang, Y.-Q., Chen, W.-L., Su, B.-A., Liu, F., Tao, C.-J., Jiang, N., et al., 2019. Deep learning for automated contouring of primary tumor volumes by mri for nasopharyngeal carcinoma. Radiology 291 (3), 677–686.
Lin, T.-Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S., 2017. Feature pyramid networks for object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2117–2125.
Liu, L., Chen, S., Zhang, F., Wu, F.-X., Pan, Y., Wang, J., 2020. Deep convolutional neural network for automatically segmenting acute ischemic stroke lesion in multi-modality mri. Neural Computing and Applications 32 (11), 6545–6558.
Liu, L., Hodgins, J., 2017. Learning to schedule control fragments for physics-based characters using deep q-learning. ACM Transactions on Graphics (TOG) 36 (3), 1–14.
Man, Y., Huang, Y., Feng, J., Li, X., Wu, F., 2019. Deep q learning driven ct pancreas segmentation with geometry-aware u-net. IEEE Trans Med Imaging 38 (8), 1971–1980.
Men, K., Chen, X., Zhang, Y., Zhang, T., Dai, J., Yi, J., Li, Y., 2017. Deep deconvolutional neural network for target segmentation of nasopharyngeal cancer in planning computed tomography images. Front Oncol 7, 315.
Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., et al., 2015. Human-level control through deep reinforcement learning. Nature 518 (7540), 529–533.
Oktay, O., Schlemper, J., Folgoc, L.L., Lee, M., Heinrich, M., Misawa, K., Mori, K., McDonagh, S., Hammerla, N.Y., Kainz, B., et al., 2018. Attention u-net: learning where to look for the pancreas. arXiv preprint arXiv:1804.03999.
Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In: Advances in Neural Information Processing Systems, pp. 91–99.
Rezatofighi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I., Savarese, S., 2019. Generalized intersection over union: A metric and a loss for bounding box regression. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 658–666.
Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 234–241.
Shelhamer, E., Long, J., Darrell, T., 2017. Fully convolutional networks for semantic segmentation. IEEE Trans Pattern Anal Mach Intell 39 (4), 640–651.
Sims, R., Isambert, A., Grégoire, V., Bidault, F., Fresco, L., Sage, J., Mills, J., Bourhis, J., Lefkopoulos, D., Commowick, O., et al., 2009. A pre-clinical assessment of an atlas-based automatic segmentation tool for the head and neck. Radiotherapy and Oncology 93 (3), 474–478.
Sutton, R.S., Barto, A.G., 2011. Reinforcement learning: an introduction. Cambridge, MA: MIT Press.
Tang, P., Zu, C., Hong, M., Yan, R., Peng, X., Xiao, J., Wu, X., Zhou, J., Zhou, L., Wang, Y., 2021. Da-dsunet: dual attention-based dense su-net for automatic head-and-neck tumor segmentation in mri images. Neurocomputing 435, 103–113.
Wang, Y., Zu, C., Hu, G., Luo, Y., Ma, Z., He, K., Wu, X., Zhou, J., 2018. Automatic tumor segmentation with deep convolutional neural networks for radiotherapy applications. Neural Processing Letters 48 (3), 1323–1334.
Wei, W.I., Sham, J.S., 2005. Nasopharyngeal carcinoma. The Lancet 365 (9476), 2041–2054.
Wen, Z., Zhang, W., Qian, M., 2021. A comprehensive review of deep reinforcement learning for object detection. In: 2021 International Symposium on Artificial Intelligence and its Application on Media (ISAIAM). IEEE, pp. 146–150.
Zhao, L., Lu, Z., Jiang, J., Zhou, Y., Wu, Y., Feng, Q., 2019. Automatic nasopharyngeal carcinoma segmentation using fully convolutional networks with auxiliary paths on dual-modality pet-ct images. J Digit Imaging 32 (3), 462–470.
Zhao, Z.-Q., Zheng, P., Xu, S.-t., Wu, X., 2019. Object detection with deep learning: a review. IEEE Trans Neural Netw Learn Syst 30 (11), 3212–3232.
Zhou, J., Chan, K.L., Xu, P., Chong, V.F., 2006. Nasopharyngeal carcinoma lesion segmentation from mr images by support vector machine. In: 3rd IEEE International Symposium on Biomedical Imaging: Nano to Macro, 2006.. IEEE, pp. 1364–1367.
Zhuo, E.-H., Zhang, W.-J., Li, H.-J., Zhang, G.-Y., Jing, B.-Z., Zhou, J., Cui, C.-Y., Chen, M.-Y., Sun, Y., Liu, L.-Z., et al., 2019. Radiomics on multi-modalities mr sequences can subtype patients with non-metastatic nasopharyngeal carcinoma (npc) into distinct survival subgroups. Eur Radiol 29 (10), 5590–5599.