

APRIL: Anatomical Prior-guided ReInforcement Learning for Accurate Carotid Lumen Diameter and Intima-media Thickness Measurement

Sheng Lian^{a,c,d}, Zhiming Luo^{a,*}, Cheng Feng^b, Shaozi Li^{a,*}, Shuo Li^{c,d,*}

^aDepartment of Artificial Intelligence, Xiamen University, Xiamen, Fujian, China

^bDepartment of Ultrasound, The Second Affiliated Hospital, Southern University of Science and Technology, Shenzhen Third Peoples Hospital, Shenzhen, Guangdong, China

^cDigital Image Group (DIG), London, ON, Canada

^dSchool of Biomedical Engineering, Western University, London, ON, Canada

Abstract

Carotid artery lumen diameter (CALD) and carotid artery intima-media thickness (CIMT) are essential factors for estimating the risk of many cardiovascular diseases. The automatic measurement of them in ultrasound (US) images is an efficient assisting diagnostic procedure. Despite the advances, existing methods still suffer the issue of low measuring accuracy and poor prediction stability, mainly due to the following disadvantages: 1) ignore anatomical prior and prone to give anatomically inaccurate estimation; 2) require carefully designed post-processing, which may introduce more estimation errors; 3) rely on massive pixel-wise annotations during training; 4) can not estimate the uncertainty of the predictions. In this study, we propose the **Anatomical Prior-guided ReInforcement Learning** model (**APRIL**), which innovatively formulate the measurement of CALD & CIMT as an RL problem and dynamically incorporate anatomical prior (AP) into the system through a novel reward. With the guidance of AP, the designed keypoints in APRIL can avoid various anatomy impossible mis-locations, and accurately measure the CALD & CIMT based on their corresponding locations. Moreover, this formulation significantly re-

* Corresponding author

Email addresses: zhiming.luo@xmu.edu.cn (Zhiming Luo), szlig@xmu.edu.cn (Shaozi Li), slishuo@gmail.com (Shuo Li)

duces human annotation effort by only using several keypoints and can help to eliminate the extra post-processing steps. Further, we introduce an uncertainty module for measuring the prediction variance, which can guide us to adaptively rectify the estimation of those frames with considerable uncertainty. Experiments on a challenging carotid US dataset show that APRIL can achieve MAE (in pixel/mm) of 3.02 ± 2.23 / 0.18 ± 0.13 for CALD, and 0.96 ± 0.70 / 0.06 ± 0.04 for CIMT, which significantly surpass popular approaches that use more annotations.

Keywords: CALD and CIMT Measurement, Anatomical Prior (AP), Reinforcement Learning (RL).

1. Introduction

Carotid artery lumen diameter (CALD) and carotid artery intima-media thickness (CIMT) have been reported as critical factors for measuring the risk of many cardiovascular diseases, including atherosclerotic, myocardial infarction, and stroke (Bauer et al., 2012; Mathai et al., 2019; O’Leary & Bots, 2010; Saba et al., 2019). Besides, CALD and CIMT are also used for evaluating the effect of treatment or assessing the potential impact of certain factors (such as smoking) on cardiovascular disease (Tell et al., 1994; Lonn et al., 2001). So far, the most popular imaging technology for evaluating carotid artery diseases is B-mode ultrasound (US) imaging, for its low costing, non-invasive, painless examination, and non-radiation (Barth, 2002; Stein et al., 2008). The common carotid artery is close to the skin surface and is approximately parallel to the skin, making it easy to be measured through US imaging. As illustrated in the typical carotid artery ultrasound (CCA-US) (Figure 1(a)), the CALD refers to the distance between the near-wall (NW) and the far-wall (FW). The CIMT is the distance between intima-lumen interface (LI) and the media-adventitia interface (MA) on FW. At present, the CALD and CIMT in CCA-US images are manually measured by well-trained cardiologists, which is extremely time-consuming and tedious. Given the scarcity of our medical resources, it is with an urgent need

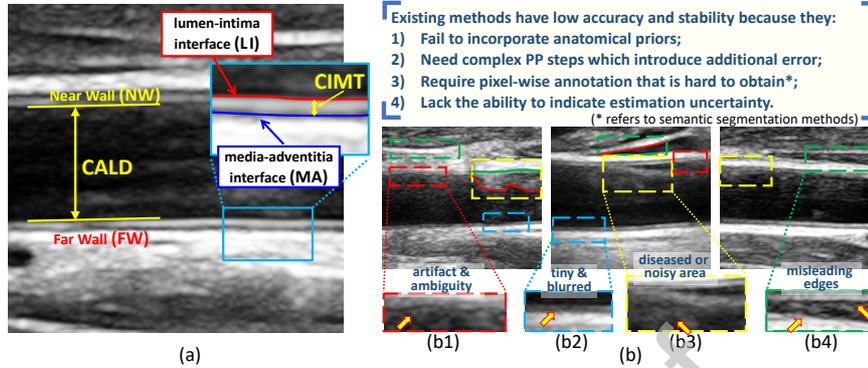


Figure 1: (a) The schematic diagram of CALD, CIMT, and other tissue structures in the CCA-US image. (b) Several difficult samples show different kinds of challenges, including (b1) artifact and ambiguous areas in the edge area, (b2) tiny intima-media interface structure, (b3) diseased or noisy areas, (b4) irrelevant areas look similar to artery edges. GT and predictions of some edges are marked with solid green and red lines, which are easily mis-located with anatomically incorrect, and thus lead to low accuracy and stability. *PP* is short for post-processing.

20 to build an automatic method for accurately measuring the CALD and CIMT. However, there are some difficulties during the measurement in CCA-US. 1) Periodic carotid artery motion is often accompanied by artifacts, and will bring many ambiguous areas; 2) The intima-media interfaces are close to each other and have a fuzzy boundary, making it difficult to distinguish the two lines; 3) 25 There are many noisy or uncertain areas in CCA-US image, especially in the diseased areas; 4) Some tissues have similar appearances as the vessel wall in CCA-US images, which are easily misidentified as NW or FW. We illustrated these difficulties in the dashed boxes in Fig. 1(b1 ~ b4).

30 During the past decades, several different types of approaches have been proposed for this task by locating the edges of NW, LI, and MA, including the gradient-based edge detection methods (Pignoli & Longo, 1988; Liguori et al., 2001; Stein et al., 2005; Golemati et al., 2007; Faita et al., 2008), the active contour-based methods (Gutierrez et al., 2002; Cheng et al., 2002; Loizou et al., 2007; Petroudi et al., 2012; Zhao et al., 2017b), the machine learning-based

35 methods (Menchón-Lara et al., 2014; Menchón-Lara & Sancho-Gómez, 2015; Shin et al., 2016; Qian & Yang, 2018; Biswas et al., 2018; Xie et al., 2019; Zhou et al., 2019; Zhao et al., 2020; Vila et al., 2020). Although these methods have made significant advances, they still suffer the issue of low measuring accuracy (measured by mean absolute error) and poor stability (measured by the magnitude of change in absolute error), mainly due to the following disadvantages: 1) These methods are mostly appearance-based, which do not consider the carotid anatomical information, and prone to give anatomical incorrect estimation and thus lower the accuracy and stability. 2) These methods are susceptible to various noises, and a carefully designed post-processing step is needed to transform the estimated edges or segmentation masks to CALD or CIMT, which may introduce more estimation errors. 3) Although modern deep learning-based segmentation methods can improve the measurement accuracy to some extent, they require a large amount of pixel-wise labeled data for training, which is usually hard to obtain. 4) Due to the complexity of the data and the limitation of algorithms, these methods do not have a suitable way to estimate the uncertainty of the predictions, especially for those results deviate significantly from the ground truth (GT).

In the context of medical image analyzing, anatomical priors (AP) can provide more constraints in terms of the shape and locations. As shown in Fig. 1(a), the carotid artery in CCA-US images shows several distinct anatomical priors, *i.e.*, **1)** NW, LI, and MA are straight-line-like edges arranged from top to bottom; **2)** LI and MA are close to each other and have similar directions. However, these AP factors have not been incorporated into the previously mentioned methods for improving accuracy. On the other hand, reinforcement learning (RL) allows the agents to accomplish a complex sequential task through interactions with the environment. This sequential learning mechanism enables us to integrate above AP factors into the reward as feedback during the step-wise interaction between agent and environment. As such, we could formulate the CALD & CIMT estimation task as an RL problem, in which a set of keypoints represents the edges of NW, LI, and MA. The main goal is to perform a sequen-

tial movement to locate the keypoints into their correct positions guided by an anatomical prior sensitive reward function.

In this study, we propose the **A**natomical **P**rior-guided **R**e**I**nforcement **L**earning model (**APRIL**) for CALD & CIMT measurement. We innovatively formulate the localization of NW, LI, and MA as a multi-agent RL problem, that each agent contains several keypoints for representing the edge, and the main goal is to find the correct positions in the US frame. To incorporate the AP into the learning procedure, we propose an anatomical prior sensitive reward function to guide the multiple agents' mutual movements. Besides, we further introduce an extra uncertainty estimation module to adaptively rectify the estimation errors in those frames with more substantial prediction uncertainty. Finally, we conduct experiments on a challenging CCA-US dataset containing 4351 CCA-US frames to evaluate our proposed method's effectiveness.

To summarize, the main contributions of this study are as follows:

- We formulate the CALD & CIMT measurement as a multi-agent RL task and propose a novel AP guided reward function for the optimization. Based on this formulation, we can avoid the anatomically inaccurate estimations encountered in previous methods and achieve better performance.
- Our formulation significantly reduces human annotation effort by only using several keypoints representing the carotid edges, and help to get rid of carefully designed post-processing steps and the resulting additional errors.
- We introduce an extra uncertainty estimation module, which enables us to adaptively rectify the estimation errors in those frames with higher uncertainty. Experimental results on a challenging dataset demonstrated that our APRIL outperforms other representative comparing methods by a large margin in both accuracy and stability.

2. Related Work

During the past decades, researchers proposed various methods for tackling the CALD or CIMT estimation tasks from the following three different categories: gradient-based edge detection methods (Liguori et al., 2001; Stein et al., 2005; Golemati et al., 2007; Faita et al., 2008), active contour-based methods (Gutierrez et al., 2002; Cheng et al., 2002; Loizou et al., 2007; Petroudi et al., 2012), and machine learning-based methods (Menchón-Lara et al., 2014; Menchón-Lara & Sancho-Gómez, 2015; Shin et al., 2016; Qian & Yang, 2018; Biswas et al., 2018; Zhou et al., 2019; Zhao et al., 2020; Vila et al., 2020). In this section, we will briefly review these three different types of methods.

Gradient-based edge detection methods aim to locate edges representing carotid walls by modeling the intensity profile distribution, or computing gradient of CCA-US images. Liguori et al. (2001) assumed that the artery is horizontally placed in the CCA-US images, and then detected the carotid wall based on horizontal gradients of the intensity profile. Golemati et al. (2007) applied a series of image processing techniques and the Hough transform to locate lines representing NW, LI, and MA edges. Faita et al. (2008) proposed an improved gradient-based method by using a first-order absolute moment edge operator (FOAM), together with a heuristic search and thresholding process. However, these methods are not robust enough to deal with the noisy CCA-US images and need to adjust the hyper-parameters manually.

Active contour-based methods are iterative region-growing image segmentation algorithms (Kass et al., 1988), which transform the task of carotid wall location into an energy optimization problem. The study by Loizou et al. (2007) initialized the contour through a series of graphic operations, and adopted the active contour method proposed by Williams & Shah (1992) for locating carotid walls. Besides the internal energy and external energy in original active contour, Gutierrez et al. (2002) further utilized the damping force introduced by Lobregt & Viergever (1995) to ensure the smoothness and stability during the deformation process. Delsanto et al. (2007) proposed a combined approach of

local intensity statistics and active contour to perform the CIMT measurement. However, active contour-based methods heavily rely on initialization and the
 125 corresponding energy functions. They also tend to be easily affected by image noises and find local optimal solutions.

Machine learning-based methods. Benefited from the large-scale labeled dataset and the rapid development of computation power, various machine learning methods have been proposed for the measurement of the CALD
 130 or CIMT. Menchón-Lara et al. (2014) utilized a multi-layer perceptron to classify the pixels inside the FW regions as either '*carotid wall boundary*' pixels or not. Shin et al. (2016) combined the active contour model with a patch-based CNN for locating the LI and MA edges in the CCA-US images. Biswas et al. (2018) and Vila et al. (2020) adopt the fully connected network (FCN)
 135 and DenseNet to segment the whole carotid image into different regions, and then perform the estimation of CALD or CIMT using a series of post-processing (PP) steps. Such PP steps introduce complex operations including (1) largest connected component extraction, (2) edge smoothing, (3) outliers removal, etc., which are with additional parameters and low robustness. These operations re-
 140 quire careful design and lack the flexibility for different data, thus may introduce additional errors. Zhao et al. (2020) combined the CNN regression network with a bidirectional optical flow model to resolve the inconsistencies between CCA-US slices caused by artery motion, however this method required continuously labeled frames. Although these recently emerged deep learning-based methods
 145 can improve measurement accuracy, they still require a large amount of pixel-wise annotations for training their model and fail to incorporate anatomical priors of CCA-US images.

3. Background

Reinforcement Learning (RL) enables the agent to learn complex tasks
 150 by interacting with the environment by taking actions exploratively (for uncharted territory) and exploitative (for current knowledge). Among the key

components of RL, the *environment* is the physical world in which the agent operates. The *state* is the current situation of the agent. The *reward* is the feedback that the agent interacts with the environment after taking an *action*.
 155 The general workflow of RL is summarized as follows, where the agent interacts with the environment over a sequence of discrete time steps. At each step t , the agent receives the representation of current state S_t , performs an action A_t , and will get a reward R_t . Generally, RL problems are modeled and solved based on Markov Decision Processes (MDP)(Kaelbling et al., 1996; Sutton & Barto,
 160 2018). Many representative solutions have been proposed during the development of RL, such as Q-learning (Watkins & Dayan, 1992) and deep Q-network (DQN) (Mnih et al., 2015).

Deep Q-network (DQN): Deep learning-based methods have achieved remarkable success in various research areas, and deep neural networks (DNN)
 165 have been integrated into the RL as a powerful function approximator. Mnih et al. (2015) proposed DQN, which approximates the action-value function by $Q(s, a; \omega) \approx Q(s, a)$, where ω represents the network parameters. The DQN combines the merits of Q learning (Watkins & Dayan, 1992) and DNN for solving complex sequential learning tasks. Like Q-learning, the agents in DQN
 170 interact with the environment by taking actions calculated by the DNN model and receiving the reward signal r . The corresponding main objective is to find the optimal policy by maximizing the cumulative future rewards. The loss function of DQN goes as Eq 1.

$$L_{DQN}(\omega) = \mathbb{E}[(r + \gamma \max_{a'} Q_{target}(s', a'; \omega^-) - Q_{net}(s, a; \omega))^2] \quad (1)$$

This function follows the form of mean squared error (MSE), where r , s , a and γ
 175 are *reward*, *state action*, and *discount factor*, respectively. Q is the approximated action-value function. The s' and a' represent the next state and action. For avoiding the instability caused by rapid policy changes, DQN uses $Q_{target}(\omega^-)$, a temporarily fixed version of $Q_{net}(\omega)$. The Q_{target} 's parameters ω^- are updated to be equal to Q_{net} periodically for every N_{target} steps. Furthermore, DQN

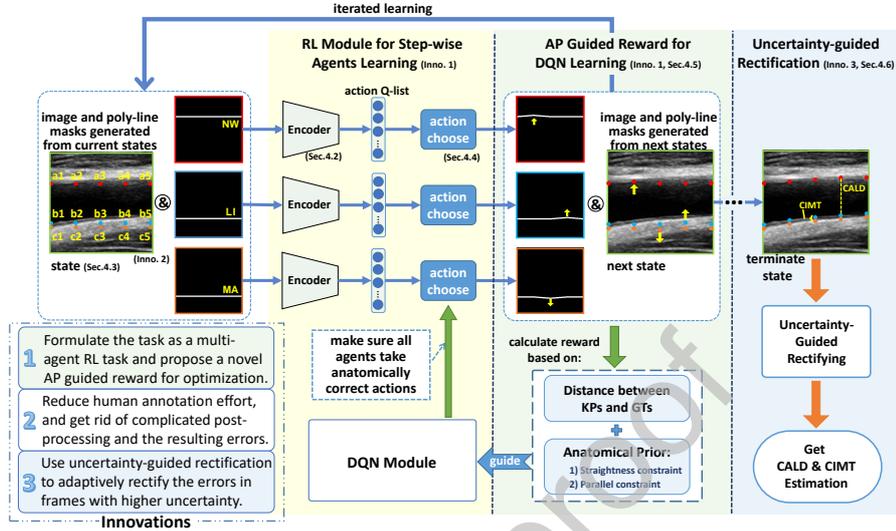


Figure 2: The framework of the proposed APRIL model including 1) RL module for step-wise agents learning; 2) AP guided reward for DQN learning; 3) Uncertainty-guided rectification. *KP* and *GT* are short for *keypoint* and *ground truth*, respectively. *Inno* is short for *innovations*, which are summarized in the bottom left corner.

180 introduces an experience replay buffer that stores transitions of (s, a, r, s') , and the buffer is sampled randomly to create the mini-batches for training. In this way, the correlation of data is broken, and this is helpful for network convergence and training stability. There are also some DQN variants to improve DQN from different perspectives, such as double DQN (Van Hasselt et al., 2016), duel
 185 DQN (Wang et al., 2016).

RL, especially deep learning integrated RL, has been proved effective in various medical image analysis tasks including segmentation (Sahba et al., 2006; Wang et al., 2011; Man et al., 2019), landmark localization (Ghesu et al., 2017; Alansary et al., 2019; Vlontzos et al., 2019), tissue detection (Luo et al., 2019;
 190 Maicas et al., 2017; Ali et al., 2018), registration (Ma et al., 2017; Krebs et al., 2017; Liao et al., 2017).

4. Anatomical Prior-guided Reinforcement Learning for CALD & CIMT Measurement

In this study, we formulate the CALD and CIMT measurement task as an Anatomical Prior-guided Reinforcement Learning (APRIL) problem, and utilize the deep Q-network (DQN) (Mnih et al., 2015) to perform the step-wise learning procedure. In this section, we describe the overall framework of APRIL in Sec. 4.1, and explain the main components of APRIL, including the environment and feature encoder (Sec.4.2), the state representation (Sec.4.3), the actions for locating keypoints (Sec.4.4), the anatomical prior-guided reward for DQN learning (Sec.4.5), and the uncertainty guided rectification module (Sec.4.6).

4.1. The overall framework of APRIL

The overall framework of APRIL is illustrated in Fig. 2. The proposed APRIL adopts three agents to control the five keypoints representing NW, LI, and MA edges, respectively. At each learning step, each agent in APRIL receives the representations of the current state, and generates the action Q-list accordingly through the encoder. Specifically, the action Q-list is a vector containing the Q value of every action, where the Q value of each action is the basis for the agent to choose the action. The reward is calculated by considering both the accuracy of predicted locations of keypoints, and the fitness correlated to AP factors. Through the process of exploration and exploitation, the agents learn to make sequential decisions to move the keypoints towards the target locations, thereby the estimations of CALD and CIMT can be calculated by simple coordinate subtraction. The uncertainty-guided rectification module further rectifies the predicted results for those frames with considerable uncertainty.

4.2. Environment and Feature Encoder

The environment is the carrier of RL operation, in which the agent takes actions and gets the response, etc. In our APRIL's implementation, the environment defines the composition of the state, the setting of action, the cal-

culation method of *reward*, and so on. In each learning step of APRIL, the environment of current state contains current CCA-US image and positions of all the keypoints. The feature encoder will extract features for representing the current input, and the agents will choose the corresponding actions for current step based on the extracted features. Therefore, the feature encoder will significantly influence the final measure accuracy, and it is crucial to design an accurate and efficient feature encoder.

In this study, we adopt the SE-ResNext (SERN) module described in (Hu et al., 2018) to construct the feature encoder, which combines the advantages of SE-Net (Hu et al., 2018) and ResNeXT (Xie et al., 2017). The structure of such encoder is illustrated in the left part of Fig. 3, and the auxiliary path for uncertainty-guided rectification in the right will be further described in Sec. 4.6. In our feature encoder, we adopt four SERN blocks, one average pooling layer, a dropout layer, and two fully-connected layers to compute the final feature vector. Each SERN block consists of the structure of split-transform-merge, shortcut connection, and the squeeze-and-excitation operation. More details of the SERN can be found in (Hu et al., 2018). Notice that it is compatible with using other different CNN model as feature encoder. At each step, this feature encoder takes the current CCA-US frames along with the corresponding poly-line masks as the input (described in Sec. 4.3), and generates the predicted Q-value of each action for further strategy selection.

4.3. State representation in APRIL

In our implementation, the state at each step contains the current CCA-US image and the current positions of all keypoints related to NW, LI, and MA edges. Given that all the keypoints are evenly distributed horizontally and all the horizontal coordinates are fixed, the state only stores the vertical coordinates of all keypoints, which goes as $[a_1, a_2, \dots, a_5]$, $[b_1, b_2, \dots, b_5]$, and $[c_1, c_2, \dots, c_5]$. To better extract features from the current state, we generate poly-line masks based on the position of all keypoints in the current step and concatenate this mask with the input CCA-US image to form a state representation. An example

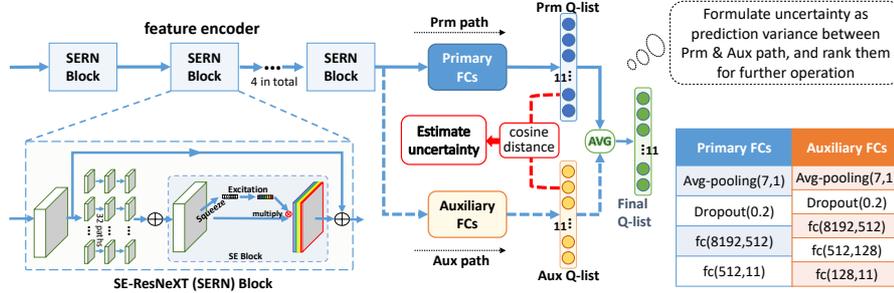


Figure 3: The SERN-based feature encoder adopted in APRIL (Sec.4.2). At the end of the traditional encoder structure, we added an auxiliary path for uncertainty-guided rectification (Sec.4.6). On the right side, we list the network structure of the primary and auxiliary path. 'Prm' and 'Aux' are short for primary and auxiliary, respectively.

of a state representation is depicted in the top left corner of Fig 2. Specifically, for generating a poly-line mask, we connect the coordinates of all current keypoints with a line segment with a width of three pixels, and set the pixel value of background and the poly-line area to 0 and 255, respectively.

255 4.4. Actions for locating keypoints

In the proposed APRIL, we have three agents to locate the positions of NW, LI, and MA, which each contains five keypoints. The horizontal coordinates of these keypoints are fixed, and our goal is to locate the carotid artery wall accurately by moving each keypoint up or down towards its correct position. Then at each step, we define the action for each agent as moving one of five keypoints up or down at a unit length, or not moving any keypoints. Therefore, we have an action space containing $2*5+1 = 11$ different actions for each agent.

4.5. Anatomical Prior-Guided Reward for DQN learning

In APRIL, the reward function calculates the benefit after moving keypoints at each step, which consists of 1) the distance reduction R_{gt} towards the annotated GT; and 2) the increment of satisfaction R_{ap} related to the anatomical

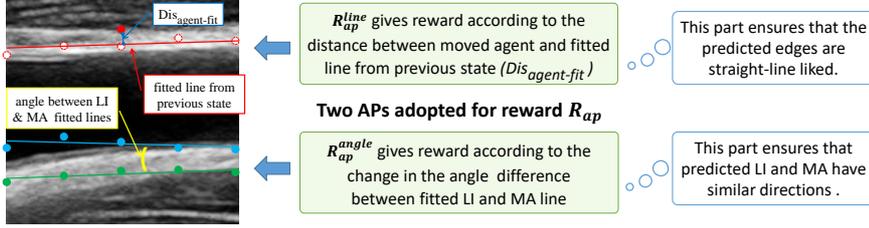


Figure 4: Some AP factors considered in APRIL model's reward, including: R_{ap}^{line} and R_{ap}^{angle} . The colored lines are fitted lines from agents of NW, LI, and MA.

prior (Fig. 4), defined as

$$R = R_{gt} + \lambda \overbrace{(R_{ap}^{line} + R_{ap}^{angle})}^{R_{ap}}, \quad (2)$$

where λ is a hyper-parameter weighting influence of AP and set to 0.2.

The first term R_{gt} computes the distance reduction of all keypoints after each agent's action at a time step, which goes as:

$$R_{gt} = \sum_{i=1}^N |y_{pred}^i - y_{gt}| - |y_{pred-}^i - y_{gt}^i|, \quad (3)$$

where N is the number of keypoints, y_{gt}^i , y_{pred}^i and y_{pred-}^i represent the i th keypoint's y -coordinate of GT, before the action, and after the action, respectively.

For the R_{ap} , we consider a straight line fitness penalty R_{ap}^{line} , and a parallel fitness gain R_{ap}^{angle} between the estimated LI and MA edges (Fig.4). After moving one keypoint, if it has a negative influence on the straightness of its corresponding line, we will assign a penalty to R_{ap}^{line} . The penalty equals to the distance between the moved keypoint and fitted line of other keypoints representing the current edge, and such penalty implicitly guarantees the smoothness between different keypoints. Moreover, since the LI and MA have similar directions, we can compute the angle difference reduction between them based on the slope of their corresponding fitting line after an action. Given the keypoints, we firstly estimate the slopes and intercepts of the fitting straight line related to NW, LI, and MA, by $k, b = \arg \min_{k, b} \sum_{i=1}^n (kx_i + b - y_i)^2$. After moving a

keypoint i , we compute the R_{ap}^{line} as using the negative distance of keypoint i
 285 to its corresponding fitting line. If the action is taken by the agents of LI and
 MA , then we compute the R_{ap}^{angle} by:

$$R_{ap}^{angle} = |\theta(k_{LI}) - \theta(k_{MA})| - |\theta(k_{LI-}) - \theta(k_{MA-})| \quad (4)$$

where the k_{LI} , k_{MA} , k_{LI-} , k_{MA-} represent the slope of the fitting line of LI
 and MA before and after taking an action, and θ is the arc-tangent function for
 computing corresponding angles.

290 In this reward function, we consider the relative position between the key-
 points & GTs, and the degree of satisfaction with AP factors. **Note that these
 constraints are not strict restrictions, and slight violations are allowed (e.g.,
 curve and pathology areas in Fig. 6 {B, C, D, E}).** The three designed agents
 interact with each other and select the appropriate actions for the related key-
 295 points. Notice that, if any keypoint reaches the edge of the image, we directly
 give a large negative value as a reward (-1000 in our implementation), and ter-
 minate this episode of exploration. Besides, as mentioned in the introduction
 that NW , LI , and MA are arranged from top to bottom, and LI and MA are
 close to each other, we directly incorporate this prior information by initializing
 300 the keypoints into appropriate positions.

4.6. Uncertainty guided rectification module

When the carotid artery is in the state of contraction/relaxation changes,
 there will be much noise in the CCA-US frame. It will blur the edge of vessel
 walls and cause inaccurate predictions of our model. Estimating the uncertainty
 305 of neural network predictions is an effective way for dealing with noise in data,
 and quantifying the reliability of predictions, especially for life-critical appli-
 cations such as medical image analysis and autonomous driving (Eaton-Rosen
 et al., 2018; Postels et al., 2019). Therefore, we introduce an uncertainty guided
 rectification module to rectify the results of those noisy frames in the testing
 310 stage. Explicitly, we assign the predicted values of the frames with high un-

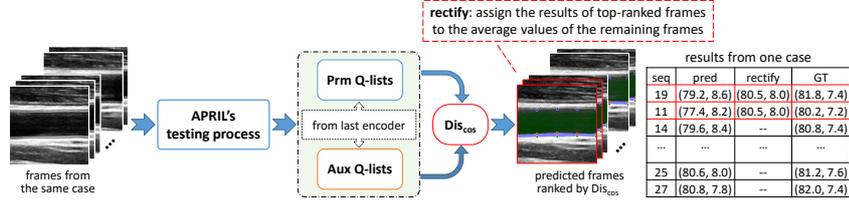


Figure 5: The process of uncertainty driven rectification for a group of frames from one case. Prm and Aux Q-lists are the outputs of primary and auxiliary paths of APRIL’s encoder, as is indicated in Fig. 3. On the right side, the red borders indicate the frames that need to be further rectified after uncertainty ranking. Specifically, we rank the calculated cosine distance of the two predictors and assign the results of the top-ranked frames to the average values of CALD and CIMT of the remaining frames.

certainty to the average values of the remaining frames in the same case (as illustrated in Fig. 5).

Inspired by the work proposed by Zheng & Yang (2020), we formulate the uncertainty in this task as the prediction variance. Specifically, on the feature encoder introduced in Sec 4.2, we add an auxiliary predictor for the same task as the primary predictor. Both of them predict the Q-values for each action. After the four SERN blocks of our feature encoder, the network is divided into two branches: the primary branch, and the auxiliary branch with a slightly different structure (Fig. 3). By leveraging the discrepancy of prediction between the primary and the auxiliary branches, we get an estimation of the degree of uncertainty. In detail, we calculate the cosine distance between the outputs from the primary branch and auxiliary branch by

$$\text{Dist}_{cos} = 1 - \cos(\theta) = 1 - \frac{\mathbf{A} \cdot \mathbf{B}}{\|\mathbf{A}\| \|\mathbf{B}\|} \quad (5)$$

where \mathbf{A} and \mathbf{B} represent the estimation of the primary and the auxiliary branches, respectively.

After getting the sequential predictions of a patient, we set a [thresholding ratio](#) α_{unc} to pick out those frames with a higher degree of uncertainty to perform a rectification. Since frames with high uncertainty often appear in the vasomo-

tion stage, it is reasonable to assign their results as the average of the remaining frames with higher confidence. Therefore, we directly assign their CALD and CIMT as the average values of the remaining $(1 - \alpha_{unc})$ frames with higher confidence. A schematic diagram of this process is shown in Fig. 5. In practice, we also can only leverage the frames with higher confidence for the final CALD and CIMT measurement. Our model can correct outliers adaptively according to the actual characteristics of different cases and the uncertainty predictions, instead of setting fixed thresholding for all cases or requiring constant manual adjustment of the thresholding value as in the existing approaches.

5. Experimental Setup

5.1. Dataset

We evaluate the proposed APRIL method on a challenging dataset, which contains 4351 annotated CCA-US frames from 95 patients collected at a medical center by a MINDRAY Resona 7. There are 54 male patients and 41 female patients with a mean age of 33 ± 9.6 years old. The dataset contains 187 CCA-US sequences from both left and right carotid arteries of 92 patients, and only left or right of the rest three patients. Each sequence has 175 frames collected in a longitudinal B-mode at a frame time of 28.6 ms. The spatial resolution is 0.060 mm/pixel. For measuring the changes of CALD and CIMT during carotid vasomotion, we randomly annotated 20 ~ 27 frames at the interval of 1 ~ 4 frames for each sequence. The annotation procedure is accomplished by annotating the NW, LI, and FW in those selected frames, and two carotid physicians with more than 10-year experiences help to double-check the annotations.

For the training and testing, we randomly select 137 sequences from 69 patients (3282 frames) for training, 17 sequences from 9 patients (404 frames) for validation, and 33 sequences from 17 patients (665 frames) for testing. All frames are uniformly cropped to $256 * 256$ to unify with clinical practice (Menchón-Lara & Sancho-Gómez, 2015; Golemati et al., 2007; Loizou et al., 2007; Biswas et al., 2018) and improve computing efficiency. Note that this

process is accomplished by our self-designed labeling tool, and for all frames of the same case, we simply locate a $256 * 256$ box once to complete ROI location process.

360 5.2. Implementation details

We implement our proposed APRIL by using the PyTorch toolbox with the OpenAI Gym platform and train the model in a device with Intel E5-2650 CPU and 4 * NVIDIA P100 Pascal GPU. During the training phase, we adopt an ϵ -greedy exploration strategy for choosing the actions of agents. Specifically, 365 the ϵ is initialized to 0.6 and is increased by 0.1 for every 25 epochs till 0.9. The experience replay buffer capacity is 600, and the discount factor (γ in Eq. 1) is 0.9. We set N_{target} to 50, where N_{target} is the update frequency of Q_{target} in Eq.1. In the training stage, for all the three encoders, we adopt the weights pre-trained with ImageNet and train them in a separate manner simultaneously. 370 We use MSE loss for measuring the distance between action Q-lists predicted by Q_{target} and Q_{net} in Eq. 1, and the MSE loss goes as

$$l(x, y) = \{l_1, \dots, l_N\}^T, l_n = \frac{1}{M} \sum_{m=1}^M (x_n^m - y_n^m)^2 \quad (6)$$

where N is the batch size, M is the length of the action Q-list. For the n_{th} batch, x_n^m is the m_{th} Q value in the action Q-list predicted by Q_{net} , and y_n^m is the m_{th} action Q-list predicted by Q_{target} . For the optimization, we use Adam 375 optimizer with an initial learning rate of $1e - 4$ and a batch size of 32.

After computing the average statistical position of the NW, LI, and MA in the training set, we initialize their corresponding five keypoints evenly along the horizontal axis at the vertical coordinates of 80, 180, 190, respectively. By using this initialization strategy, we can get the benefits of ensuring all keypoints 380 reach accurate positions in as few steps as possible, and utilizing the relative position prior information of the NW, LI, and MA.

During the training phase, we treat the process of agents perform a series of actions in one CCA-US frame as an episode, and limit the maximum number

of actions per episode to 180. In the testing phase, the agents follow the policy
 385 learned by APRIL, which chooses the action with the highest Q-value at each
 step. If all three agents select the same keypoints and oscillate in the same
 positions four times, we considered that a suitable edge had been found and
 terminated. Otherwise, our max step per episode is 120 in the testing phase,
 which takes around 4.7s for each CCA-US frame. Specifically, the setting of all
 390 DQN-specific hyper-parameters, including ϵ for ϵ -greedy exploration strategy
 and γ for discount factor in Eq. 1, refers to (Alansary et al., 2019; Man et al.,
 2019). While other task-specific parameters, including the line-width of the
 poly-line masks and the number of keypoints, are obtained based on the results
 on validating set.

395 6. Experimental result and analysis

We evaluate the CALD and CIMT measurement accuracy by computing the
 mean absolute error (MAE) between the predicted values with the ground truth
 (GT) of all testing frames, which goes as:

$$MAE = \frac{1}{N} \sum_{n=1}^N |p_n - g_n| \quad (7)$$

where N is the number of testing frames, and p_n and g_n denote the prediction
 and GT of CALD/CIMT of the n th frame, respectively. In the following, we
 compare the proposed APRIL with other representative methods in Sec.6.1, and
 perform an ablation study to validate the effectiveness of our model design and
 400 components in Sec.6.2.

Table 1: MAEs and standard deviations of CALD & CIMT from APRIL and other representative methods. PP is short for post-processing.

method	year	description	CALD		CIMT	
			in <i>pixel</i>	in <i>mm</i>	in <i>pixel</i>	in <i>mm</i>
Golemati et al. (2007)	2007	Edge detect & Hough transform	9.71 ± 5.66	0.58 ± 0.34	2.06 ± 1.79	0.12 ± 0.11
Loizou et al. (2007)	2007	Active contour (snakes)	13.67 ± 7.81	0.82 ± 0.47	2.68 ± 2.33	0.16 ± 0.14
Menchón-Lara et al. (2014)	2014	CNN-based edge detect + PP	7.08 ± 5.21	0.42 ± 0.31	1.61 ± 1.33	0.10 ± 0.08
Shin et al. (2016)	2016	CNN-based edge detect + snakes	7.79 ± 4.15	0.47 ± 0.25	1.57 ± 1.79	0.09 ± 0.11
Vila et al. (2020)	2020	CNN-based segment + PP	6.72 ± 4.96	0.40 ± 0.30	1.54 ± 1.25	0.09 ± 0.08
U-Net (Ronneberger et al., 2015)	2015	U-Net + Vila et al. (2020)'s PP	6.98 ± 5.28	0.42 ± 0.32	1.64 ± 1.83	1.10 ± 0.11
DeepLabv3+ (Chen et al., 2018)	2018	DeepLabv3+ + Vila et al. (2020)'s PP	6.56 ± 4.81	0.39 ± 0.29	1.52 ± 1.60	0.09 ± 0.10
PSPNet (Zhao et al., 2017a)	2017	PSPNet + Vila et al. (2020)'s PP	6.85 ± 5.16	0.41 ± 0.31	1.47 ± 1.62	0.09 ± 0.10
APRIL	-	proposed method	3.02 ± 2.23	0.18 ± 0.13	0.96 ± 0.70	0.06 ± 0.04

6.1. Comparison with other representative methods

We compare the proposed APRIL with the three categories of CALD/CIMT estimation methods described in Sec. 2, including gradient-based edge detection methods (Golemati et al., 2007), active contour-based methods (Loizou et al., 2007), machine learning-based methods (Menchón-Lara et al., 2014; Shin et al., 2016; Vila et al., 2020). All these comparing methods are reproduced on our proposed dataset according to the parameters described in the original literature. We summarize the results of the proposed APRIL and other representative comparing approaches in Table 1, and show the visualization results of some challenging cases in Fig. 6.

Gradient-based edge detection methods: As a gradient-based edge detection method, Golemati et al. (2007) consists of ROI extraction, morphologi-

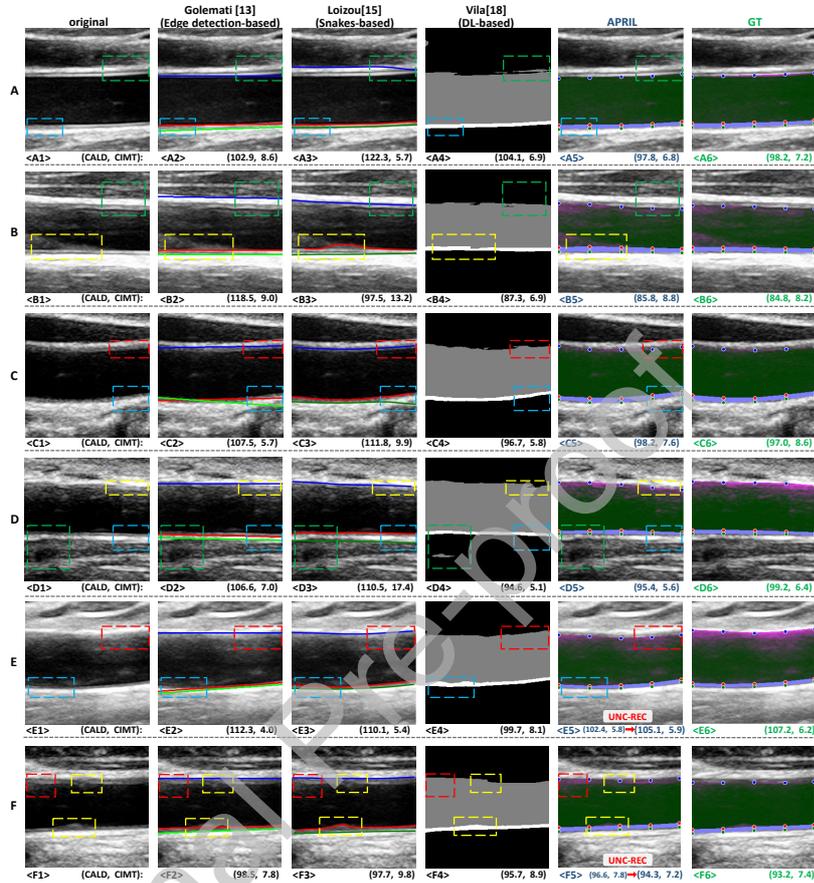


Figure 6: The visualization results of the representative models and the proposed APRIL. We mark some challenging areas of each picture with colored dashed boxes as Fig. 1, Red: artifact and ambiguous areas in the edge area, Blue: tiny intima-media interface structure, Yellow: diseased or noisy areas, Green: irrelevant areas look similar to artery edges. E1 and F1 are picked out for uncertainty guided rectification (marked with red arrows), where the detailed rectification process is illustrated in Fig. 5.

cal operations, edge detection, and Hough transform, which need to be carefully tuned for a specific frame. Besides, as shown in Fig. 6 [C2, E2], since the result of (Golemati et al., 2007) is calculated based on the distance between the straight lines predicted by Hough transform, the accuracy of this method will be further reduced when the direction of the carotid artery has a considerable

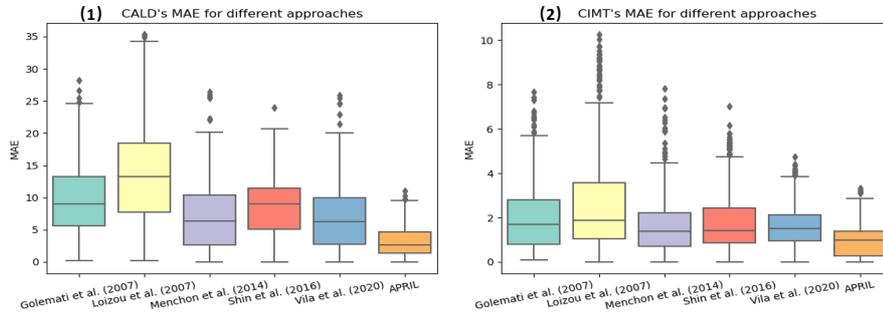


Figure 7: Boxplots of CALD & CIMT's MAEs (in pixel) for comparative approaches.

deflection.

In practice, the predictions of this kind of approaches are not ideal and fluctuate significantly among the time series, since the difference between CCA-US images is relatively significant and with much noise. On the other hand, these methods are based on low-level features such as edges and gradients and do not incorporate AP factors. They are less robust to the CCA-US image diversity and are prone to produce anatomical incorrect edge predictions, such as crossed LI & MA edges (Fig. 6 [C2, D2, E2]).

Active contour-based methods: Similar to gradient-based detection methods, the representative active contour-based method Loizou et al. (2007) also need a series of carefully tuned operations such as cropping and despeckle to initialize snake contour. As shown in the third column of Figure 6, the estimated contours can adequately fit the edges' curvature when the carotid wall is unambiguous. However, they still suffer from low accuracy and fluctuation, mainly due to 1) rely on the proper initialization of snakes, 2) lack the constraint from AP factors, 3) the snakes tend to be biased or get stuck in local minima. Therefore, active contour-based methods tend to generate prediction errors with a significant anatomical violation, such as abnormal bends due to the lesion or abnormality on carotid walls (Fig. 6 [B3]) and abnormal bends due to misidentification of unrelated tissues (Fig. 6 [A3, C3]).

Deep learning-based methods: We also compare APRIL with deep learning-based methods, including Menchón-Lara et al. (2014) and Shin et al. (2016) that directly locate carotid walls, and Vila et al. (2020) that accomplish the task through multi-class semantic segmentation. For fair comparisons, we changed the backbone networks of these methods to the same SE-ResNeXt (SERN) as our APRIL (U-shaped model with SERN as encoder for semantic segmentation tasks). The comparing results in Table. 1 indicated that the pipeline and the operations adopted in Vila et al. (2020) has better performance than Shin et al. (2016) and Menchón-Lara et al. (2014). Thus, for more comparison, we switch the segmentation model in Vila et al. (2020) to other representative models that achieved state-of-the-art performance in medical imaging tasks, including U-Net (Ronneberger et al., 2015), DeepLabv3+ (Chen et al., 2018), and PSPNet (Zhao et al., 2017a). The results listed in Table. 1 suggest that when using the same pipeline and post-processing steps, the measurement results between different backbones do not have much difference. Among them, PSPNet, SERN, and DeepLabv3+ perform better, and U-Net is slightly inferior. Although DL-based methods achieved significant improvement, they still suffer the issues of low measuring accuracy and poor stability due to (a) the inability to incorporate anatomical priors, (b) the additional errors caused by the post-processing steps, which are with complex operations and hyper-parameters, and lacking the flexibility to remove outliers as our rectification module (see Sec. 4.6) does. The main prediction error of these DL-based methods comes from the complex post-processing performed on the predicted segmentation masks, which also requires manual setting of a series of complex operations and hyper-parameters.

In Fig. 7, we show the boxplots for APRIL and the comparing methods. Boxplot is a quartiles-based graphical tool for visualization of data distribution and to point out possible outliers. From this figure, we can intuitively find that our method has better accuracy and stability than the other three categories of methods, in which more errors and outliers are introduced. Besides, the pre-deep learning methods, especially snakes-based methods, have larger prediction

Table 2: The MAE results summary of ablation study on the key components of APRIL. *Rect* stands for uncertainty-guided rectification module.

	AP factors		Rect	CALD		CIMT	
	R_{ap}^{line}	R_{ap}^{angle}		in <i>pixel</i>	in <i>mm</i>	in <i>pixel</i>	in <i>mm</i>
APRIL				3.67 ± 2.65	0.22 ± 0.16	1.20 ± 0.87	0.07 ± 0.05
	✓			3.38 ± 2.44	0.20 ± 0.15	1.14 ± 0.82	0.07 ± 0.05
		✓		3.53 ± 2.52	0.21 ± 0.15	1.11 ± 0.83	0.07 ± 0.05
	✓	✓		3.19 ± 2.33	0.19 ± 0.14	1.07 ± 0.80	0.06 ± 0.05
	✓	✓	✓	3.02 ± 2.23	0.18 ± 0.13	0.96 ± 0.70	0.06 ± 0.04

errors and higher instability, due to more parameters that need to be manually
 470 tuned. Deep learning-based methods can improve the prediction results to some
 extent, but there is still a considerable gap compared with APRIL.

6.2. Ablation Study

We perform an ablation study on the APRIL model to explore and validate
 our design choice. Specifically, we evaluate the effectiveness of the key compo-
 475 nents in APRIL in Sec. 6.2.1, and validate the performances of different learning
 schemes in Sec. 6.2.2.

6.2.1. Ablation study on the key components of APRIL

We evaluate the influence of the key components in APRIL, including:

(1) **The effectiveness of AP factors:** We conduct evaluation for the AP
 480 factors from the following aspects: (a) Removing all the AP factors and the
 rectification module (a naive version of APRIL); (b) Adding only R_{ap}^{line} to (a),
 which considers the straightness constraint of the three edges; (c) Adding only
 R_{ap}^{angle} to (a), which considers the parallel relationship between LI and MA;
 (d) Adopting all the AP factors. The comparing results of these aspects are
 485 summarized in Table 2. Comparing with the naive version of APRIL, we can
 find that the R_{ap}^{line} has a positive effect on the locating of all three edges, and
 R_{ap}^{line} plays a more significant role when locating LI and MA related to CIMT.
 By adopting all the AP factors, the performance is further improved. We also

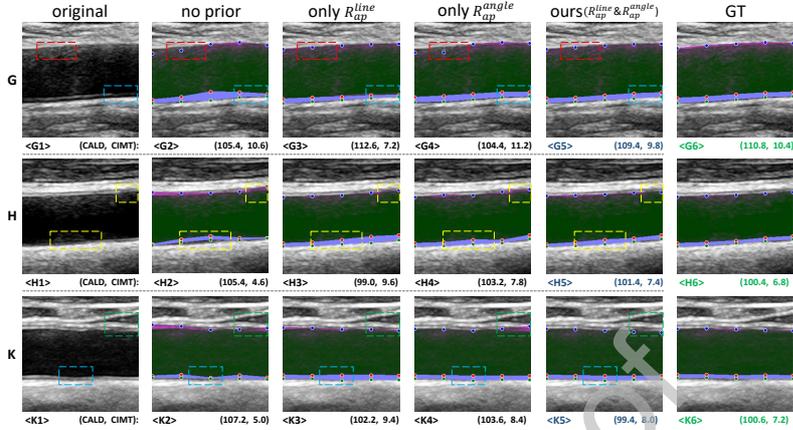


Figure 8: The visualization results of APRIL’s ablation models with/without different AP factors. We mark some challenging areas of each picture with colored dashed boxes as Fig. 1, Red: artifact and ambiguous areas in the edge area, Blue: tiny intima-media interface structure, Yellow: diseased or noisy areas, Green: irrelevant areas look similar to artery edges.

illustrate the visualization results of APRIL with/without different AP factors in Fig. 8. The results in the second column demonstrate that when without any
 490 AP factors, APRIL sometimes gives anatomical inaccurate predictions. The introduction of AP factors greatly eliminate this situation, which demonstrates the mutual benefits of the two AP factors.

(2) The impact of uncertainty-guided rectification: We further imple-
 495 ment APRIL with the uncertainty-guided rectification, and the accuracy of both CALD and CIMT are improved as reported in Table 2. The detailed rectification process is illustrated in Fig. 5. Specifically, in Fig. 6, frames *E* and *F* are under the state of vasomotion/pathology, and are selected for rectification by APRIL due to the higher predicted uncertainty. After such a rectification
 500 process, both frames give CALD&CIMT predictions with higher accuracy.

We conduct further evaluation of the hyper-parameter α_{unc} in the rectification module. Specifically, as introduced in Sec. 4.6, for the same patients’ frame sequence, we rank them according to the calculated $Dist_{cos}$ from high to low, and calculate the MAE of CALD & CIMT of all frames in every 10% interval. The

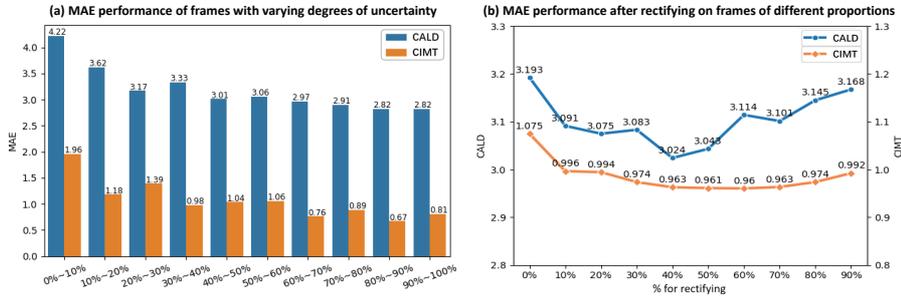


Figure 9: Schematic diagram of uncertainty guided rectification results, including: (a) MAE performance of frames with varying degrees of uncertainty in every 10% interval, and (b) MAE performance after uncertainty guided rectification on different proportions of frames.

505 results are shown in Fig. 9 (a). Then we pick out the frames with the degree of uncertainty ranked in the top α_{unc} from the same cases, and assign their CALD and CIMT to the average values of the remaining $(1 - \alpha_{unc})$ frames of the same patient (illustrated in Fig. 6 [E, F]). The results indicated in Fig. 9(b) demonstrate that when the α_{unc} is 40%, we got the best result.

510 6.2.2. Ablation study on different learning schemes

We also evaluate the performances of different learning schemes for the same task, including: (1) Changing the number of agents in APRIL from three to one, (2) Changing the prediction pipeline by directly estimating the position of all keypoints or directly predict the value of CALD/CIMT.

515 **(1) Single agent vs Three agents:** We implement the proposed APRIL in a single-agent way. Unlike the design indicated in Fig. 2, we use only a single agent to control all the 15 keypoints representing the edges of NW, LI, and MA. Besides, to further demonstrate the effectiveness of our AP factors, we implement APRIL(1 agent) both with & without AP. Under this setting, the
 520 action space becomes $15 * 2 + 1 = 31$. At each step, the model moves only one keypoint. In such a setting, the model controls more keypoints at the same time, and requires larger max steps per episode (set as 250 in our implementation),

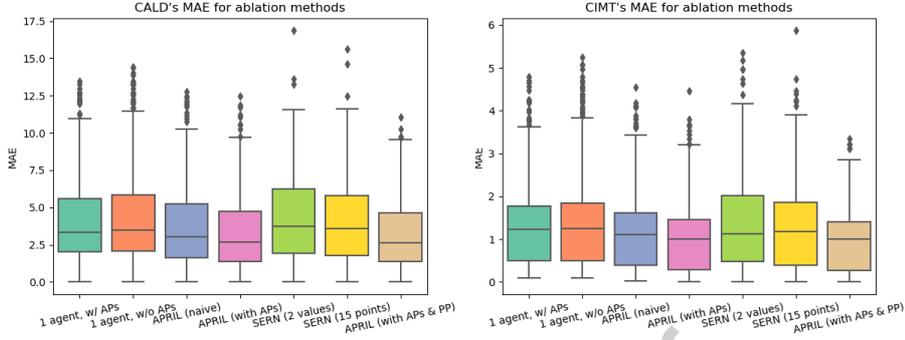


Figure 10: Boxplots of CALD & CIMT’s MAEs (in pixel) for ablation methods.

525 making the testing time of one frame more than 10s. The prediction results in Table. 3 show a non-negligible reduction when implementing APRIL in only one agent manner (CALD: 3.02 to 4.06, CIMT: 0.96 to 1.36), and discarding AP factors further degrade the results.

Table 3: MAEs and standard deviations of CALD and CIMT from APRIL, and the ablation models. *RECT* is short for uncertainty-guided rectification module.

schemes	description	CALD		CIMT	
		in <i>pixel</i>	in <i>mm</i>	in <i>pixel</i>	in <i>mm</i>
APRIL (1 agent, w/o AP)	only 1 agent to control all points (w/o AP factors)	4.39 ± 3.07	0.26 ± 0.18	1.44 ± 1.12	0.09 ± 0.07
APRIL (1 agent, w/ AP)	use only 1 agent to control all points (with AP factors)	4.06 ± 2.73	0.24 ± 0.16	1.36 ± 0.93	0.08 ± 0.06
SERN (15 points)	use SERN to directly regress to 15 keypoints	3.93 ± 2.57	0.24 ± 0.15	1.30 ± 0.96	0.08 ± 0.06
SERN (2 values)	use SERN to directly regress to 2 values	4.15 ± 2.68	0.25 ± 0.16	1.36 ± 1.11	0.08 ± 0.07
APRIL	the proposed APRIL with all AP factors and PP	3.02 ± 2.23	0.18 ± 0.13	0.96 ± 0.70	0.06 ± 0.04

(2) **Verifying the effect of direct estimation:** We adopt the same SE-ResNeXt as backbone encoder to directly regress to (a) the ordinance of 15 keypoints (SERN_(15 points)), (b) the final value of CALD and CIMT (SERN_{(2 val-}

ues)). These two methods have similar results, as indicated in Table. 3. However, direct prediction of CALD and CIMT values lacks intermediate results, so the accuracy of prediction cannot be verified conveniently in practical application. Comparing to APRIL, such approaches lack the limitation of prior knowledge, and are prone to produce anatomical-incorrect results. However, similar to our APRIL, such methods do not require complicated post-processing on the segmentation masks or edge detection results, that can avoid the additional errors and instability, compared to the approaches listed in Table. 1.

We further show boxplots of MAEs from different ablation models in Fig. 10. Basically, these models have better performance comparing with methods in Sec. 6.1, where the two direct regression methods perform slightly worse. In addition, the introduction of uncertainty-guided rectification can effectively reduce outliers, thereby reducing the overall error.

7. Conclusion

In this study, we formulate the CALD & CIMT estimation task as an RL problem and use DQN to solve it. In such setting, the estimation task is solved by learning three agents, where each agent controls five keypoints to locate the NW, LI, and MA edges, respectively. We dynamically incorporate the AP factors into the step-wise DQN learning procedure through an AP sensitive reward function, which guide the multiple agents' mutual movements. Via our innovative modeling, CALD and CIMT can be accurately measured by finding the exact locations of multiple keypoints with fewer annotations. We further optimize the prediction results by the uncertainty-guided rectification operation based on prediction variance on Q-lists. Experiments on a challenging CCA-US dataset prove the effectiveness and superiority of our method in terms of the measurement accuracy and stability. Our future work aims to incorporate the periodicity of the carotid vasomotion into the calculation process, and explore the anatomical relationships between adjacent frames.

Acknowledgements

This work is supported by the National Nature Science Foundation of China
 560 (No. 61876159, No. 61806172 & No. U1705286); the Fundamental Research
 Funds for the Central Universities, Xiamen University (No. 20720200030);
 the China Postdoctoral Science Foundation (No.2019M652257). The compu-
 tations are performed using the data analytical cloud at SHARCNET (<http://www.sharcnet.ca>) provided through the Southern Ontario Smart Computing
 565 Innovation Platform (SOSCIP).

References

- Alansary, A., Oktay, O., Li, Y., Le Folgoc, L., Hou, B., Vaillant, G., Kamnitsas, K., Vlontzos, A., Glocker, B., Kainz, B. et al. (2019). Evaluating reinforcement learning agents for anatomical landmark detection. *Medical*
 570 *image analysis*, 53, 156–164.
- Ali, I., Hart, G. R., Gunabushanam, G., Liang, Y., Muhammad, W., Nartowt, B., Kane, M., Ma, X., & Deng, J. (2018). Lung nodule detection via deep reinforcement learning. *Frontiers in oncology*, 8, 108.
- Barth, J. D. (2002). An update on carotid ultrasound measurement of intima-media thickness. *The American journal of cardiology*, 89, 32–38.
 575
- Bauer, M., Caviezel, S., Teynor, A., Erbel, R., Mahabadi, A. A., & Schmidt-Trucksäss, A. (2012). Carotid intima-media thickness as a biomarker of sub-clinical atherosclerosis. *Swiss medical weekly*, 142.
- Biswas, M., Kuppili, V., Araki, T., Edla, D. R., Godia, E. C., Saba, L., Suri,
 580 H. S., Omerzu, T., Laird, J. R., Khanna, N. N. et al. (2018). Deep learning strategy for accurate carotid intima-media thickness measurement: an ultrasound study on japanese diabetic cohort. *Computers in biology and medicine*, 98, 100–117.

- Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., & Adam, H. (2018). Encoder-
585 decoder with atrous separable convolution for semantic image segmentation.
In *Proceedings of the European conference on computer vision (ECCV)* (pp.
801–818).
- Cheng, D.-c., Schmidt-Trucksäss, A., Cheng, K.-s., & Burkhardt, H. (2002).
Using snakes to detect the intimal and adventitial layers of the common
590 carotid artery wall in sonographic images. *Computer methods and programs
in biomedicine*, *67*, 27–37.
- Delsanto, S., Molinari, F., Giustetto, P., Liboni, W., Badalamenti, S., & Suri,
J. S. (2007). Characterization of a completely user-independent algorithm for
carotid artery segmentation in 2-d ultrasound images. *IEEE Transactions on*
595 *Instrumentation and Measurement*, *56*, 1265–1274.
- Eaton-Rosen, Z., Bragman, F., Bisdas, S., Ourselin, S., & Cardoso, M. J. (2018).
Towards safe deep learning: accurately quantifying biomarker uncertainty in
neural network predictions. In *International Conference on Medical Image
Computing and Computer-Assisted Intervention* (pp. 691–699). Springer.
- 600 Faita, F., Gemignani, V., Bianchini, E., Giannarelli, C., Ghiadoni, L., & Demi,
M. (2008). Real-time measurement system for evaluation of the carotid
intima-media thickness with a robust edge operator. *Journal of ultrasound in
medicine*, *27*, 1353–1361.
- Ghesu, F.-C., Georgescu, B., Zheng, Y., Grbic, S., Maier, A., Hornegger, J., &
605 Comaniciu, D. (2017). Multi-scale deep reinforcement learning for real-time
3d-landmark detection in ct scans. *IEEE transactions on pattern analysis and
machine intelligence*, *41*, 176–189.
- Golemati, S., Stoitsis, J., Sifakis, E. G., Balkizas, T., & Nikita, K. S. (2007).
Using the hough transform to segment ultrasound images of longitudinal and
610 transverse sections of the carotid artery. *Ultrasound in medicine & biology*,
33, 1918–1932.

- Gutierrez, M. A., Pilon, P. E., Lage, S., Kopel, L., Carvalho, R., & Furuie, S. (2002). Automatic measurement of carotid diameter and wall thickness in ultrasound images. In *Computers in Cardiology* (pp. 359–362). IEEE.
- 615 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).
- Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4, 237–285.
- 620 Kass, M., Witkin, A., & Terzopoulos, D. (1988). Snakes: Active contour models. *International journal of computer vision*, 1, 321–331.
- Krebs, J., Mansi, T., Delingette, H., Zhang, L., Ghesu, F. C., Miao, S., Maier, A. K., Ayache, N., Liao, R., & Kamen, A. (2017). Robust non-rigid registration through agent-based action learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 344–352). Springer.
- 625 Liao, R., Miao, S., de Tournemire, P., Grbic, S., Kamen, A., Mansi, T., & Comaniciu, D. (2017). An artificial agent for robust image registration. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- 630 Liguori, C., Paolillo, A., & Pietrosanto, A. (2001). An automatic measurement system for the evaluation of carotid intima-media thickness. *IEEE Transactions on instrumentation and measurement*, 50, 1684–1691.
- Lobregt, S., & Viergever, M. A. (1995). A discrete dynamic contour model. *IEEE transactions on medical imaging*, 14, 12–24.
- 635 Loizou, C. P., Pattichis, C. S., Pantziaris, M., Tyllis, T., & Nicolaides, A. (2007). Snakes based segmentation of the common carotid artery intima media. *Medical & biological engineering & computing*, 45, 35–49.

- Lonni, E. M., Yusuf, S., Dzavik, V., Doris, C. I., Yi, Q., Smith, S., Moore-Cox, A., Bosch, J., Riley, W. A., & Teo, K. K. (2001). Effects of ramipril and vitamin e on atherosclerosis: the study to evaluate carotid ultrasound changes in patients treated with ramipril and vitamin e (secure). *Circulation*, *103*, 919–925.
- Luo, G., Dong, S., Wang, K., Zhang, D., Gao, Y., Chen, X., Zhang, H., & Li, S. (2019). A deep reinforcement learning framework for frame-by-frame plaque tracking on intravascular optical coherence tomography image. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 12–20). Springer.
- Ma, K., Wang, J., Singh, V., Tamersoy, B., Chang, Y.-J., Wimmer, A., & Chen, T. (2017). Multimodal image registration with deep context reinforcement learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 240–248). Springer.
- Maicas, G., Carneiro, G., Bradley, A. P., Nascimento, J. C., & Reid, I. (2017). Deep reinforcement learning for active breast lesion detection from dce-mri. In *International conference on medical image computing and computer-assisted intervention* (pp. 665–673). Springer.
- Man, Y., Huang, Y., Feng, J., Li, X., & Wu, F. (2019). Deep q learning driven ct pancreas segmentation with geometry-aware u-net. *IEEE transactions on medical imaging*, *38*, 1971–1980.
- Mathai, T. S., Gorantla, V., & Galeotti, J. (2019). Segmentation of vessels in ultra high frequency ultrasound sequences using contextual memory. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 173–181). Springer.
- Menchón-Lara, R.-M., Bastida-Jumilla, M.-C., Morales-Sánchez, J., & Sancho-Gómez, J.-L. (2014). Automatic detection of the intima-media thickness in ultrasound images of the common carotid artery using neural networks. *Medical & biological engineering & computing*, *52*, 169–181.

- Menchón-Lara, R.-M., & Sancho-Gómez, J.-L. (2015). Fully automatic segmentation of ultrasound common carotid artery images based on machine learning. *Neurocomputing*, *151*, 161–167.
- 670 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*, 529–533.
- O’Leary, D. H., & Bots, M. L. (2010). Imaging of atherosclerosis: carotid
675 intima–media thickness. *European Heart Journal*, *31*, 1682–1689.
- Petroudi, S., Loizou, C., Pantziaris, M., & Pattichis, C. (2012). Segmentation of the common carotid intima-media complex in ultrasound images using active contours. *IEEE transactions on biomedical engineering*, *59*, 3060–3069.
- Pignoli, P., & Longo, T. (1988). Evaluation of atherosclerosis with b-mode
680 ultrasound imaging. *The Journal of nuclear medicine and allied sciences*, *32*, 166–173.
- Postels, J., Ferroni, F., Coskun, H., Navab, N., & Tombari, F. (2019). Sampling-free epistemic uncertainty estimation using approximated variance propagation. In *Proceedings of the IEEE International Conference on Computer Vi-*
685 *sion* (pp. 2931–2940).
- Qian, C., & Yang, X. (2018). An integrated method for atherosclerotic carotid plaque segmentation in ultrasound image. *Computer methods and programs in biomedicine*, *153*, 19–32.
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks
690 for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention* (pp. 234–241). Springer.
- Saba, L., Biswas, M., Suri, H. S., Viskovic, K., Laird, J. R., Cuadrado-Godia, E., Nicolaides, A., Khanna, N., Viswanathan, V., & Suri, J. S. (2019).

- 695 Ultrasound-based carotid stenosis measurement and risk stratification in di-
 abetic cohort: a deep learning paradigm. *Cardiovascular diagnosis and ther-*
apy, *9*, 439.
- Sahba, F., Tizhoosh, H. R., & Salama, M. M. (2006). A reinforcement learning
 framework for medical image segmentation. In *The 2006 IEEE International*
Joint Conference on Neural Network Proceedings (pp. 511–517). IEEE.
- 700 Shin, J., Tajbakhsh, N., Todd Hurst, R., Kendall, C. B., & Liang, J. (2016).
 Automating carotid intima-media thickness video interpretation with convo-
 lutional neural networks. In *Proceedings of the IEEE Conference on Computer*
Vision and Pattern Recognition (pp. 2526–2535).
- Stein, J. H., Korcarz, C. E., Hurst, R. T., Lom, E., Kendall, C. B., Mohler,
 705 E. R., Najjar, S. S., Rembold, C. M., & Post, W. S. (2008). Use of carotid
 ultrasound to identify subclinical vascular disease and evaluate cardiovascular
 disease risk: a consensus statement from the american society of echocardi-
 ography carotid intima-media thickness task force endorsed by the society for
 vascular medicine. *Journal of the American Society of Echocardiography*, *21*,
 710 93–111.
- Stein, J. H., Korcarz, C. E., Mays, M. E., Douglas, P. S., Palta, M., Zhang, H.,
 LeCaire, T., Paine, D., Gustafson, D., & Fan, L. (2005). A semiautomated
 ultrasound border detection program that facilitates clinical measurement of
 ultrasound carotid intima-media thickness. *Journal of the American Society*
 715 *of Echocardiography*, *18*, 244–251.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*.
 MIT press.
- Tell, G. S., Polak, J. F., Ward, B. J., Kittner, S. J., Savage, P. J., & Robbins, J.
 (1994). Relation of smoking with carotid artery wall thickness and stenosis in
 720 older adults. the cardiovascular health study. the cardiovascular health study
 (chs) collaborative research group. *Circulation*, *90*, 2905–2908.

- Van Hasselt, H., Guez, A., & Silver, D. (2016). Deep reinforcement learning with double q-learning. In *Thirtieth AAAI conference on artificial intelligence*.
- Vila, M., Remeseiro, B., Grau, M., Elosua, R., Betriu, À., Fernandez-Giraldez, E., & Igual, L. (2020). Semantic segmentation with densenets for carotid artery ultrasound plaque segmentation and cimt estimation. *Artificial Intelligence in Medicine*, 103, 101784.
- Vlontzos, A., Alansary, A., Kamnitsas, K., Rueckert, D., & Kainz, B. (2019). Multiple landmark detection using multi-agent reinforcement learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 262–270). Springer.
- Wang, L., Merrifield, R., & Yang, G.-Z. (2011). Reinforcement learning for context aware segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 627–634). Springer.
- Wang, Z., Schaul, T., Hessel, M., Hasselt, H., Lanctot, M., & Freitas, N. (2016). Dueling network architectures for deep reinforcement learning. In *International Conference on Machine Learning* (pp. 1995–2003).
- Watkins, C. J., & Dayan, P. (1992). Q-learning. *Machine learning*, 8, 279–292.
- Williams, D. J., & Shah, M. (1992). A fast algorithm for active contours and curvature estimation. *CVGIP: Image understanding*, 55, 14–26.
- Xie, M., Li, Y., Xue, Y., Shafritz, R., Rahimi, S. A., Ady, J. W., & Roshan, U. W. (2019). Vessel lumen segmentation in internal carotid artery ultrasounds with deep convolutional neural networks. In *2019 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)* (pp. 2393–2398). IEEE.
- Xie, S., Girshick, R., Dollár, P., Tu, Z., & He, K. (2017). Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 1492–1500).

- 750 Zhao, C., Feng, C., Li, D., & Li, S. (2020). Of-msrn: Optical flow-auxiliary multi-task regression network for direct quantitative measurement, segmentation and motion estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence* (pp. 1218–1225). volume 34.
- 755 Zhao, H., Shi, J., Qi, X., Wang, X., & Jia, J. (2017a). Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2881–2890).
- Zhao, S., Gao, Z., Zhang, H., Xie, Y., Luo, J., Ghista, D., Wei, Z., Bi, X., Xiong, H., Xu, C. et al. (2017b). Robust segmentation of intima-media borders with different morphologies and dynamics during the cardiac cycle. *IEEE journal of biomedical and health informatics*, *22*, 1571–1582.
- 760 Zheng, Z., & Yang, Y. (2020). Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. *arXiv preprint arXiv:2003.03773*, .
- 765 Zhou, Z., Shin, J., Feng, R., Hurst, R. T., Kendall, C. B., & Liang, J. (2019). Integrating active learning and transfer learning for carotid intima-media thickness video interpretation. *Journal of digital imaging*, *32*, 290–299.