



Delft University of Technology

Deep reinforcement learning for cerebral anterior vessel tree extraction from 3D CTA images

Su, Jiahang; Li, Shuai; Wolff, Lennard; van Zwam, Wim; Niessen, Wiro J.; van der Lugt, Aad; van Walsum, Theo

DOI

[10.1016/j.media.2022.102724](https://doi.org/10.1016/j.media.2022.102724)

Publication date

2023

Document Version

Final published version

Published in

Medical Image Analysis

Citation (APA)

Su, J., Li, S., Wolff, L., van Zwam, W., Niessen, W. J., van der Lugt, A., & van Walsum, T. (2023). Deep reinforcement learning for cerebral anterior vessel tree extraction from 3D CTA images. *Medical Image Analysis*, 84, Article 102724. <https://doi.org/10.1016/j.media.2022.102724>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Deep reinforcement learning for cerebral anterior vessel tree extraction from 3D CTA images

Jiahang Su ^{a,*}, Shuai Li ^a, Lennard Wolff ^a, Wim van Zwam ^b, Wiro J. Niessen ^{a,c}, Aad van der Lugt ^a, Theo van Walsum ^a

^a Department of Radiology & Nuclear Medicine, Erasmus MC, University Medical Center Rotterdam, The Netherlands

^b Department of Radiology & Nuclear Medicine, Maastricht UMC, Cardiovascular Research Institute Maastricht, The Netherlands

^c Faculty of Applied Sciences, Delft University of Technology, The Netherlands

ARTICLE INFO

MSC:

41A05

41A10

65D05

65D17

Keywords:

Deep reinforcement learning

Tracking

Brain vessel

Bifurcation detection

3D CTA

CNN

ABSTRACT

Extracting the cerebral anterior vessel tree of patients with an intracranial large vessel occlusion (LVO) is relevant to investigate potential biomarkers that can contribute to treatment decision making. The purpose of our work is to develop a method that can achieve this from routinely acquired computed tomography angiography (CTA) and computed tomography perfusion (CTP) images.

To this end, we regard the anterior vessel tree as a set of bifurcations and connected centerlines. The method consists of a proximal policy optimization (PPO) based deep reinforcement learning (DRL) approach for tracking centerlines, a convolutional neural network based bifurcation detector, and a breadth-first vessel tree construction approach taking the tracking and bifurcation detection results as input. We experimentally determine the added values of various components of the tracker. Both DRL vessel tracking and CNN bifurcation detection were assessed in a cross validation experiment using 115 subjects. The anterior vessel tree formation was evaluated on an independent test set of 25 subjects, and compared to interobserver variation on a small subset of images.

The DRL tracking result achieves a median overlapping rate until the first error (1.8 mm off the reference standard) of 100, [46, 100] % on 8032 vessels over 115 subjects. The bifurcation detector reaches an average recall and precision of 76% and 87% respectively during the vessel tree formation process. The final vessel tree formation achieves a median recall of 68% and precision of 70%, which is in line with the interobserver agreement.

1. Introduction

1.1. Clinical background

The cerebral vessel network is a complex network that feeds the brain tissue. Diseases affecting the cerebral vessel network, such as an intracranial ischemic stroke caused by a large vessel occlusion, may have severe consequences. In patients with an intracranial large vessel occlusion (LVO), the occlusion often occurs in the anterior circulation, which consists of the middle cerebral artery (MCA) vessel tree and anterior cerebral artery (ACA) vessel tree. Knowing the precise anterior vessel tree with its topology could provide additional information for the treatment planning of LVO and potentially contribute to the treatment decision making process. The complete cerebral vessel centerline annotation with anterior vessel trees is shown in Fig. 1.

In clinical practice, 3D modalities such as magnetic resonance angiography (MRA) and computed tomography angiography (CTA) are

commonly used to visualize cerebral vasculature. In this manuscript, we are focusing on CT-based imaging, which is a common modality for the workup of stroke patients. CT imaging techniques for cerebral vessels in clinical practice are CTA (both single phase and multiphase) and computed tomography perfusion (CTP) (Potter et al., 2019). CTP is an imaging protocol where a series of 3D CTA images is acquired.

The purpose of our work is to extract the anterior vessel trees as distal as possible given the initial direction vectors of each tree.

1.2. Related work

Traditionally, vessel tree extraction starts with obtaining a vessel segmentation or geometric model (Damseh et al., 2020). Extensive reviews on vessel segmentation and vessel shape extraction methods have been presented by Kirbas and Quek (2004), Lesage et al. (2009),

* Corresponding author.

E-mail address: j.su@erasmusmc.nl (J. Su).

<https://doi.org/10.1016/j.media.2022.102724>

Received 17 June 2022; Received in revised form 24 November 2022; Accepted 2 December 2022

Available online 9 December 2022

1361-8415/© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Nomenclature

2D	Two dimensional
3D	Three dimensional
A1	Proximal segment of Anterior Cerebral Artery tree
A2C	Advantage Actor Critic
ACA	Anterior Cerebral Artery
Bf	Bifurcation
CCS	Curve to Curve Similarity
CNN	Convolutional Neural Network
CT	Computed Tomography
CTA	Computed Tomography Angiography
CTP	Computed Tomography Perfusion
CTT	Correct Tree Topology
DQN	Deep Q-network
DRL	Deep Reinforcement Learning
GAE	Generalized Advantage Estimation
HU	Hounsfield Units
IQR	Interquartile Range
JS	Jiahang Su
LSTM	Long Short-Term Memory
LVO	Large Vessel Occlusion
M	Medical student
M1	Proximal segment of Middle Cerebral Artery tree
MCA	Middle Cerebral Artery
MDP	Markov Decision Process
MRA	Magnetic Resonance Angiography
OR	Overlap Rate
PPO	Proximal Policy Optimization
PReLU	Parametric Rectified Linear Unit
RCNN	Region-based Convolutional Neural Networks
ReLU	Rectified Linear Unit
RNN	Recurrent Neural Network
SLO	Scanning Laser Ophthalmoscopy
TRPO	Trust Region Policy Optimization

and Moccia et al. (2018). After obtaining a vessel segmentation from the input image, various methods can be used to obtain the vessel tree. Lee et al. (1994) used a thinning algorithm, Moriconi et al. (2018) used geodesics in combination with a minimum spanning tree, Robben et al. (2016) used a graph-based tracking approach to obtain the Circle of Willis and the proximal segment of each arterial trees from 3D MRA images and Zhang et al. (2021) utilize the general confluence constraint with minimal arborescence on a directed graph. Others build the tree in a more explicit way from a set of bifurcations and connected centerlines, an approach followed by Zhang et al. (2020) and Jeon (2021). In these approaches, vessel tracking and bifurcation detection are two essential ingredients.

1.2.1. Tracking

In conventional methods, vessel segmentation approaches start with centerline tracking. Many tracking methods have been described in the previously mentioned reviews. Nowadays, deep learning based methods have replaced conventional approaches for many image processing problems. Also, combinations of convolutional neural networks (CNNs) and conventional approaches have become popular. For example, Wolterink et al. (2019) extract the coronary artery centerline with an iterative tracking approach utilizing a 3D CNN to estimate the

orientation and radius from cardiac CT angiography, Guo et al. (2019) use a multi-task CNN method to obtain an estimate of a centerline distance map and an endpoint confidence map from coronary CTA images and subsequently a minimal path method is used to compute the centerline.

Reinforcement learning is a branch of machine learning algorithms that have been studied in domains such as gaming and control. Deep reinforcement learning (DRL), which is the fusion of reinforcement learning and deep learning, has been shown to be able to solve complex image processing challenges by exploiting image features (Arulku-maran et al., 2017). Such DRL based methods are gradually entering the medical imaging field (Zhou et al., 2021). Tracking is a preferred application direction for DRL approaches, as DRL learns the optimal policy by maximizing the accumulated reward over time, which maps nicely to the problem of finding a path through an nD image. The optimal policy determines the best actions for an agent in each state over time.

There are two main streams in DRL methods, value-based and policy gradient based. In value-based methods e.g. deep Q network (Mnih et al., 2013) the policy π is directly obtained by approximating the state-action value $Q_{\theta}(s, a)$ as a function of state s and action a or state value $V_{\theta}(s)$ as a function of a state s . Here, θ represents the network parameters. For example, Zhang et al. (2018) apply a deep Q network (Mnih et al., 2013) (DQN) approach with a discrete action space given by a six-connected neighborhood as action space and a point-to-curve reward function for thoracic aorta centerline tracking in 3D contrast-enhanced and none contrast-enhanced CT images. Zhang et al. (2020) similarly utilizes a double-DQN (Van Hasselt et al., 2016) with a discrete 26-connected neighborhood as the action space to track the coronary centerline. In this work, the bifurcations were detected using a standalone detector. Li et al. (2021) utilize the DQN approach with predefined orientation as discrete action space to track coronary centerlines. In addition, the value-based methods are known to often lead to an overestimation of the action value (Van Hasselt et al., 2016).

In policy gradient methods, a parameterized policy is learned directly using a stochastic optimization algorithm (Sutton and Barto, 2018). Trust region policy optimization (TRPO) (Schulman et al., 2015a) and proximal policy optimization (PPO) (Schulman et al., 2017) are two well-known methods for such optimization. Dai et al. (2019) utilize PPO for learning the policy for neuronal tracking in 2D neuronal microscopy images. The above-mentioned DRL-based tracking approach has demonstrated superior performance in comparison with a benchmark using CNN based methods.

1.2.2. Bifurcation detection

Bifurcation detection plays an important role in determining the vessel tree. Some feature based machine learning methods have been presented in the past. For example, Baboiu and Hamarneh (2012) utilized the scale-space features of vessel structures to construct a bifurcation detector on 2D and 3D synthetic images. Kalaie and Gooya (2017) used Gaussian profiles in cross sections of bifurcation points to discriminate vessels from bifurcation points in 2D retinal images. Jeon (2021) performs bifurcation detection using intensity based spatial clustering methods for 3D coronary arteries extraction from CTA images. Recently, deep learning methods have been presented by Zheng et al. (2015), who utilized a two-stage CNN approach to detect carotid artery bifurcations from 3D CTA images, and Hervella et al. (2020), who used a multi-instance heat map with a U-net architecture Ronneberger et al. (2015) to estimate vessel crossings and bifurcation points from 2D retina images. The latter application was also addressed by Zhao et al. (2020) utilizing a two-stage refinement approach, in which an RCNN (Girshick et al., 2014) based network first obtained a rough estimation of bifurcations and in a 2nd stage, the same network architecture was applied to address the errors from the first network.

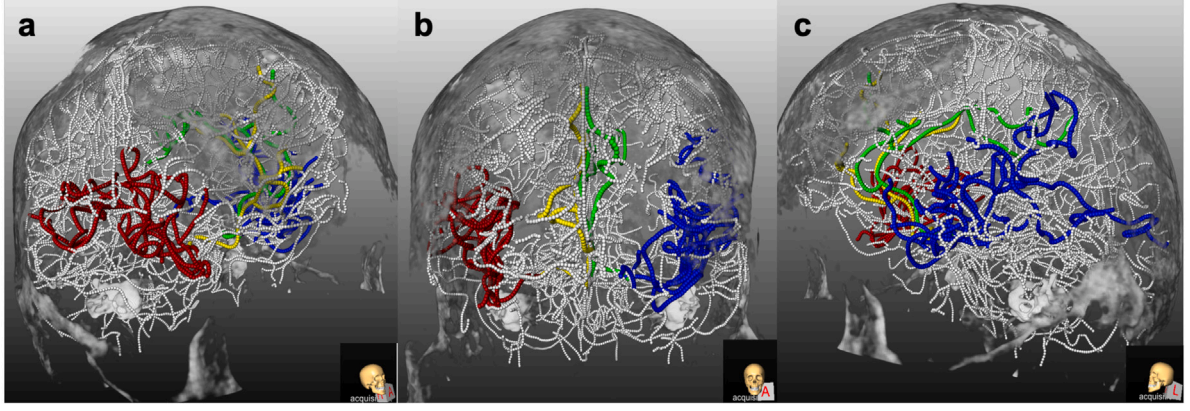


Fig. 1. Example of brain vessel annotation with colored anterior vessel trees for a subject with a large vessel occlusion in the left M2 segment. White represents veins and posterior cerebral artery trees, red is the right MCA tree, yellow is the right ACA tree, green is the left ACA tree, blue is the left MCA tree, a: the right sagittal view; b: the coronal view; c: the left sagittal view.

1.3. Contribution and organization of our work

Our aim is to extract the anterior vessel tree from 3D CT images, independent of the acquisition protocol. For such a task, to the best of our knowledge, most vessel tree extraction methods and the subsequent tracking and bifurcation steps have focused on coronary artery tree extraction from 3D cardiac CTA images and vessel tree extraction from 2D retina scanning laser ophthalmoscope (SLO) images. None of the existing works presents a method to extract the cerebral anterior vessel tree from 3D CTA images or methods that perform tree extraction beyond the proximal anterior vessel tree.

Our work has several contributions. First, we assess a DRL approach that, to the best of our knowledge, has not been applied yet for cerebral vessel tracking approaches. Second, the DRL method was adapted to this application. More specifically, we introduce a novel curve-to-curve distance based reward function and a network architecture (CNN with recurrent neural network (RNN)) for cerebral vessel tracking, and we assess the impact of various features of the methods in an ablation study. In addition, we performed a comparison with an existing baseline method. Finally, all experiments have been performed with a large set of clinical data.

The remainder of this manuscript is organized as follows. In Section 2, we first present a DRL PPO based directed vessel centerline tracking model, followed by a CNN-based bifurcation detection, and breadth-first tree formation methods. The data set, data annotations, and preprocessing are described in Section 3. The implementation, hyperparameter optimizations and experimental results can be found in Section 4, followed by discussion in Section 5 and conclusions in Section 6.

2. Method

The proposed cerebral anterior vessel tree extraction method starts with a DRL centerline tracking approach starting from the root of the tree (with an initial direction vector). Along the tracked path (agent path), a CNN based bifurcation detector is used to identify the bifurcation points. New tracks are generated from the bifurcation points detected along the tracks. In the tree formation, stopping criteria are applied to remove the spurious parts of the tracked paths, and a breadth-first approach is used to construct the tree from the tracked paths and the bifurcation points.

2.1. Directed vessel tracking

A reinforcement learning problem can be modeled as a Markov Decision Process (MDP) (Bertsekas et al., 2011). We, therefore, introduce the key ingredients of our DRL tracker as the components of an MDP. The environment in our case is the normalized 3D CTA image. The other components, such as state and action, are introduced below.

2.1.1. State and action

The agent at time t has an associated state, denoted as s_t . Next to the current position in the image, the agent state contains image information from previous steps and the corresponding displacement vector array. Let p be any discrete 3D position in a CTA volume. The image information in the state consists of three $21 \times 21 \times 21$ sub-volumes centered at respectively the agent position at p_t and two previous agent positions p_{t-1} , p_{t-2} . This size is sufficient to capture the local neighborhood, even for large vessels and bifurcations. In addition, the state contains a sub-volume of the binary agent path centered at p_t , which represents the voxels the agent went through. As a vessel can be viewed as a tubular structure with varying radius, any independent point in the track through a vessel can be tracked in two directions. Using the image information from the three previous steps and the binary agent path can help the agent track the vessel in a directed way. The dynamic displacement vector array at time t consists of the directional vectors (steps of the agent) from the current position till the initial position $[a_t, a_{t-1}, \dots, a_0]$. The dynamic displacement vector provides additional directional information on previous steps.

The agent discrete action space is defined as the 26 connected neighboring voxels of the current position p_t in an agent path.

We represent the set of actions by scaled displacement vectors as follows:

$$\Delta p_t = \alpha \times \{x, y, z \mid x, y, z \in \{-1, 0, 1\}\}, \quad (1)$$

As a consequence, an agent action a_t can make a step of at least two voxels per episode step if $\alpha = 2$.

2.1.2. Reward

The reward function is an essential component for the network to achieve the optimal policy. The ultimate goal of the vessel tracking training task is to learn the agent tracking the vessel centerline by maximizing R , the accumulated discounted instant reward r_t :

$$R = \sum_{i=0}^{T-1} \gamma^i r_{t+i} \quad (2)$$

The instant reward r_t is used to measure the policy performance of the agent at time t , and thus is relevant for policy optimization. We would like the agent to follow the reference path G_t , which in our case is a voxelized representation of the annotated vessel track.

We, therefore, use the difference of the curve-to-curve distance between the agent path and reference standard between time t and $t-1$ to measure the reward of the corresponding action. The curve-to-curve distance was introduced in Walsum et al. (2008) to quantify the distance between coronary centerlines. The curve-to-curve distance is computed by approximating the surface area spanned between two




Fig. 2. Example of curve-to-curve distance and CCS metric, the red dotted line is the reference standard and the blue dotted line is the actual track, the summation of lengths of the red lines between the reference standard and track is the computed surface distance; a: the anterior vessel tree with highlighted color-coded vessel centerline (corresponding to the colored frame). b: example of L_0 . c: example of success tracked case, CCS = 0.87, OR = 1. d: example track at time t , CCS = 0.58, OR = 0.5. e: example of miss tracked case, CCS = 0.58, OR = 0.09.

curves via the summation of the Euclidean length of all corresponding point-to-point connections. The optimal (minimal area) correspondence between two curves was determined via a Dijkstra minimum cost algorithm (Dijkstra et al., 1959). Fig. 2 shows examples of surface area computation (the summation of red lines) in various scenarios.

In addition, we would like to enforce the agent to follow the reference standard G_t . Therefore, a binary overlap metric, which measures the overlap between the reference path G_t and the agent path, is part of the instant reward.

We propose an instant reward function that combines binary overlap and the curve-to-curve distance. The binary overlap metric provides extra positive (+1) feedback when the agent path and ground truth are overlapping at position p_t . Denoting the surface distance in position p_t as L_t , and the binary overlap at position p_t as B_t , the proposed instant reward is defined as follows:

$$r_t = \begin{cases} B_t - \log(\epsilon + |L_t - L_{t-1}|), & \text{if } L_t - L_{t-1} < 0; \\ B_t + \log(\epsilon + |L_t - L_{t-1}|), & \text{otherwise.} \end{cases} \quad (3)$$

2.1.3. Policy

In our vessel tracking application, the optimal policy is obtained using the PPO method with the advantage actor critic (A2C) (Mnih et al., 2016) framework. The A2C architecture consists of two networks: an actor network and a critic network. The critic network $V_v(s)$ estimates the state value and actor network $Q_w(s, a)$ learns the state-action value suggested by the critic. The advantage function $A(s_t, a_t)$ (Eq. (4)) is introduced to measure the state action pair at time t :

$$A(s_t, a_t) = Q_w(s_t, a_t) - V_v(s_t) \quad (4)$$

It determines how much better a selected action is compared to the expected value of all possible actions. In this application, generalized advantage estimation (GAE) (Schulman et al., 2015b) is applied to further reduce the variance of the temporal difference error of the advantage function.

In policy gradient DRL methods, stochastic optimization algorithms such as gradient ascent are sensitive to the gradient update step size. To accommodate for the step size issue, trust region policy optimization (TRPO) (Schulman et al., 2015a) was introduced using KL-divergence (Kullback and Leibler, 1951) constraints to stabilize the gradient update step between old and new policies. However, the use of KL-divergence yields costly computations on multiple Hessian-vector products (Engstrom et al., 2020). The PPO method (Schulman et al., 2017) which was later developed to approximate the KL-divergence

constraints by a simple but efficient regularization mechanism in objective function as follows:

$$J(\theta) = \mathbb{E}[\min(\rho(t)A(s_t, a_t), \text{clip}(\rho(t), 1 - \epsilon, 1 + \epsilon)A(s_t, a_t))] \quad (5)$$

where $\rho(t)$ is defined as the likelihood ratio between policies at two time points:

$$\rho(t) = \frac{\pi_t(s_t, a_t)}{\pi_{t-1}(s_{t-1}, a_{t-1})} \quad (6)$$

2.1.4. Network architecture

The proposed network architecture is shown in Fig. 3. The architecture consists of an actor network and a critic network. The critic network helps the agent (actor) with learning the optimal state-action value during the training process. Assuming that the policy is optimal after training, only actor network is required for obtaining the tracks in the inference stage.

The agent state consists of sub-volumes sampled from three consecutive points along the agent path and one sub-volume containing the binary agent path. In the CNN part, the three sub-volumes generally cannot cover the whole path of the agent. A recurrent neural network (RNN) (Hochreiter and Schmidhuber, 1997) was therefore used to provide extra temporal information along the agent path.

The architecture of the actor and critic networks are similar, they both consist of a CNN and an RNN. The CNN model is similar to the architecture proposed by Wolterink et al. (2019) but with a parametric rectified linear activation function (PReLU) (He et al., 2015) instead of a rectified linear activation function (ReLU) (Maas et al., 2013). The proposed CNN architecture is based on a dilated network (Yu and Koltun, 2015) architecture. Each CNN block consists of a series of dilated convolutions, with dilation kernels of size 1. For each convolutional layer, instance normalization was applied. The detailed CNN architecture of the actor and critic can be found in Table 1. The critic CNN takes a ternary ground truth sub-volume centered at p_t as additional input. The dimension of this sub-volume is $21 \times 21 \times 21$ voxels, and each voxel has a value of -1, 0, or 1: the paired ground truth centerline for the current episode has value 1, other centerlines are labeled as -1, and the remaining voxels have value 0. This setup ensures a decrease in instant reward if the agent moves to the wrong vessel.

The RNN in both the actor and critic network consists of two LSTM cells. The input of the RNN in the actor network is the action history vector; the input vector x of the RNN in the critic network at time

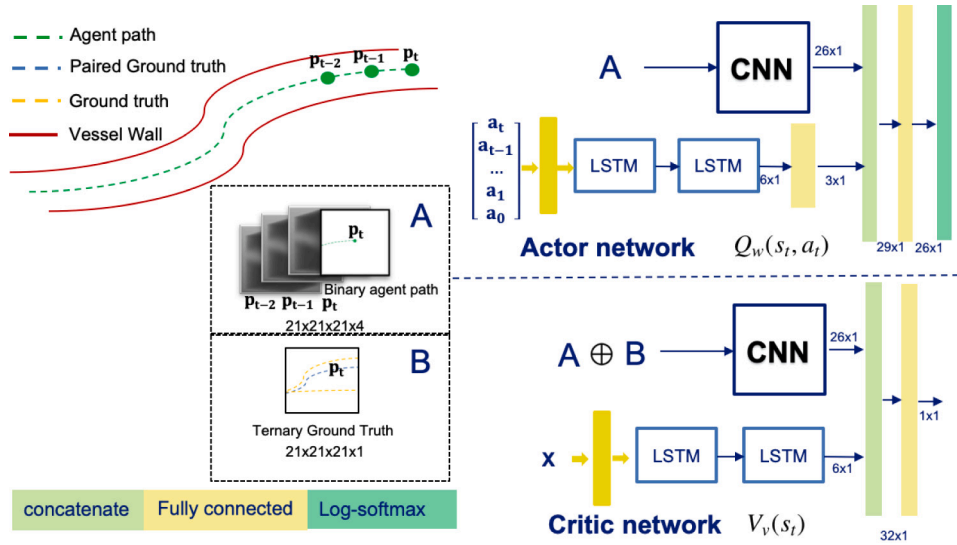


Fig. 3. The architecture of the advantage actor critic (A2C) network. Both actor and critic networks consist of CNN and an RNN. The RNN part consists of two LSTM cells. The input of the actor CNN (A) consists of three CTA cubes and one binary agent path cube ($21 \times 21 \times 21$ voxels). The three CTA cubes are centered at the current position p_t and two previous positions p_{t-1} , p_{t-2} . The binary agent path cube is centered in the current position p_t . The CNN input of the critic network ($A \oplus B$) consists of A and a ternary ground truth cube (B) that is centered at the current position p_t . The RNN input of the critic network is a vector array (x) that consists of paired local features along the track. The RNN input of the actor network is a vector $[a_t, a_{t-1}, \dots, a_1, a_0]$ that contains the action history.

Table 1

The network architecture of actor/critic CNN. The actor's input is $21 \times 21 \times 21 \times 4$ and the critic's input is $21 \times 21 \times 21 \times 5$. Each convolutional layer is followed by an instance normalization and PReLU.

Layer	1	2	3	4	5	6	7
Kernel width	3	3	3	3	1	1	7
Dilation rate	1	1	2	3	4	1	1
Channels	32	32	32	32	64	64	26

point t consists of the paired instant rewards (r_t), the log of instant reward ($\log(r_t)$), the intensity value of agent path (I_t) and the mean intensity value of a two voxel wide region around the agent path (\bar{I}_t):

$$x = \begin{pmatrix} r_t & \log(r_t) & I_t & \bar{I}_t \\ r_{t-1} & \log(r_{t-1}) & I_{t-1} & \bar{I}_{t-1} \\ \vdots & \vdots & \vdots & \vdots \\ r_0 & \log(r_0) & I_0 & \bar{I}_0 \end{pmatrix} \quad (7)$$

The outputs of the CNN and RNN are concatenated and fed into a fully connected layer. The final output of the actor network is a (log) softmax layer, which determines the direction of the agent in the next state. The final output of the critic network is a scalar value, which approximates the value of the input state.

The final fully connected layer in the actor network is absent in the critic network. As the latter network needs to learn the state value, the final fully connected layer is not necessary.

2.2. Bifurcation detection

To be able to build a vascular tree, a bifurcation detector was developed that can be applied to each path that was generated by the tracker; from bifurcations, new paths can be tracked.

We approach bifurcation detection as a binary classification problem. The structure of the model used for bifurcation detection is similar to the dilated network architecture in the CNN blocks of the A2C networks, but with only one image cube p_t as input. Each convolutional layer is followed by switchable normalization (Luo et al., 2018) and a dropout rate of 0.2; we use binary cross-entropy as the loss function. The detailed architecture of the bifurcation detector can be found in Table 2.

2.3. Tree formation

The anterior vessel tree can be viewed as an arborescence with a given root. We apply a breadth-first tree formation for the anterior vessel tree extraction. The tracking agent and bifurcation detector described in the previous paragraph can be applied to obtain a series of tracks and bifurcation points. Starting from the root, we obtain the first track. Subsequently, we prune the tracked path based on stopping criteria, and recursively and in a breadth-first manner, find bifurcations along this pruned tracked path, and start new tracks from these bifurcations. In this latter step, new paths are tracked from the bifurcation point in several directions, and then failed or overlapping tracks are removed.

In an attempt to get a more robust approach, increasing the true positive tracks, while reducing the missed ones and false positives, we also introduce an ensemble method with five different vessel tracking models.

2.3.1. Stopping criteria

Stopping criteria are used to determine the termination point of a track; i.e. the track will be pruned if one of the following conditions is met:

- the track runs into the brain tissue
- the track runs into the dilated skull,
- the length of the track exceeds a maximal length.
- the track reaches the border of images in z-axis (for half brain coverage cases).
- the track runs into a different arterial territory. e.g. forming ACA tree but running to MCA territory.

Fig. 4b shows the example of the relationship between the anterior vessel trees and the arterial territory map.

Table 2

The network architecture of bifurcation CNN. The input size is $21 \times 21 \times 21$. Each convolutional layer is followed by a switchable normalization, a dropout rate of 0.2, and a PReLU activation function.

Layer	1	2	3	4	5	6	7
Kernel width	3	3	3	3	1	1	7
Dilation rate	1	1	2	3	4	1	1
Channels	32	32	32	32	64	64	26




Fig. 4. Processing of the anterior vessel trees data preparation; a. shows the example of the anterior vessel tree structure, the MCA tree is presented in blue and ACA trees are presented in green; b. shows the anterior vessel trees and thresholded (0.5) MCA and ACA mask, the green mask is the MCA mask, the red mask is the ACA mask; c. shows the segment labels. The proximal segments are in red, the middle segments are in blue, the distal segments are in green and the skull segments are in yellow.

2.3.2. Bifurcation inference

Application of the bifurcation detection model (Section 2.2) for all the voxels of a tracked path yields a probability of a bifurcation being present at each voxel of the track. The detector may already give some probability to voxels that are close to a bifurcation, The probabilities in voxels from near a bifurcation to the bifurcation increase, and have a peak at the bifurcation. We, therefore, extract the bifurcation points by taking the max probability voxel from a consecutive series of voxels that consists of at least two voxels with a probability larger than 0.5. The threshold allows for smaller bifurcations to be detected.

2.3.3. Bifurcation path tracking

After bifurcation detection, new paths must be tracked from the bifurcation. This may be done using an estimate of the candidate directions based on the vessel features. We take a two-step approach. In the first step, we start a tracker in all candidate directions regardless of the presence of vessels. Then, we remove unwanted tracks that do not satisfy the eligibility criteria. The candidate directions are determined from 26 connected neighborhoods. In order to prevent the tracked paths from following the original track backward, tracks are only started in the forward direction with regard to the current position (i.e. if the inner product of the current path and direction vector of the new path is positive). After tracking these candidate directions, spurious tracks will be removed based on the amount of overlap with existing tracks and the length of the track. If the track has an overlap with a previously tracked path that is larger than or equal to 95 %, the tracked path will be ignored. Short tracked paths (less than 4.5 mm) will also be ignored.

2.3.4. Ensemble method

In addition to tracking with one model, the combined policy and models ensemble method aims to further improve the performance by combining results from multiple tracking models in two different ways. The first track of each tree is essential for tree formation. To ensure a good initial track, we use majority voting for the first track: at each tracking step, we initiate five agents to generate the candidate actions and select the action (i.e. next voxel) that is closest to the average of the actions of each of the models. After the first track, the output of all models is used to reduce missed vessels in bifurcations. The trackers in the ensemble method use the same stopping criteria as the other trackers (see Section 2.3.1).

3. Data

3.1. Data overview

In our study, we used CTA and CTP image data from two sources: MR CLEAN registry (Jansen et al., 2018) and Erasmus MC. The MR CLEAN registry is an ongoing multi-center registry for stroke patients that underwent endovascular treatment for LVO in The Netherlands since March 2014. Seventeen centers are involved in this study. Data selection criteria applied when selecting the images were: (i) slice thickness ≤ 1.5 mm; (ii) slice spacing ≤ 1.5 mm; (iii) for CTA, the contrast acquisition phase has to be peak arterial, equilibrium or early venous phase according to the definition of Rodriguez-Luna et al. (2014); (iv) brain coverage has to be at least half of the brain, the insula region needs to be present in both hemispheres; (v) no large motion artifact along the time axis for CTP data; (vi) CTP data has to contain full cardiac cycle from early arterial phase to late venous phase.

1594 subjects were included in the MR CLEAN registry from March 2014 to June 2016. In a previous study (Su et al., 2020), we selected 270 images from this set that matched the inclusion criteria. Of these, 49 were manually selected such that there is variation in vendor, image quality, and acquisition phase. In our study, we used these 49 images (with annotations), and we randomly selected another 26 from the 270 images.

From the 63 registry subjects with CTP images matching our selection criteria, we randomly selected 35 subjects. In addition, from 58 stroke patients with an LVO that were admitted from Jan. 2018 to March 2020 at the Erasmus MC, we randomly selected 15 subjects from the 58 subjects (out of a total of 335) that matched our selection criteria.

More detailed information on the data selected is presented in Table 3.

3.2. Data annotation

The tracks of the anterior vessel tree for training the DRL were manually annotated by the first author of this manuscript (JS) for all 125 subjects. This annotation was done under the supervision of an experienced radiologist (20 years of experience) and a physician (5 years of experience). The annotation task was defined as labeling the track of every anterior vessel from the ICA-top to the most distal part

Table 3

The data distribution of 125 subjects.

Properties		Data division (n = 125)						
		Set A (n = 20)	Set B (n = 20)	Set C (n = 20)	Set D (n = 20)	Set E (n = 20)	Validation (n = 10)	Test (n = 15)
Image modalities	Numbers of CTA (n = 75)	13	10	14	15	12	2	9
	Numbers of CTP (n = 50)	7	10	6	5	8	8	6
Brain coverage	More than half (n = 25)	2	6	3	5	5	2	2
	Complete (n = 100)	18	14	17	15	15	8	13
Slice thickness	[0.50-0.75]mm (n = 59)	10	9	9	11	8	6	6
	[0.75-1.00]mm (n = 32)	5	5	7	3	6	1	5
	[1.00-1.50]mm (n = 34)	5	6	4	6	6	3	4
Acquisition phase	Early arterial phase (n = 19)	4	3	3	1	3	0	5
	Peak arterial phase (n = 37)	6	6	4	11	6	1	3
	Equilibrium phase (n = 69)	10	11	13	8	11	9	7
Occlusion location	ICA (n = 20)	4	4	5	2	2	2	1
	M1 (n = 60)	9	12	11	7	10	4	7
	M2 and above (n = 38)	6	3	3	9	8	3	6
	Others (n = 7)	1	1	1	2	0	1	1
Complexity	Low: < 50 (n = 40)	7	5	8	7	7	1	5
	Medium:[50-100] (n = 65)	12	12	10	9	9	6	7
	High: ≥ 100 (n = 20)	1	3	2	4	4	3	3
Proximal segments	Numbers (n = 417)	67	65	65	70	64	34	52
	Intensity (HU)	311 ±94	306 ±84	280 ±87	312 ±92	321 ±120	256 ±105	274 ±77
	Length (mm)	26 ±17	24 ±15	26 ±14	25 ±15	24 ±13	21 ±12	24 ±16
Middle segments	Number (n = 3814)	607	693	577	650	564	288	435
	Intensity (HU)	161 ±76	161 ±70	152 ±68	153 ±68	167 ±80	153 ±73	143 ±61
	Length (mm)	29 ±23	29 ±26	29 ±26	31 ±25	31 ±24	29 ±25	31 ±24
Distal segments	Number (n = 2436)	343	390	336	432	410	200	325
	Intensity (HU)	129 ±66	134 ±61	130 ±64	121 ±54	148 ±84	116 ±51	127 ±65
	Length (mm)	27 ±22	29 ±26	25 ±22	28 ±24	28 ±23	27 ±25	28 ±27
Near skull segments	Number (n = 1404)	170	237	216	220	255	112	194
	Intensity (HU)	153 ±60	149 ±61	144 ±47	139 ±45	149 ±56	128 ±46	127 ±38
	Length (mm)	63 ±40	64 ±40	61 ±39	62 ±38	61 ±37	61 ±40	57 ±39
Near skull segments (Added extra)	Numbers (n = 595)	121	171	147	94	62	0	0
	Intensity(HU)	204 ±81	140 ±56	173 ±83	148 ±63	181 ±95	NA	NA
	Length(mm)	65 ±30	60 ±30	67 ±39	63 ±30	57 ±26	NA	NA

of the MCA and ACA vessel tree in the CTA or CTP maximal intensity projection (MIP) images. Each tree always starts from the ICA-top of the corresponding side, if there is no ICA-top occlusion. The annotation of a track was discontinued when the artery was not clearly visible anymore. In our study, the vessel diameter varies from 3.6 ± 0.45 mm at the ICA top (Rai et al., 2013) to 0.45 mm in the most distal vessels. For this annotation task, we used an in-house developed annotation tool used in Su et al. (2020). In order to increase the number of segments near the skull (which are difficult vessels to track), veins near the skull were annotated in 10 out of 75 subjects (randomly selected). These veins have an appearance that is similar to arteries.

3.3. Interobserver variation on annotation

For such a difficult annotation task, it is relevant to assess the interobserver variation in the annotation. To this end, a second observer, a medical student (M) was asked to perform the same annotation. For this, three subjects from the CTP category were selected, two subjects with middle complexity and one with high complexity. The tree overlap was calculated using a dilated binary spherical shape with a radius of four voxels with observer M as the reference standard, in the same way as the overlap is computed for the method. For the three subjects, the tree overlap was 0.50, 0.70, and 0.88 respectively, with an average overlap of 0.69. The vessels where the annotation differed were mainly low intensity vessels, and vessels close to the skull.

3.4. Data preparation

In this study, we use both the CTA and the frame with maximal vessel volume from the CTP series. In order to minimize the variation

in spatial resolution, all CTA and CTP images were resampled to a 0.45 mm isotropic grid using cubic B-spline interpolation and were normalized to a range of 0 to 1 using min-max normalization.

The annotated vessel tracks were transformed into an arborecence structure from the root based on the connectivity as shown in Fig. 4a. In this way, all the bifurcations and the connected segments (vessel segment) are known. During this transformation process, two adjacent bifurcations with a distance of less than 1.8 mm were merged into one trifurcation. The vessel segment representation was subsequently transformed into a voxelized 26-connected representation for use in the training process of the directed vessel tracking model. The bifurcations were used for the bifurcation models. For the tree formation, the initial direction vectors were generated from two manually annotated starting points from the ICA top to the proximal points of each tree. The user has to click two points for the initial direction vector.

3.5. Segment label generation

To permit an analysis of the results (i.e. distal versus proximal), we define four categories (labels) for the vessels. proximal, near skull segments, middle segments (majority of the segment inside the MCA/ACA mask), and distal segments (remaining segments). The proximal segments consisted of M1 and A1, which are defined as the segments between the given root till the first bifurcation. Near skull segments were defined as the vessel segments falling into the dilated skull (thresholded at 1000 HU, dilated with a $5 \times 5 \times 5$ kernel) for consecutive five points. The middle segments were defined using the MCA and ACA mask, i.e. those segments where the majority of the points in the segment were inside these masks. The above-mentioned segments labels and the thresholded MCA and ACA maps are shown in Figs. 4b,

4c. The segment labels are only used in the analysis of the proposed method.

3.6. Experimental setup

The data in this study is used for training and testing DRL directed vessel tracking and bifurcation detection model, and for assessing tree formation. In the directed vessel tracking and bifurcation detection experiments, we use the same data distribution: in both cases, we use 100 images in a five-fold cross validation setup, where the data was stratified on complexity (defined as the number of segments per subject). The ten subjects with extra annotations of segments near the skull were evenly distributed over each fold. The validation set in the cross-validation consisted of another ten subjects (from the remaining 25), and this set was kept the same for all cross-validation experiments. The remaining fifteen subjects were used as an independent test set. The tree formation assessments were performed on the validation and test sets. 1.8 mm means 4 voxels, and with a scaling (stepsize) of 2, we thus allow an error of maximal two steps.

The image characteristics and segment labels of each fold, validation, and test set are listed in Table 3. In the following, we will name the models based on the fold. For instance, Set A is the test set of tracking model Tr-A and Bifurcation model (Bf-A) in 5 fold cross validation setting. The tracking model Tr-S and bifurcation model (Bf-S) use set A to E for training, and a test set ($n = 15$) for testing. The validation set ($n = 10$) is the same for all models.

4. Experiments and results

4.1. Implementation

4.1.1. Directed vessel centerline tracking

The proposed PPO based directed vessel tracking approach was implemented in PyTorch (Paszke et al., 2019). The model training and validation were done on NVIDIA A40 GPUs.

The PPO training uses episodic learning with a random start position from the first position to five voxels before a bifurcation point of a segment. The length of each episode was therefore arbitrary with a minimal length of five voxels. Each mini batch consists of 8 episodes, in which one episode corresponds to one segment. The training was stopped when one of the following criteria was met: (1) the total length of the track exceeds $1.5 \times$ the length of the reference standard, (2) the agent is off the reference standard segment for a distance of 1.8 mm, (3) the target was reached (4) at the beginning of the training, if the agent has gone into the direction that is opposite to the direction of the initial vector.

The hyper-parameters of the PPO were: discount factor $\gamma = 0.9$; PPO clip value $\epsilon = 0.2$; GAE parameter $\lambda = 0.95$. The model weights were updated 10 times per mini batch. The initial learning rate for Adam was $1e-5$. The learning rate is halved when the validation score does not improve for five epochs in a row. The lower bound of the learning rate was $1e-6$. During training, the only augmentation applied is the random flipping of both 3D CTA image and centerline segments along the x- and y-axis with a probability of 0.5. We use a two-stage training: first, the data from the anterior vessel tree for learning the general anterior vessel tree track is used until the model stops converging; in this stage, the target of a track was a bifurcation. In the second stage, we added the extra segments close to the skull to better learn to track vessels running in the vicinity of the skull. Also, the target is set to five points beyond the bifurcation with both branches with equal probability. In this way, the agent could learn to track beyond bifurcations.

To monitor the episodic training process, we use a curve-to-curve similarity (CCS) metric, defined as follows:

$$CCS = 1 - \frac{L_t}{L_0} \quad (8)$$

with L_0 as the max surface distance at the initial point of the reference standard curve, see also Fig. 2b. This metric is 1 in case of complete overlap and is negative when the track runs in the opposite position.

4.1.2. Bifurcation detection

The model implementation and training for bifurcation detection were similar to the PPO, both using Pytorch and the same GPU. The training samples are sub-volumes that are at randomly shifted positions along the reference centerlines. This random offset is to simulate a track resulting from the vessel tracking. The random offsets were ranging from 1 to 3 voxels. The training label was obtained by dilating the ground-truth bifurcation points with a binary spherical kernel with a radius of four voxels, in order to increase the size of true positives and make the classification problem less imbalanced. During training, the number of true positive samples was equal to the number of true negative samples. In the training phase, the initial learning rate was $1e-3$ for the Adam optimizer. The learning rate reduction scheme was identical to the PPO training. The learning rate lower bound was $1e-6$ as well.

4.1.3. Tree formation

The tree formation method was implemented in python. For the stopping criteria, tracking into brain tissue was defined as an intensity value along the track less than 50 HU for three consecutive points, and tracking into the skull is defined as tracking into a dilated skull (threshold at 1000 HU, dilated with $5 \times 5 \times 5$ kernel) for three consecutive points. The maximum length for a track was 330 mm, which is 1.2 times the maximal length of root-to-leaf distance in the training set of 100 subjects.

Probability density maps of the MCA and ACA regions, as well as a hemisphere map, were obtained using earlier described atlases (Peter et al., 2017) that were registered to the images. After transformation to the CTA or CTP images, the MCA and ACA arterial territory maps were thresholded at 0.5 to obtain a binary mask. The hemisphere map consists of three values, indicating the left and right brain hemispheres and the background.

4.2. Directed vessel tracking assessment

The assessment of the directed vessel tracking consists of an ablation study, comparisons with a baseline DQN model and the baseline DQN with our own reward function. In addition, we assess the impact of using all data. For the final model, we also present an analysis per vessel category. We use overlap rate and average distance as metrics, where the overlap rate (OR) is defined as the part of the track that matches the reference standard until the first error (distance larger than 1.8 mm). The average distance refers to the average curve-to-curve distance in world coordinates between the overlapping part of the tracking path and the reference standard. Statistical significance was assessed using a paired t-test on all metrics, and a value lower than $p = 0.05$ was considered statistically significant. The vessel tracking results are reported using a median with an interquartile range (IQR).

4.2.1. Ablation study

In the ablation study, we investigate the importance of our network architecture design choices, training scheme, and reward function. In addition, we compare our reward function with an existing reward from Li et al. (2021). We, therefore, divided our ablation study into two parts. In part one, the ablation study focused on investigating the added value of network architecture and training schemes. In part two, the ablation study investigates the impact of the reward functions.

The network architecture may have a substantial impact on the performance of DRL methods, and similarly, the activation function has been shown to have an impact on the performance (Henderson et al., 2018). The ablation study part one focuses therefore on investigating the added value of:

Table 4

Evaluation result of ablation in network architecture and training scheme (part one) and proposed method on set A.

Model	CCS (%)	Average distance (mm)	Overlap rate (%)
CNN only	75 [29, 91]	0.64 [0.57, 0.74]	100 [36, 100]
One time point	78 [40, 93]	0.63 [0.56, 0.72]	100 [46, 100]
No scaling	38 [3, 80]	0.56 [0.49, 0.66]	72 [18, 100]
ReLU	79 [40, 93]	0.62 [0.55, 0.72]	100 [44, 100]
Single stage	80 [36, 93]	0.61 [0.54, 0.72]	100 [41, 100]
Model Tr-A	82 [44, 93]	0.65 [0.58, 0.75]	100 [50, 100]

Table 5

Evaluation result of reward function ablation study(CC distance and Li reward), baseline DQN, DQN with our reward function(DQN + ours) and proposed method on set A.

Model	CCS (%)	Average distance (mm)	Overlap rate (%)
CC distance	71 [34, 88]	0.75 [0.66, 0.85]	100 [40, 100]
Li reward	53 [21, 85]	0.64 [0.56, 0.73]	61 [28, 100]
Baseline DQN	26 [13, 53]	0.69 [0.59, 0.78]	29 [13, 58]
DQN + ours	44 [15, 79]	0.73 [0.63, 0.85]	44 [16, 100]
Model Tr-A	82 [44, 93]	0.65 [0.58, 0.75]	100 [50, 100]

- the RNN module, by comparing the complete model with a model that uses only a CNN part in the A2C network (*CNN only model*),
- the RNN module with less image information, by training a model that uses image information from only one time point with RNN module (*One time point model*),
- scaled action space, by removing the scaling factor, (*No scaling model*)
- PReLU instead of ReLU (*ReLU model*),
- two-stage training, by adding all data and bifurcation extension throughout the training process (*Single stage model*).

For reference, we also train a model that contains all components *model Tr-A*.

The ablation study part two investigates the impact of various reward functions with the same model configuration as model Tr-A, by comparing our model with the following rewards:

- curve-to-curve distance only, by using the log curve-to-curve distance function (*CC distance model*),
- the point to curve distance reward of Li et al. (2021) (*Li reward*)

Fig. 5 shows the learning curves with CCS metric and Table 4 contains the test results for the ablation study of network architecture and the training scheme. The median OR values in the ablation study are 100 %, but the variation in Q1 values demonstrates that the distributions are different. The Q3 CCS of most models is above 90 %, except the *No scaling model* (83%). The difference between Tr-A and the other configurations is statistically significant for both CCS and OR metrics. In addition, statistically significant differences were found between the No scaling configuration and the other configurations in terms of average distance.

Fig. 6 shows the learning curves and Table 5 contains the results of the reward functions ablation study. Statistically significant differences were found between the Tr-A and the other configurations in all of the cases except for the average distance of *Li reward*.

4.2.2. Comparison with DQN baseline

We also compared our method with a popular DQN method. Our baseline model was introduced by Li et al. (2021). This method uses the same network architecture as Wolterink et al. (2019), which is similar to the CNN architecture of our actor network but with one time point and ReLU activation function. To do a fair comparison, we limit the action space of the baseline DQN methods to the same action space (26 connected neighborhoods). The hyperparameters and optimizer settings are the same as for the PPO training. The data used is the same as in the




Fig. 5. Curve to curve similarity (CCS) metrics of the ablation study models on network architecture and training scheme (part one) and model Tr-A (All models are trained using set B to E, validated on set validation and tested on set A).

ablation study. In addition, we also used a baseline architecture with our own reward function(*DQN + ours*) to further assess the added value of this reward function in a different DRL method. The performance of baseline DQN and (*DQN + ours*) are shown in Table 5 and the learning curves are shown in Fig. 6.

4.2.3. Amount of training data

We also investigated whether the current amount data training data is sufficient for our application. For this purpose, we trained a model (*model S*) with all 100 subjects from set A to E with the same training scheme as was used in the five fold cross validation setup. Table 6 shows the test results of model S and five models on the same independent test set. None of the differences is statistically significant.

4.2.4. Directed vessel tracking performance

The last tracking experiment focuses on the generalizability of the tracking model and further analysis of the tracking performance. For

Table 6

The evaluation result of the five models and model S from bifurcation to bifurcation on 15 independent test subjects.

Model Tr	CCS (%)	Average distance (mm)	Overlap rate (%)
Model Tr-A	83 [44, 93]	0.60 [0.54, 0.70]	100 [46, 100]
Model Tr-B	82 [40, 93]	0.60 [0.54, 0.72]	100 [46, 100]
Model Tr-C	84 [49, 93]	0.60 [0.54, 0.70]	100 [47, 100]
Model Tr-D	82 [34, 92]	0.60 [0.54, 0.70]	100 [41, 100]
Model Tr-E	84 [41, 93]	0.61 [0.55, 0.71]	100 [46, 100]
Model Tr-S	85 [50, 93]	0.59 [0.54, 0.70]	100 [46, 100]

Table 7

The evaluation result of the directed tracking from bifurcation to bifurcation using five fold cross validation over 100 subjects.

Model Tr	CCS (%)	Average distance (mm)	Overlap rate (%)
Model Tr-A	82 [44, 93]	0.65 [0.58, 0.75]	100 [50, 100]
Model Tr-B	76 [32, 92]	0.65 [0.56, 0.76]	100 [38, 100]
Model Tr-C	81 [37, 93]	0.65 [0.56, 0.76]	100 [44, 100]
Model Tr-D	78 [35, 92]	0.66 [0.57, 0.75]	100 [44, 100]
Model Tr-E	85 [46, 94]	0.64 [0.56, 0.74]	100 [55, 100]




Fig. 6. Curve to curve similarity (CCS) metrics of the rewards ablation study, baseline DQN, DQN with our reward function(DQN + ours) and model Tr-A. All models are trained using set B to E, validated on set validation and tested on set A.

this, we used 100 images from the five fold cross validation. The learning curves of the five fold cross validation are shown in Fig. 7. The test result of the models with their corresponding set is listed in Table 7. Variations in the results may be caused by differences in the validation datasets. Therefore, we compute the result for all models with the same independent test set. These results are listed in Table 6. It shows only minor differences between the results, and there are no statistically significant differences between all listed models (including model S).

For the same test set and models, we also analyze the tracking performance with regard to the segment labels. The result can be found in Table 8. They show that the performance of directed vessel tracking depends on how distal the vessels are.

4.3. Bifurcation detection assessment

The bifurcation detection was tested in two ways. First, the performance was assessed in the training setup with a dilated reference




Fig. 7. Curve to curve similarity metrics for 5 fold cross validation.

standard, with random samples from the randomly shifted voxels along the annotated tracks (see. Section 4.1.2). This was done in a five fold cross validation setup (models Bf-A – Bf-E), and the results are shown in Fig. 8. The test accuracy of all five models with different sets converges to about 0.82. The accuracy on the validation set was 0.79, slightly lower than the test performance.

The second assessment focused on bifurcation detection on paths that were obtained from the DRL-based tracker. For this, 100 subjects from the five fold cross validation and 15 subjects from the independent test set were used.

The directed vessel tracking model was combined with the corresponding bifurcation detection model. e.g. for the testing set A, tracking model A and bifurcation model A were used. In addition to the five fold cross validation models, we trained a bifurcation model Bf-S with 100 subjects. The average precision and recall for the different sets with corresponding models are listed in Table 9. The recall and precision of bifurcation detection during tree formation are around 76% and 87% respectively, and there is no statistically significant difference between the models.

Table 8


Analysis based on the different segment labels on 115 independent subjects.

Segments (type)	CCS (%)	Average distance (mm)	Overlap rate (%)
Proximal	91 [84, 94]	0.61 [0.55, 0.69]	100 [100, 100]
Middle	88 [62, 94]	0.63 [0.55, 0.73]	100 [84, 100]
Distal	80 [44, 91]	0.65 [0.56, 0.74]	100 [59, 100]
Near skull	44 [12, 82]	0.68 [0.59, 0.79]	46 [18, 81]

Table 9

The evaluation result of bifurcation detection and bifurcation inference for tree formation over 115 subjects.

Model Bf	Bf-A	Bf-B	Bf-C	Bf-D	Bf-E	Bf-S
Precision (%)	86	87	88	87	86	87
Recall (%)	78	77	76	76	75	77

**Fig. 8.** Learning curve of bifurcation model 5 fold cross validation with 100 subjects.

4.4. Tree formation

In the tree formation evaluation, we assess the performance of the combined tracker and bifurcation detector. In addition, we investigate whether the ensemble method would improve over the single model method, and the consistency of the tree formation method and we assess the impact of image characteristics. We performed a visual analysis of the false positive and false negative branches. Also, we performed interobserver analysis of tree formation between two annotators as shown in Table 10.

Two features are relevant for the tree formation method: (1) the completeness of the tree and (2) the topology of the tree (whether connectivity is correct). The completeness of the tree is measured with precision and recall. To this end, the reference standard and extracted tree were dilated using a binary spherical shape with a radius of four voxels (allowing a maximum distance of 1.8 mm). True positives are points along the extracted tree that are in the dilated reference standard, false positives are points of the extracted tree that are not being covered in dilated reference standard, and false negatives are points along the reference standard that are not in the dilated extracted tree.

The topology was assessed using the correct tree topology (CTT) ratio, which quantifies the fraction of points of the tree for which the path to the root is similar (i.e. within 1.8 mm distance everywhere) to the path to the root of the corresponding point in the reference standard.

4.4.1. Single model vs ensemble method

For the comparison of the single model (Model Tr-S, trained on 100 images) and the ensemble method (using the five cross-validation

models) both the 10 validation and 15 test images were used. The tree formation results are shown in the top row of Table 11. The precision and recall are reported with median and IQR. However, the CTT was reported using the mean and standard deviation since the median and IQR in almost all items are 100, [100, 100], except near skull segments. In general, the performance of the ensemble model is better than the performance of the single model, in terms of precision and recall. When compared with a single model, the true positive rate of the ensemble model increases by 14, [6, 30] %.

In addition, we determined the intensity distribution of false positive and false negative vessels with regard to our reference standard for the ensemble method over 25 subjects. The example results of the ensemble method can be found in Fig. 9. The false positives have a mean and standard deviation of 134 ± 91 HU, and the false negatives have a mean and standard deviation of 114 ± 62 HU.

4.4.2. Consistency test of ensemble method

We also perform a consistency test with our final tree formation ensemble method, to assess the impact of changes in the starting positions. For this experiment, we used images with moderate complexity ($n = 13$) from the test and validation set. The input of the methods was varied by randomly shifting the starting position within a range of five voxels along the given starting centerline. We ran the methods five times with a randomly shifted initial vector. The first result is used as the baseline of the tree formation consistency test. We then compute the precision and recall for the other four results. Both average precision and recall are above 99%.

4.4.3. Interobserver analysis on tree formation

For the interobserver analysis part, we would like to compare the tree formation result for the three subjects that were annotated by both observers with the annotated versions. The precision and recall regarding to the different observers were shown in Table 10.

4.4.4. Impact of imaging characteristics on tree formation

To get more insight into the performance of the methods, and how these depend on the properties of the images (CTA vs CTP, acquisition phase, number of vessels, etc.), we present the tree formation results for each of these categories in Table 11. It shows that the method performs better in the CTA images compared to CTP and that proximal vessels are better detected than distal vessels. For the other categories (acquisition phase, tree type, or complexity), there are no apparent differences in the results.

5. Discussion

In this work, we developed and assessed a method to construct tree models of the anterior vessel arteries from 3D CTA and CTP images of the brain. The method consists of a deep reinforcement learning tracking approach, a CNN-based classifier to detect bifurcations along the tracks, and a breadth-first tree construction algorithm that uses the




Fig. 9. Examples of ensemble method tree formation result in the test set. Those subjects were selected based on the median value of recall. Green denotes true positive segments, blue is false positive, and yellow is false negative. a&b: frontal and sagittal view of a subject with middle complexity, a precision of 0.85, and recall of 0.77. c&d: frontal and sagittal view of a subject with high complexity, a precision of 0.83, and recall of 0.76.

Table 10
The tree formation result using different reference standard.

	Observer JS		Observer M		Agreement
	Precision (%)	Recall (%)	Precision (%)	Recall (%)	
Subject 1	0.78	0.57	0.45	0.49	0.50
Subject 2	0.59	0.72	0.53	0.68	0.70
Subject 3	0.71	0.76	0.71	0.54	0.88

Table 11

The evaluation result of tree formation. The precision and recall are reported with median and IQR. However, the CTT were reported using the mean and standard deviation since the median and IQR in almost all items are 100, [100, 100], except near skull segments.

Properties		Single Model			Ensemble		
		Precision (%)	Recall (%)	CTT (%)	Precision (%)	Recall (%)	CTT (%)
Overall		63 [52, 74]	59 [46, 70]	90 ±31	70 [57, 81]	68 [54, 81]	87 ±34
Image modalities	CTA	66 [52, 79]	64 [51, 73]	95 ±22	72 [58, 83]	77 [63, 86]	92 ±27
	CTP	62 [52, 71]	51 [41, 63]	85 ±35	69 [57, 72]	61 [47, 78]	82 ±38
Acquisition phase	Early arterial phase	77 [67, 87]	67 [64, 73]	94 ±23	82 [77, 86]	77 [70, 81]	94 ±23
	Peak arterial phase	66 [57, 78]	58 [49, 68]	95 ±22	70 [58, 84]	80 [58, 85]	90 ±30
	Equilibrium phase	60 [51, 69]	52 [42, 67]	87 ±34	66 [55, 74]	62 [48, 81]	83 ±38
Complexity	Low	57 [48, 71]	73 [62, 83]	95 ±22	62 [56, 73]	77 [63, 87]	88 ±33
	Middle	64 [56, 78]	62 [52, 81]	86 ±35	72 [56, 82]	66 [48, 81]	85 ±36
	High	68 [58, 72]	58 [41, 64]	94 ±25	71 [67, 78]	66 [57, 76]	90 ±31
Tree type	MCA	62 [51, 72]	57 [43, 67]	88 ±33	73 [64, 83]	60 [52, 78]	84 ±37
	ACA	66 [56, 79]	59 [48, 72]	92 ±28	65 [55, 80]	70 [60, 86]	89 ±31
Segment label	Proximal	99 [68, 100]	99 [68, 100]	100 ±0	92 [71, 100]	90 [69, 100]	100 ±0
	Middle	60 [42, 76]	59 [45, 70]	97 ±16	69 [54, 79]	67 [57, 84]	98 ±16
	Distal	61 [45, 82]	47 [38, 64]	80 ±41	74 [59, 88]	67 [49, 79]	78 ±42
	Near skull	39 [27, 59]	43 [18, 75]	63 ±49	66 [48, 87]	64 [35, 75]	50 ±51

tracker and bifurcation detector to construct the tree from the tracking and bifurcation detection results. When evaluated on a test set of 25 subjects, the median precision and recall w.r.t. the manual annotation is 68% and 69%.

For the tracking, we investigate the policy gradient based proximal policy optimization DRL approach to perform the directed tracking. The average tracking performance on directed tracking from bifurcation to bifurcation over 100 subjects with 7026 vessels was 100 [46, 100]. Bifurcation detection is another essential ingredient, it was tested both in the training scenario and during the single model tree formation with the same 100 subjects. The accuracy of bifurcation detection over 100 test subjects was 82%. During the single model tree formation setting, the precision and recall were on average 87% and 76% respectively. At the final ensemble tree formation method test on 25 independent subjects, the overall precision and recall were 68 [54, 81] and 70 [57, 81].

An ablation study was performed to assess the added value of various design choices made. From the results, it follows that most of the choices (including time information via RNN and multiple time points, activation function, and training strategy) have a small but statistically significant impact on the final tracking results. Both RNN

and multiple time points serve as a means to include information from previous steps in the model. From the results it is clear that using RNN with one CNN time point has a better result than CNN only with more time points; apparently, these elements of the model permit to focus on different temporal aspects relevant for the tracking. Adding an RNN module would enhance the tracking ability.

Excluding the scaling of the step size (i.e. taking smaller steps) yields improved accumulated reward but lower CCS and OR. An explanation could be that, first, the frequency of calculating the instant reward is at least two times more than the rest of the configuration. Furthermore, in the no scaling variant, the distance between p_t and p_{t-2} is 0.9 mm. This leads to very similar image information at the three time points, which hampers the full exploitation of temporal information. The performance of this no scaling model is similar to the DQN with our reward function(DQN+ours), which is the configuration using CNN with one time point.

In the reward ablation study, CC distance performed second best, by comparing the result of CC distance and our proposed method Tr-A model, adding binary overlap in the reward function would improve the average distance. The point-to-curve distance reward has a higher

average distance but lower CCS and OR in comparison with our proposed reward. This result can be observed from both DQN methods and PPO methods. The curve-to-curve distance reward aims to optimize the surface area between the reference standard and agent path at t and $t - 1$, whereas the point-to-curve distance used by the work of Zhang et al. (2018) and Li et al. (2021) aim to optimize the distance between the agent position and corresponding curves. Using the curve-to-curve distance reward function, the instant reward at t is approximated by the difference in surface area between the agent path at t and $t - 1$ with regard to the complete reference standard. Therefore, the agent moves along the centerline because it minimizes the surface distance between the agent path and reference standard. This is particularly handy in the cases of sharp turning vessel structures, e.g. helix shape, where the correspondence between agent position and corresponding reference standard point, such as used in other approaches, cannot be found in an accurate way using the closest Euclidean distance (e.g. using point to curve distance). Therefore, the curve-to-curve distance fits better with our application.

Most existing DRL tracking methods use a DQN approach with a customized reward function; for instance, Li et al. (2021) utilize a DQN tracking approach with the same network architecture as the state-of-art method of Wolterink et al. (2019) and a point-to-curve distance based instance reward for coronary tracking in a 3D CTA image. The tracking performance of the method by Li et al. (2021) outperforms the CNN tracking result of Wolterink et al. (2019). We, therefore, used this method as our baseline method. The baseline methods perform less well in our application. Zhang et al. (2018) utilize the CNN with one time point (different architecture) to track the thoracic aorta. The reward function is similar to Li et al. (2021). Zhang et al. (2020) utilize the same CNN architecture with double-DQN methods and the dot product between the reference standard and agent path of t and $t - 1$ as the reward function for coronary centerline tracking. The corresponding point between the reference standard and the corresponding agent location was determined based on the shortest distance. The above-mentioned architectures are less complex and are sufficient for tracking the coronary arteries and aorta. The brain vasculature, as shown in Fig. 1, is different from coronary arteries and aorta, e.g. there is more bending and there is a large variety in curvatures in cerebral vessels compared to coronary vessels, and they are much smaller than the aorta. As a consequence, during tracking, a CNN with one time point only might not be able to provide sufficient information for the network to make accurate direction estimation. In addition, the agent might have difficulty finding the correspondence between the agent's position and the reference standard.

While we were finalizing this manuscript, the work of Chen et al. (2021) demonstrated the possibility of walking outside of the classical reinforcement learning problem formulation, which is either value-based or policy-based. In contrast, they regard the reinforcement learning problem as a sequence prediction problem. Similarly, we could use our trained agent and our reward with the sequential framework in an offline reinforcement learning fashion, which may improve our tracking performance. To what extent and whether the transformer would be the best sequential model for our reward function could be further investigated.

When training DL models, data is essential: both the amount of data and the quality. To investigate whether adding may further improve the results, we compared the results of the cross-validation models with a model trained on the full dataset. The results show that the additional 25% of data does not significantly improve the results, suggesting that the amount of data is sufficient for the task.

Annotation of anterior vessel trees in brain CT images is a tedious and difficult task, among others caused by the coexistence of arterial and venous structures in the whole brain. In Fig. 1, the anterior vessel tree is only a small part of the complete brain vessel annotation. Under such circumstances, kissing vessels, i.e. locations (mostly distally) where vessels (almost) touch are common. This explains the moderate

average interobserver overlap of 0.69. The intensity distribution in the branches where the observers disagree is in the range of the low intensity vessels that consist of distal and near skull segments. Though the annotations are not perfect, we assume that for the training, because of the large number of vessels, the errors in the annotations would not greatly impact the models, which is also suggested by the results. Errors in the annotation also impact the quantitative results. We therefore also compare the tree formation results with the annotations from two different observers in three datasets. The differences again are mainly in the vessels that have a low intensity, which is harder to annotate, and often are distal vessels. Also, there was a variation in vessels close to the skull. These errors are similar to the errors of the automated method, suggesting that these low-intensity vessels are difficult for both humans and the automated method.

From the analysis of the tree formation result with respect to several imaging characteristics, it followed that only the imaging modality has a significant impact on the quantitative results: the tree formation method on CTP performs less than on CTA. The difference between the validation and test set (the performance on the validation set was always worse than on the test set), may therefore be caused by the higher percentage of CTP images in the validation set. The worse performance may be caused by the use of the maximum volume image for the tracking, whereas the annotation was done in the MIP image, and the blurring was caused by the interpolation after the alignment to the first frame. After annotating the vessels of the CTP in the MIP image, we decided to use the maximum volume CTP image for training and tracking instead of the MIP. Whereas the MIP is convenient for annotation, its appearance is different from CTA images, because of the higher background intensities and more noisy appearance caused by the MIP. This is a limitation of our study, we however preferred including the (suboptimal) annotations over leaving out the CTP, or redoing the annotation.

The final result of the method is a tree representation of the brain vasculature, and this tree representation will be the basis for subsequent works. We intend to use this tree to e.g. find lesions (occlusions in more distal arteries, such as M2, which is still a challenging task Luijten et al., 2021), better quantification of collateral status, and possibly also linking/registering the 3D vascular information from CTA and CTP to interventional DSA images for improving image guidance during endovascular treatment. The tree formation method performed reasonably well in the proximal and middle segments. In the distal and near skull segments, the performances were also in the interobserver ranges. For the DSA to CTA mapping and distal occlusion detection, it is likely that such tree formation performance is sufficient.

6. Conclusion

In this study, we developed and assessed a method to construct a brain vessel tree from a CTA or CTP image, using a starting point and direction vector. The method consists of a DRL-based tracker, a CNN-based bifurcation detector, and a breadth-first tree building. The tracker performance on segments from bifurcation to bifurcation has a median overlap of 100 %, and the bifurcation detector has a mean accuracy of 82 %. The combination of both components in an ensemble tree building algorithm yields trees that have a 69 % overlap with manual annotations, which is in the interobserver variation range.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Wim van Zwam reports financial support was provided by Cerenovus. Wim van Zwam reports financial support was provided by Stryker Neurovascular. Aad van der Lugt reports financial support was provided by Siemens Healthineers. Aad van der Lugt reports financial support was provided by GE Healthcare. Aad van der Lugt reports financial support was provided by Philips Healthcare.

Data availability

The authors do not have permission to share data.

Acknowledgments

We thank all MR CLEAN (Multicenter Randomized Clinical Trial of Endovascular Treatment for Acute Ischemic Stroke in the Netherlands) Registry investigators. MR CLEAN Registry was funded through the CONTRAST consortium, which acknowledges the support from the Netherlands Cardiovascular Research Initiative, an initiative of the Dutch Heart Foundation (CVON2015-01: CONTRAST), and from the Brain Foundation Netherlands (HA2015.01.06). The collaboration project is additionally financed by the Ministry of Economic Affairs by means of the PPP Allowance made available by the Top Sector Life Sciences & Health to stimulate public-private partnerships (LSHM17016). This work was funded in part through unrestricted funding by Stryker, Medtronic and Cerenovus. The funding sources were not involved in study design, monitoring, data collection, statistical analyses, interpretation of results, or manuscript writing.

References

- Arulkumaran, K., Deisenroth, M.P., Brundage, M., Bharath, A.A., 2017. A brief survey of deep reinforcement learning. *arXiv preprint arXiv:1708.05866*.
- Baboiu, D.-M., Hamarneh, G., 2012. Vascular bifurcation detection in scale-space. In: 2012 IEEE Workshop on Mathematical Methods in Biomedical Image Analysis. IEEE, pp. 41–46.
- Bertsekas, D.P., et al., 2011. *Dynamic Programming and Optimal Control*, third ed. vol. ii, Athena Scientific, Belmont, MA.
- Chen, L., Lu, K., Rajeswaran, A., Lee, K., Grover, A., Laskin, M., Abbeel, P., Srinivas, A., Mordatch, I., 2021. Decision transformer: Reinforcement learning via sequence modeling. *Adv. Neural Inf. Process. Syst.* 34, 15084–15097.
- Dai, T., Dubois, M., Arulkumaran, K., Campbell, J., Bass, C., Billot, B., Uslu, F., De Paola, V., Clopath, C., Bharath, A.A., 2019. Deep reinforcement learning for subpixel neural tracking. In: *International Conference on Medical Imaging with Deep Learning*. PMLR, pp. 130–150.
- Damseh, R., Delafontaine-Martel, P., Pouliot, P., Cherié, F., Lesage, F., 2020. Laplacian flow dynamics on geometric graphs for anatomical modeling of cerebrovascular networks. *IEEE Trans. Med. Imaging* 40 (1), 381–394.
- Dijkstra, E.W., et al., 1959. A note on two problems in connexion with graphs. *Numer. Math.* 1 (1), 269–271.
- Engstrom, L., Ilyas, A., Santurkar, S., Tsipras, D., Janoo, F., Rudolph, L., Madry, A., 2020. Implementation matters in deep policy gradients: A case study on PPO and TRPO. *arXiv preprint arXiv:2005.12729*.
- Girshick, R., Donahue, J., Darrell, T., Malik, J., 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 580–587.
- Guo, Z., Bai, J., Lu, Y., Wang, X., Cao, K., Song, Q., Sonka, M., Yin, Y., 2019. Deep-centerline: A multi-task fully convolutional network for centerline extraction. In: *International Conference on Information Processing in Medical Imaging*. Springer, pp. 441–453.
- He, K., Zhang, X., Ren, S., Sun, J., 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 1026–1034.
- 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*. 32, (1).
- Hervella, Á.S., Rouco, J., Novo, J., Penedo, M.G., Ortega, M., 2020. Deep multi-instance heatmap regression for the detection of retinal vessel crossings and bifurcations in eye fundus images. *Comput. Methods Programs Biomed.* 186, 105201.
- Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural Comput.* 9 (8), 1735–1780.
- Jansen, I.G., Mulder, M.J., Goldhoorn, R.-J.B., 2018. Endovascular treatment for acute ischaemic stroke in routine clinical practice: prospective, observational cohort study (MR CLEAN Registry). *Bmj* 360.
- Jeon, B., 2021. Deep recursive Bayesian tracking for fully automatic centerline extraction of coronary arteries in CT images. *Sensors* 21 (18), 6087.
- Kalaie, S., Gooya, A., 2017. Vascular tree tracking and bifurcation points detection in retinal images using a hierarchical probabilistic model. *Comput. Methods Programs Biomed.* 151, 139–149.
- Kirbas, C., Quek, F., 2004. A review of vessel extraction techniques and algorithms. *ACM Comput. Surv.* 36 (2), 81–121.
- Kullback, S., Leibler, R.A., 1951. On information and sufficiency. *Ann. Math. Stat.* 22 (1), 79–86.
- Lee, T.-C., Kashyap, R.L., Chu, C.-N., 1994. Building skeleton models via 3-D medial surface axis thinning algorithms. *CVGIP, Graph. Models Image Process.* 56 (6), 462–478.
- Lesage, D., Angelini, E.D., Bloch, I., Funka-Lea, G., 2009. A review of 3D vessel lumen segmentation techniques: Models, features and extraction schemes. *Med. Image Anal.* 13 (6), 819–845.
- Li, Z., Xia, Q., Hu, Z., Wang, W., Xu, L., Zhang, S., 2021. A deep reinforced tree-traversal agent for coronary artery centerline extraction. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 418–428.
- Luijten, S.P., Wolff, L., Duvekot, M.H., van Doormaal, P.-J., Moudrouts, W., Kerkhoff, H., a Nijeholt, G.J.L., Bokkers, R.P., Lonneke, S., Hofmeijer, J., et al., 2021. Diagnostic performance of an algorithm for automated large vessel occlusion detection on CT angiography. *J. Neurointerventional Surg.*
- Luo, P., Ren, J., Peng, Z., Zhang, R., Li, J., 2018. Differentiable learning-to-normalize via switchable normalization. *arXiv preprint arXiv:1806.10779*.
- Maas, A.L., Hannun, A.Y., Ng, A.Y., et al., 2013. Rectifier nonlinearities improve neural network acoustic models. In: *Proc. Icml*. 30, (1), Atlanta, Georgia, USA, p. 3.
- Mnih, V., Badia, A.P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., Kavukcuoglu, K., 2016. Asynchronous methods for deep reinforcement learning. In: *International Conference on Machine Learning*. PMLR, pp. 1928–1937.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., Riedmiller, M., 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.
- Moccia, S., De Momi, E., El Hadji, S., Mattos, L.S., 2018. Blood vessel segmentation algorithms—review of methods, datasets and evaluation metrics. *Comput. Methods Programs Biomed.* 158, 71–91.
- Moriconi, S., Zuluaga, M.A., Jäger, H.R., Nachev, P., Ourselin, S., Cardoso, M.J., 2018. Inference of cerebrovascular topology with geodesic minimum spanning trees. *IEEE Trans. Med. Imaging* 38 (1), 225–239.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al., 2019. Pytorch: An imperative style, high-performance deep learning library. *Adv. Neural Inf. Process. Syst.* 32.
- Peter, R., Emmer, B.J., van Es, A.C., van Walsum, T., 2017. Cortical and vascular probability maps for analysis of human brain in computed tomography images. In: 2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017). IEEE, pp. 1141–1145.
- Potter, C.A., Vagal, A.S., Goyal, M., Nunez, D.B., Leslie-Mazwi, T.M., Lev, M.H., 2019. CT for treatment selection in acute ischemic stroke: a code stroke primer. *Radiographics* 39 (6), 1717–1738.
- Rai, A.T., Hogg, J.P., Cline, B., Hobbs, G., 2013. Cerebrovascular geometry in the anterior circulation: an analysis of diameter, length and the vessel taper. *J. Neurointerventional Surg.* 5 (4), 371–375.
- Robben, D., Türetken, E., Sunaert, S., Thijs, V., Wilms, G., Fua, P., Maes, F., Suetens, P., 2016. Simultaneous segmentation and anatomical labeling of the cerebral vasculature. *Med. Image Anal.* 32, 201–215.
- Rodriguez-Luna, D., Dowlatshahi, D., Aviv, R.I., Molina, C.A., Silva, Y., Dzialowski, I., Lum, C., Czlonkowska, A., Boulanger, J.-M., Kase, C.S., et al., 2014. Venous phase of computed tomography angiography increases spot sign detection, but intracerebral hemorrhage expansion is greater in spot signs detected in arterial phase. *Stroke* 45 (3), 734–739.
- 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.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., Moritz, P., 2015a. Trust region policy optimization. In: *International Conference on Machine Learning*. PMLR, pp. 1889–1897.
- Schulman, J., Moritz, P., Levine, S., Jordan, M., Abbeel, P., 2015b. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O., 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Su, J., Wolff, L., van Es, A.C.M., Van Zwam, W., Majoie, C., Dippel, D.W., Van Der Lugt, A., Niessen, W.J., Van Walsum, T., 2020. Automatic collateral scoring from 3D CTA images. *IEEE Trans. Med. Imaging* 39 (6), 2190–2200.
- Sutton, R.S., Barto, A.G., 2018. *Reinforcement Learning: An Introduction*. MIT Press.
- Van Hasselt, H., Guez, A., Silver, D., 2016. Deep reinforcement learning with double q-learning. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 30, (1).
- Walsum, T.v., Schaap, M., Metz, C.T., Giessen, A.G., Niessen, W.J., 2008. Averaging centerlines: mean shift on paths. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 900–907.
- Wolterink, J.M., van Hamersvelt, R.W., Viergever, M.A., Leiner, T., Išgum, I., 2019. Coronary artery centerline extraction in cardiac CT angiography using a CNN-based orientation classifier. *Med. Image Anal.* 51, 46–60.
- Yu, F., Koltun, V., 2015. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*.
- Zhang, Y., Luo, G., Wang, W., Wang, K., 2020. Branch-aware double DQN for centerline extraction in coronary ct angiography. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 35–44.

- Zhang, Z., Marin, D., Drangova, M., Boykov, Y., 2021. Confluent vessel trees with accurate bifurcations. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9573–9582.
- Zhang, P., Wang, F., Zheng, Y., 2018. Deep reinforcement learning for vessel centerline tracing in multi-modality 3D volumes. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 755–763.
- Zhao, H., Sun, Y., Li, H., 2020. Retinal vascular junction detection and classification via deep neural networks. *Comput. Methods Programs Biomed.* 183, 105096.
- Zheng, Y., Liu, D., Georgescu, B., Nguyen, H., Comaniciu, D., 2015. 3D deep learning for efficient and robust landmark detection in volumetric data. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 565–572.
- Zhou, S.K., Le, H.N., Luu, K., Nguyen, H.V., Ayache, N., 2021. Deep reinforcement learning in medical imaging: A literature review. *Med. Image Anal.* 73, 102193.