
Nonlinear adaptively learned optimization for object localization in 3D medical images

Mayalen Etcheverry Bogdan Georgescu Benjamin Odry Thomas J. Re
Shivam Kaushik Bernhard Geiger Nadar Mariappan Sasa Grbic

Dorin Comaniciu

Medical Imaging Technologies, Siemens Healthineers, Princeton NJ, USA

Abstract

When training models for medical imaging tasks such as image registration, volumetric organ segmentation, lesion quantification and abnormality detection, prior localization of the target anatomy is an important prerequisite to enhance consistency in the local context. Using deep reinforcement learning, we demonstrate a novel method to actively learn how to localize an object in the volumetric scene. We show the applicability of our method by localizing boxes (9 degrees of freedom) on a set of acquired MRI scans of the brain region.

1 Context

Automatic localization of anatomical structures in the context of 3D data can be approached using several types of method such as atlas-based registration methods [Ranjan (2011)], regression-based methods [Criminisi et al. (2013); Cuingnet et al. (2012)] or classification-based methods [Zheng et al. (2009); Ghesu et al. (2016)]. More recently, Akselrod-Ballin et al. (2016) proposed to apply the Faster R-CNN technique developed by Ren et al. (2015) to medical imaging analysis. Object classification and object localization is jointly performed in a single forward pass to decrease the processing time. We propose a novel method to automatically estimate the 9 parameters (position, rotation and scale) of an anatomical bounding box. Unlike previous methods where the search is performed onto a set of independent object proposals, we propose an active search strategy to learn the optimal convergence path. We build upon the work of Ghesu et al. (2017) that uses reinforcement learning to identify the location of anatomical landmarks in a set of image data. We extend the decision-based search strategy framework, limited to finding a set of coordinates (x,y,z), to a wider range of image analysis applications by expanding the search space to a nonlinear multi-dimensional parametric space.

2 Method

The sought object is modeled with a set of D independent parameters $\{x_i\}_{i=1}^D$. An intelligent agent is deployed into the D -dimensional parametric space and can navigate, through a sequence of simple control actions, with the goal of reaching the optimal parameter vector $x^* = (x_1^*, \dots, x_D^*)$.

2.1 Autonomous learning of the control strategy through deep reinforcement learning

We model the problem as a Markov Decision Process (MDP). At each time step t , the agent observes the content of the currently attended volume region. We resample it to match a fixed-




Figure 1: Schematic illustration of the proposed control strategy. Measurement from the image (state) drive the output of the deep-Q-network which itself drive the agent decisions. In the proposed MDP, the agent follows a multi-scale progressive control strategy and adjusts the box 9 degrees of freedom.

size grid of voxels that we use as input state s_t of the network. In return, the agent selects an action a_t to either modify the current object geometry x_t (by adjusting one of the D parameters by ± 1 in the discretized space) or to terminate the search with the *stop* action. Accordingly to the consequences of its action, the agent receives a dense signal reward $r_t =$

$$\begin{cases} \text{dist}(x_t, x^*) - \text{dist}(x_{t+1}, x^*) & \text{if } a_t \text{ modifies the object geometry} \\ \left(\frac{\text{dist}(x_t, x^*) - d_{min}}{d_{max} - d_{min}} - 0.5 \right) * 6 & \text{if } a_t \text{ is the stop action} \\ -1 & \text{if } x_{t+1} \text{ not legal} \end{cases}$$

where dist defines a metric distance in the parametric space and x^* is the annotated ground-truth parameter set. Intuitively, the reward is positive when the agent gets closer to the ground truth target and negative otherwise. If an action leads to a *non-legal* set of parameters (outside of a predefined allowed search range), the agent receives a negative reward -1. If the agent decides to stop, the closer it is from the target the greater reward it gets and reversely.

We use a deep Q-network (DQN), as introduced by Mnih et al. (2013), to estimate the optimal action-value: $Q^*(s, a) \approx Q(s, a, \theta)$. The training uses Q-learning to update the network by minimizing a sequence of loss functions $L_i(\theta_i)$ expressing how far $Q(s, a; \theta_i)$ is from its target y_i : $L_i(\theta_i) = \mathbb{E}_{s, a, r, s'} (y_i - Q(s, a; \theta_i))^2$. For effective training of the DQN, the proposed concepts of experience replay, ϵ -greedy exploration and loss clipping are incorporated. The exploration is constrained to actions leading to positive reward to accelerate the agent’s discovery of good trajectory. We also use double Q-learning as proposed by Van Hasselt et al. (2016) with a “frozen” version of the online network as target network $Q_{target} = Q(\theta_{i'}), i' < i$.

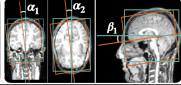
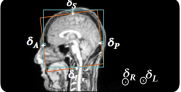
2.2 Multi-scale progressive control strategy

The global image context in which the agent evolves is downsampled to a multi-scale image pyramid with increasing image resolution L_1, L_2, \dots, L_N . The agent starts the search with both coarse field-of-view and coarse control. Following the sampling scheme of the global image context, the agent gains finer control over the parameter each time it transitions to a finer scale level L_{i+1} . This scheme goes on until the finest scale level, where the final region attended by the agent is taken as estimated localization result. The *stop* action triggers the transition between subsequent scale levels and acts as a stopping criterion at the finest scale level L_N . At inference time, if the maximum number of steps is exceeded or if the agent enters in a loop, the stop action is forcefully triggered.

3 Application to localize a standard box from Scout / Localizer images

MRI scans of the head are acquired along some specific brain anatomical regions to standardize orientations of acquisitions in order to facilitate the reading and assessment of clinical follow-up studies. We therefore propose to localize a standard box that covers the brain and is aligned along specific orientations to show the applicability of our method. Ground-truth boxes have been annotated based on specific brain anatomical structures. The orientation of the box is defined by the brain midsagittal plane (MSP) and by two anatomical points determining the rotational alignment within

Table 1: Absolute mean and maximal errors of the 30 test cases with respect to ground truth boxes.

	Inter-rater	Landmark-based	Our approach		
			(4mm)	(2mm)	
	$\alpha_1(^{\circ})$	0.99(≤ 3.50)	0.92(≤ 3.45)	1.28(≤ 3.78)	0.92(≤ 3.23)
	$\alpha_2(^{\circ})$	1.04(≤ 4.71)	0.99(≤ 4.93)	1.20(≤ 4.46)	0.97(≤ 2.11)
	$\beta_1(^{\circ})$	1.47(≤ 5.19)	2.00(≤ 6.86)	1.62(≤ 6.35)	1.39(≤ 5.86)
	$\delta_R(\text{mm})$	1.32(≤ 3.54)	2.06(≤ 5.78)	2.65(≤ 7.54)	1.45(≤ 3.30)
	$\delta_L(\text{mm})$	1.45(≤ 4.75)	1.89(≤ 5.03)	2.20(≤ 8.68)	1.83(≤ 4.95)
	$\delta_A(\text{mm})$	2.00(≤ 3.36)	1.65(≤ 4.93)	2.46(≤ 6.07)	1.94(≤ 6.08)
	$\delta_P(\text{mm})$	1.48(≤ 3.89)	1.86(≤ 9.62)	3.31(≤ 9.68)	1.65(≤ 5.68)
	$\delta_I(\text{mm})$	3.33(≤ 3.61)	2.22(≤ 6.00)	3.12(≤ 11.5)	2.74(≤ 8.21)
	$\delta_S(\text{mm})$	1.3(≤ 3.28)	2.13(≤ 5.74)	3.04(≤ 7.46)	2.16(≤ 6.31)




Figure 2: Four samples of the box evolution (blue) during inference on a challenging case.

the MSP. Given this orientation, the lower margin of the box is defined to intersect the center of C1-vertebrae arches points. The other box extremities define an enclosing bounding box of the brain. Following the annotation protocol, the box orientation is parametrized by three angles: α_1 and α_2 which control respectively the yaw and pitch of the MSP, and β_1 which controls the inplane roll around \vec{i} . The center position is parameterized by its cartesian coordinates $C = (C_x, C_y, C_z)$ and the scale by the width w , depth d and height h of the box.

In our experiments, the scale space is discretized into 4 levels: 16mm (L_1), 8mm (L_2), 4mm (L_3) and 2mm (L_4). At inference, the very first box is set to cover the whole image at the coarsest scale and is sequentially refined following the agent’s decisions. 500 annotated MRI scans of the head region were used for training and 30 for testing. The 30 test cases were annotated twice by different experts to compute the inter-rater variability. 15 additional challenging test cases with pathologies (tumors or fluid swelling in brain tissue), in plane rotation of the head, thick cushion of the head rest, or cropped top of the skull were selected to evaluate the robustness of the method.

Table 1 shows comparison between the proposed method, human performances (inter-rater variability) and a previous landmark-based method. The landmark-based method uses the proposed algorithm of Ghesu et al. (2017) to detect 14 landmarks carefully chosen after the box definition. Then, a box is robustly fitted to minimize angular and positional errors with respect to the detected landmarks. Better results are achieved for angles α_1 and α_2 than β_1 because more landmarks are associated with them. Also, a landmark only method performance degrades in the presence of larger orientation and scale differences of the target, however for this test we included target variations that can be handled by a landmark only approach. The proposed method however, does not rely on the previous detection of specific points achieving performances in the range of the inter-observer variability for every measure. The finer scale level is set to 2mm, meaning that our method achieves an average accuracy of 1-2 voxels precision. Moreover, we did not observe any major failure over the 15 “difficult” test cases, showing robustness of the method. The average detection time is 0.6s on GPU and if the search is stopped at L_3 scale-level (4mm) the detection time is 0.15s.

4 Conclusion

This paper proposes a novel approach, based on deep reinforcement learning, to sequentially search for a target object inside 3D medical images. The method can robustly localize the target object and achieves high speed and high accuracy results. The methodology can learn optimization strategies eliminating the need for exhaustive search or for complex generic nonlinear optimization techniques. The proposed object localization method can be applied to any given parametrization and imaging modality type.

Disclaimer: This feature is based on research, and is not commercially available. Due to regulatory reasons, its future availability cannot be guaranteed.

References

- Ayelet Akselrod-Ballin, Leonid Karlinsky, Sharon Alpert, Sharbell Hasoul, Rami Ben-Ari, and Ella Barkan. A region based convolutional network for tumor detection and classification in breast mammography. In *Deep Learning and Data Labeling for Medical Applications*, pages 197–205. Springer, 2016.
- Antonio Criminisi, Duncan Robertson, Ender Konukoglu, Jamie Shotton, Sayan Pathak, Steve White, and Khan Siddiqui. Regression forests for efficient anatomy detection and localization in computed tomography scans. *Medical image analysis*, 17(8):1293–1303, 2013.
- Rémi Cuingnet, Raphael Prevost, David Lesage, Laurent D Cohen, Benoît Mory, and Roberto Ardon. Automatic detection and segmentation of kidneys in 3d ct images using random forests. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 66–74. Springer, 2012.
- Florin C Ghesu, Edward Krubasik, Bogdan Georgescu, Vivek Singh, Yefeng Zheng, Joachim Hornegger, and Dorin Comaniciu. Marginal space deep learning: efficient architecture for volumetric image parsing. *IEEE transactions on medical imaging*, 35(5):1217–1228, 2016.
- Florin Cristian Ghesu, Bogdan Georgescu, Yefeng Zheng, Sasa Grbic, Andreas Maier, Joachim Hornegger, and Dorin Comaniciu. Multi-scale deep reinforcement learning for real-time 3d-landmark detection in ct scans. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Sohan R Ranjan. Organ localization through anatomy-aware non-rigid registration with atlas. In *Applied Imagery Pattern Recognition Workshop (AIPR), 2011 IEEE*, pages 1–5. IEEE, 2011.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *AAAI*, volume 2, page 5. Phoenix, AZ, 2016.
- Yefeng Zheng, Bogdan Georgescu, and Dorin Comaniciu. Marginal space learning for efficient detection of 2d/3d anatomical structures in medical images. In *International Conference on Information Processing in Medical Imaging*, pages 411–422. Springer, 2009.