



Communicative Reinforcement Learning Agents for Landmark Detection in Brain Images

Guy Leroy, Daniel Rueckert, and Amir Alansary^(✉)

Imperial College London, London, UK
a.alansary14@imperial.ac.uk

Abstract. Accurate detection of anatomical landmarks is an essential step in several medical imaging tasks. We propose a novel communicative multi-agent reinforcement learning (C-MARL) system to automatically detect landmarks in 3D medical scans. C-MARL enables the agents to learn explicit communication channels, as well as implicit communication signals by sharing certain weights of the architecture among all the agents. The proposed approach is evaluated on two brain imaging datasets from adult magnetic resonance imaging (MRI) and fetal ultrasound scans. Our experiments show that involving multiple cooperating agents by learning their communication with each other outperforms previous approaches using single agents.

1 Introduction

Robust and fast landmark localization is an essential step in medical imaging analysis applications including biometric measurements of anatomical structures [13], registration of 3D volumes [9] and extraction of 2D clinical standard planes [8]. Manual labeling of such landmarks is often a time-consuming and tedious task, which is also error-prone and requires human experts. Developing accurate and automatic detection methods will help reduce the human error and speed the diagnosis process. Recent advances in reinforcement learning (RL) have shown a significant contribution to clinical applications such as automated medical diagnosis, object localization, and landmark detection [21]. RL enables learning from reward signals that guide the agent towards the target solution in sequential steps during training. It learns to perform a non-exhaustive search without using the full 3D image as an input. RL can be data efficient by using the same 3D image for training with different starting points and states. RL has proven to achieve the best performance for landmark detection outperforming supervised methods [2, 6, 7].

Related Work: Previous works detecting anatomical landmarks have examined approaches including statistical shape priors, regression forests [5, 12], Hough voting [3], supervised convolutional neural network (CNN) [8] and attention-based autoencoder [22]. With the recent advances of deep RL, Ghesu et al. [6]

introduced the application of RL to detect anatomical landmarks by learning sequential actions towards the target landmark, while outperforming supervised methods. Alansary et al. [2] then evaluated multiple deep Q-network (DQN) variants for the detection task, namely DQN [10], double DQN [16], dueling DQN [19], and double dueling DQN. They also incorporated hierarchical steps with the multi-scale search strategy, which significantly decreased the search time. Multi-scale agents have proven to outperform fixed-scale agents for detecting the majority of landmarks [2, 7]. Vlontzos et al. [17] proposed the first multi-agent system for landmark detection, where the agents communicate efficiently by sharing the convolutional weights of the CNN model. Furthermore, RL has been utilized in various medical applications such as the detection of standardized view planes in MRI scans [1], organ localization in CT scans [11], and re-identifying the location of brain landmarks in pre- and post-operative images [18].

Contributions: (I) We propose a novel communicative multi-agent reinforcement learning for multiple landmarks detection. (II) Experiments are evaluated on two different brain imaging datasets from adult MRI and fetal ultrasound, outperforming previously published RL state-of-the-art results. (III) The implementation of the code is publicly available.

2 Background

Reinforcement learning (RL) is a sub-field of machine learning (ML), which lies under the bigger umbrella of artificial intelligence (AI). Inspired from behavioral psychology and neuroscience [15], an RL agent takes actions within an environment and receives updated states with associated rewards during training. These reward signals guide the agent to take correct actions towards the target solution, and penalize otherwise. Thus, the agent learns a policy π directly from high-dimensional inputs. In most modern applications, including ours, agents will not have total knowledge of all environment’s states. This is referred to as a partially observable Markov decision process (MDP). RL offers an efficient solution to deal with the MDP by learning a policy that maximizes the total rewards. For instance, Q-learning [20] seeks to find a q-value that measures the quality of taking an action a given a current state s by learning a policy π that maximizes the total reward during training. Mnih et al. [10] proposed to approximate these q-values using a deep neural network (θ), named DQN. The Q -function is based on the Bellman equation [4], and defined as the expected discounted cumulative rewards:

$$Q^\pi(s_t, a_t) = E_\pi \left[\sum_{k=0}^{\infty} \gamma^k r_{t+k+1} | s_t, a_t \right], \quad (1)$$

where s_t and a_t represent the state and action at step t . γ^k is the discount factor at k -th future state. DQN introduces another target network \hat{Q} that stabilizes the training, and reduce the overestimation of the maximum Q-value [10]. Whereas

at every predefined interval during training, the weights θ of the Q-network are copied to the target network $\hat{\theta}$. The DQN loss function is defined as:

$$L_i(\theta_i) = E_{s,a,r,s'} \left[\left(r + \gamma \max_{a'} \hat{Q}(s', a'; \hat{\theta}_i) - Q(s, a; \theta_i) \right)^2 \right], \quad (2)$$

where s' and a' are the next state and action. Van Hasselt et al. [16] introduced a modification to the DQN loss function to decouple the selected action from the target network, known as double DQN. This changes the loss function to,

$$L_i(\theta_i) = E_{s,a,r,s'} \left[\left(r + \gamma \hat{Q}(s', \operatorname{argmax}_{a'} Q(s', a'; \theta); \hat{\theta}_i) - Q(s, a; \theta_i) \right)^2 \right]. \quad (3)$$

The dueling network [19] uses the hypothesis that Q-values are only important in key states. It has two sequences of fully connected layers to separately estimate state-values and the advantages for each action as scalars.

Alansary et al. [2] have shown that the optimal DQN architecture depends on each landmark, where there was no overall best architecture for all landmarks. Thus, we use the double DQN as a baseline architecture.

3 Methods

In this work, we propose a communicative DQN-based RL agents for the detection of anatomical landmarks in brain images. These agents are designed to learn by communication during their search for different landmarks in 3D medical scans. This is motivated by the fact that anatomical landmarks are usually spatially correlated in the brain. Figure 1 demonstrates a schematic visualization of these navigating agents in a 3D scan or environment E .




Fig. 1. A schematic diagram of the proposed multi-agents interacting with the 3D image environment E . At each step, each agent takes an action towards a target landmark. The learned policy is formed by the path between the starting points and the target landmarks after taking the sequential actions.

States: Each state s is defined as a region of interest (ROI) of size $45 \times 45 \times 45$ voxels, and centered around each agent. To improve the network’s stability and convergence, it takes as an input a history of the last 4 states [10]. Each agent starts at a random location within the 80% of the inner region of the image at the beginning of each episode. An agent terminates navigating when it finds the target landmark. During inference the terminal state is triggered when the agent oscillates around a target point.

Action Space: It is defined based on the six directions in the 3D Cartesian coordinates, namely left, right, up, down, forward or backward. Similar to [2], we adopt a multi-scale search strategy with hierarchical steps by reducing the step and ROI size when the agent oscillates around a target point. We use three levels of scales $\{3, 2, 1\}$ mm. The episode is terminated when all agents reach their terminal states at the 1mm scale.

Rewards: First, we calculate the Euclidean distance between the current point of interest and target landmark d_t , and between the point of interest of the previous step and the target landmark d_{t-1} . The reward signal is then calculated using the difference between d_{t-1} and d_t , and clipped between -1 and 1 . This ensures that positive rewards are given, if the movements of the agent are towards the target solution.

Communicative Agents: We leverage two types of communications between the agents. Implicit communication is learned by sharing the convolutional layers of the model among all the agents [17]. Besides, communication signals are learned explicitly by sharing communication channels in the fully connected (FC) layers [14]. This is implemented by averaging the output of each FC layer for each agent, which is then concatenated with the input of the next FC layer, as seen in Fig. 2.

Network Architecture: Figure 2 shows the architecture of the proposed C-MARL model, which takes as an input a tensor of size `number_agents` $\times 4 \times 45 \times 45 \times 45$. It consists of four 3D convolutional and three 3D max pooling layers, followed by four FC layers. Whereas the convolutional layers are shared between all the agents, and each agent has its own FC layer. The output of all FC layers of each agent are averaged and concatenated with the input of the next FC layer. The size of the last FC layer is the same size of the action space. Finally, the model is trained using Eq. 3.

4 Experiments

The performance of the proposed C-MARL agents for anatomical landmark detection is tested on two brain imaging datasets, and evaluated against a single RL agent [2] and multi-agents that share only their convolutional layers

(Collab-DQN) [17]. Clinical experts manually annotated all selected landmarks using three orthogonal views. We have randomly split both datasets into train (70%), validation (15%) and test (15%) subsets. Best model is selected during training based on the best accuracy on the validation subset. The Euclidean distance error between the detected and target landmarks is used to measure the reported accuracy. The agents follow an ϵ -greedy policy, where each agent can take a random action step uniformly sampled from the action space with an initial probability of $\epsilon = 1$ to $\epsilon = 0.1$, instead of selecting the step with the highest Q-value. During testing, agents follow a full greedy policy with $\epsilon = 0$. The episode ends when all agents oscillate at the smallest scale, or after a predefined maximum number of 200 steps. Figure 3 shows C-MARL performing with five agents to detect five different landmarks from a brain MRI scan.

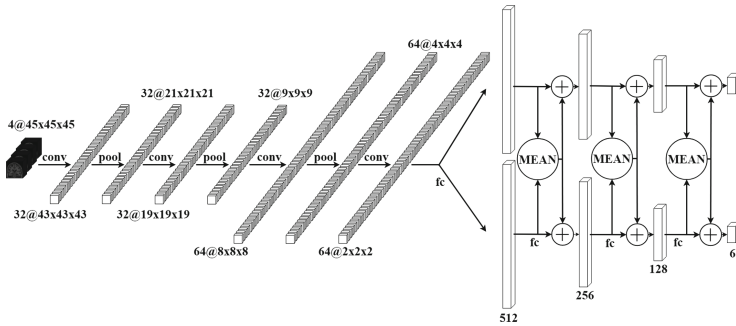


Fig. 2. The proposed C-MARL architecture for anatomical landmark detection. Here is an example of two agents sharing the same convolutional layers. They learn to communicate by averaging the output of the FC layer of each agent, which is then concatenated to the input of the next FC layer.




Fig. 3. An example of our proposed C-MARL system consisting of 5 agents. These agents are looking for 5 different landmarks in a brain MRI scan. Each agent's ROI is represented by a yellow box and centered around a blue point, while the red point is the target landmark. ROI is sampled with 3 mm spacing at the beginning of every episode. The length of the circumference of red disks denotes the distance between the current and target landmarks in z -axis. (Color figure online)

4.1 Results

Experiment (I): We use 832 T1-weighted 1.5T MRI brain scans from the Alzheimer’s disease neuroimaging initiative (ADNI)¹. All brain images are skull-stripped, and have an isotropic 1 mm³ voxel size. The selected subjects include patients with cognitively normal (CN), mild cognitive impairment (MCI), and early Alzheimer’s disease (AD). We select 8 landmarks, namely the anterior commissure (AC), the posterior commissure (PC), the outer aspect, the inferior tip and inner aspect of the splenium of the corpus callosum (SCC), the outer and inner aspect of the Genu of corpus callosum (GCC), and the superior aspect of pons.

Table 1 demonstrates the performance of the different approaches, whereas C-MARL with three agents achieves the best accuracy for all the three selected landmarks. The table also shows experiments using larger number of agents (five and eight). These experiments results in a decrease in the accuracy in most of the landmarks compared to the results using three agents. Thus, intuitively, increasing the number of agents may require architectures with a bigger capacity to be able to learn more communications. Another explanation can be that adding more landmarks, that are not strongly correlated, may affect the detection accuracy.

Table 1. Comparison between single, multiple, and communicative agents for landmark detection in brain MRIs. Distance errors are in *mm*.

Landmark	Single agent [2]	Collab-DQN [17]			C-MARL		
		3 agents	5 agents	8 agents	3 agents	5 agents	8 agents
AC	1.14 ± 0.53	1.16 ± 0.59	1.13 ± 0.64	1.21 ± 0.92	1.04 ± 0.58	1.12 ± 0.65	1.84 ± 0.91
PC	1.18 ± 0.55	1.25 ± 0.57	1.19 ± 0.61	1.22 ± 0.93	1.13 ± 0.66	1.25 ± 0.55	1.38 ± 0.64
Outer SCC	1.47 ± 0.64	1.38 ± 0.75	1.51 ± 0.77	1.46 ± 0.90	1.35 ± 0.66	1.62 ± 0.79	5.20 ± 13.49
Inferior SCC	2.40 ± 1.13	–	1.39 ± 0.85	1.53 ± 0.87	–	1.50 ± 0.89	1.87 ± 1.28
Inner SCC	1.46 ± 0.73	–	1.53 ± 0.97	2.09 ± 3.65	–	1.53 ± 0.76	3.56 ± 9.42

Experiment (II): We use 72 subjects of 3D fetal head ultrasound scans from the iFIND project². All images are resampled to isotropic voxel size with average dimensions of 324 × 207 × 279 voxels. We select the right and left cerebellum (RC and LC respectively), the cavum septum pellucidum (CSP) and the center and anterior head (CH and AH respectively) landmarks.

Table 2 shows multiple agents have a lower distance error across all fetal landmarks, while C-MARL significantly outperforms the other methods for detecting the CSP and CH. Similar to the previous experiment, increasing the number of agents did not necessarily improve the detection accuracy. However, the AH landmark has significantly benefited from increasing the number of agents. In this

¹ <http://adni.loni.usc.edu>.

² <http://www.ifindproject.com>.

experiment, results show that multi-agent system is superior in all landmarks, but rather suggest the best architecture depends on the landmark.

Table 2. Comparison between single, multiple, and communicative agents for landmark detection in fetal head ultrasound. Distance errors are in *mm*.

Landmark	Single agent [2]	Collab-DQN [17]			C-MARL		
		3 agents	5 agents	8 agents	3 agents	5 agents	8 agents
RC	7.23 ± 3.54	2.73 ± 1.71	4.20 ± 3.76	3.39 ± 2.36	6.53 ± 4.21	4.06 ± 2.95	4.75 ± 3.28
LC	4.37 ± 1.45	4.20 ± 2.87	5.98 ± 8.58	5.42 ± 4.50	5.10 ± 3.66	4.43 ± 32.26	4.64 ± 3.16
CSP	9.90 ± 3.13	5.18 ± 2.05	8.02 ± 5.34	5.74 ± 5.07	5.78 ± 3.04	5.13 ± 3.51	7.08 ± 4.13
CH	29.43 ± 17.83	–	14.45 ± 5.25	16.83 ± 12.54	–	13.00 ± 4.97	16.29 ± 8.94
AH	5.73 ± 2.88	–	8.11 ± 5.22	4.10 ± 2.26	–	4.33 ± 2.96	8.89 ± 4.91

Experiment (III): The previous experiments are conducted in the scenario of using a single agent for the detection of one landmark. In this experiment, we proceed to evaluate the performance of using multi-agents for detecting the same single landmark. The final location of the agents are averaged at the end of an episode. To give a baseline, we include a column for five single agents looking for the same landmark in parallel. We report the results on a selected landmark from each dataset used in the previous two experiments, namely AC and CSP. Table 3 shows C-MARL’s results are much better than in any of the previous methods. Parallel single agents are not significantly better than the results with only one agent.

Table 3. Results from using five agents looking for the same landmark. Distance error are in *mm*.

Landmarks	Single agents [2]	Collab-DQN [17]	C-MARL
AC	0.97 ± 0.40	0.81 ± 0.36	0.75 ± 0.34
CSP	10.43 ± 4.28	6.66 ± 4.19	5.10 ± 4.25

Experiment (IV): We further evaluate using multi agents for detecting multiple landmarks, where each single landmark have multiple agents. In this experiment, we train four agents to detect the AC and PC landmarks, where each landmark has two dedicated agents. Similar to the previous experiment, to give a baseline, we compare with four non communicating agents as a baseline. Table 4 shows that C-MARL agents perform better than the baseline, but worse than using five agents for a single landmark from Experiment (III). Finally, these

experiments show that multiple cooperative agents trained to detect one single landmark can outperform the same number of agents detecting different landmarks.

Table 4. Results from using two pairs of agent looking for two landmarks (four agents in total). Distance error are in *mm*.

Landmarks	Single agents [2]	C-MARL
AC	1.17 ± 0.61	0.95 ± 0.43
PC	1.12 ± 0.55	0.97 ± 0.46

Implementation: We run each experiment for four days, but each would converge usually after one or two days. We used Nvidia Tesla or Nvidia GeForce GTX Titan Xp with 12 GB RAM, using CUDA v10.0.130 and Torch v1.4. A 24-core/48 thread Intel Xeon CPU was used with 256 GB RAM. In four days, collab-DQN ran 30k episodes while our proposed method only ran 20k episodes. The memory space during training is mostly driven up by the memory buffer, which we set to $\frac{100,000}{\#agents}$ episodes. As for the model’s size, more agents take up more space and communication channels are added on the collab-DQN’s architecture. More precisely, our model size is 5,504,759 and 8,144,365 bytes for three and five agent respectively, while for collab-DQN it is 3,529,451 and 4,852,185 bytes. For comparison, three single agents working independently have model size $2,206,723 \times 3 = 6,620,169$ bytes and for five single agents it is $2,206,723 \times 5 = 11,033,615$ bytes. This shows multi-agent models greatly reduce the models’ trainable parameters. For the testing speed, our method takes around 2.5 and 4.9s per episode for three and five agents respectively and those figures are 2.2 and 4.2s for collab-DQN. The code is publicly available on Github, <https://github.com/gml16/rl-medical>.

5 Conclusion

We introduced a communicative multi-agent reinforcement learning (C-MARL) system for detecting multiple anatomical landmarks from brain medical images. Multi-agents share the weights of the convolutional layers to learn implicit communications. They also learn explicit communication channels calculated from the output of their fully connect layers, which are then shared among them by concatenating to the input of the following fully connected layers. C-MARL was evaluated on adult brain MRI and fetal head ultrasound, outperforming single- and multi-agents approaches.

Future Work: The optimal number of agents and combination of landmarks will be further investigated. It will be also interesting to research weighted communication channels based on nearby agents to reduce noise from distant landmarks. We will incorporate more complex communication channels, e.g. skip connections and temporal units. Another direction is to investigate competitive approaches for communication instead of collaboration between the agents.

References

1. Alansary, A., et al.: Automatic view planning with multi-scale deep reinforcement learning agents. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11070, pp. 277–285. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00928-1_32
2. Alansary, A., et al.: Evaluating reinforcement learning agents for anatomical landmark detection. *Med. Image Anal.* **53**, 156–164 (2019)
3. Basher, A., et al.: Hippocampus localization using a two-stage ensemble Hough convolutional neural network. *IEEE Access* **7**, 73436–73447 (2019)
4. Bellman, R.: Dynamic programming. *Science* **153**(3731), 34–37 (1966)
5. Gauriau, R., Cuingnet, R., Lesage, D., Bloch, I.: Multi-organ localization with cascaded global-to-local regression and shape prior. *Med. Image Anal.* **23**(1), 70–83 (2015)
6. Ghesu, F.C., Georgescu, B., Mansi, T., Neumann, D., Hornegger, J., Comaniciu, D.: An artificial agent for anatomical landmark detection in medical images. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9902, pp. 229–237. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46726-9_27
7. Ghesu, F.C., et al.: Multi-scale deep reinforcement learning for real-time 3D-landmark detection in CT scans. *IEEE Trans. Pattern Anal. Mach. Intell.* **41**(1), 176–189 (2017)
8. Li, Y., et al.: Fast multiple landmark localisation using a patch-based iterative network. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11070, pp. 563–571. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00928-1_64
9. Lian, C., Liu, M., Zhang, J., Shen, D.: Hierarchical fully convolutional network for joint atrophy localization and Alzheimer’s disease diagnosis using structural MRI. *IEEE Trans. Pattern Anal. Mach. Intell.* **42**, 880–893 (2018)
10. Mnih, V., et al.: Human-level control through deep reinforcement learning. *Nature* **518**(7540), 529–533 (2015)
11. Navarro, F., Sekuboyina, A., Waldmannstetter, D., Peeken, J.C., Combs, S.E., Menze, B.H.: Deep reinforcement learning for organ localization in CT. arXiv preprint [arXiv:2005.04974](https://arxiv.org/abs/2005.04974) (2020)
12. Oktay, O., et al.: Stratified decision forests for accurate anatomical landmark localization in cardiac images. *IEEE Trans. Med. Imaging* **36**(1), 332–342 (2016)
13. Payer, C., Štern, D., Bischof, H., Urschler, M.: Regressing heatmaps for multiple landmark localization using CNNs. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9901, pp. 230–238. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46723-8_27
14. Sukhbaatar, S., Fergus, R., et al.: Learning multiagent communication with back-propagation. In: *Advances in Neural Information Processing Systems*, pp. 2244–2252 (2016)

15. Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction. MIT Press, Cambridge (2018)
16. Van Hasselt, H., Guez, A., Silver, D.: Deep reinforcement learning with double q-learning. In: Thirtieth AAAI Conference on Artificial Intelligence (2016)
17. Vlontzos, A., Alansary, A., Kamnitsas, K., Rueckert, D., Kainz, B.: Multiple landmark detection using multi-agent reinforcement learning. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11767, pp. 262–270. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32251-9_29
18. Waldmannstetter, D., et al.: Reinforced redetection of landmark in pre- and post-operative brain scan using anatomical guidance for image alignment. In: Špiclin, Ž., McClelland, J., Kybic, J., Goksel, O. (eds.) WBIR 2020. LNCS, vol. 12120, pp. 81–90. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-50120-4_8
19. Wang, Z., Schaul, T., Hessel, M., Hasselt, H., Lanctot, M., Freitas, N.: Dueling network architectures for deep reinforcement learning. In: International Conference on Machine Learning, pp. 1995–2003 (2016)
20. Watkins, C.J., Dayan, P.: Q-learning. *Mach. Learn.* **8**(3–4), 279–292 (1992)
21. Yu, C., Liu, J., Nemati, S.: Reinforcement learning in healthcare: a survey. arXiv preprint [arXiv:1908.08796](https://arxiv.org/abs/1908.08796) (2019)
22. Zhong, Z., Li, J., Zhang, Z., Jiao, Z., Gao, X.: An attention-guided deep regression model for landmark detection in cephalograms. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11769, pp. 540–548. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32226-7_60