

Deep Reinforcement Learning for L3 Slice Localization in Sarcopenia Assessment

Othmane Laousy, Guillaume Chassagnon, Edouard Oyallon, Nikos Paragios, Marie-Pierre Revel, Maria Vakalopoulou

► To cite this version:

Othmane Laousy, Guillaume Chassagnon, Edouard Oyallon, Nikos Paragios, Marie-Pierre Revel, et al.. Deep Reinforcement Learning for L3 Slice Localization in Sarcopenia Assessment. Machine Learning in Medical Imaging @MICCAI, Sep 2021, Strasbourg, France. pp.317-326, 10.1007/978-3-030-87589-3_33. hal-03532954

HAL Id: hal-03532954 https://hal.science/hal-03532954

Submitted on 18 Jan2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Deep Reinforcement Learning for L3 Slice Localization in Sarcopenia Assessment

Othmane Laousy^{1,2,3}, Guillaume Chassagnon², Edouard Oyallon⁴, Nikos Paragios⁵, Marie-Pierre Revel², and Maria Vakalopoulou^{1,3}

 1 Centrale Supélec, Paris-Saclay University, Mathématiques et Informatique pour la

Complexité et les Systèmes, Gif-sur-Yvette, France.

² Radiology Department, Hôpital Cochin AP-HP, Paris, France ³ Inria Saclay, Gif-sur-Yvette, France

⁴ Centre National de Recherche Scientifique, LIP6, Paris, France
⁵ Therapanacea, Paris, France

Abstract. Sarcopenia is a medical condition characterized by a reduction in muscle mass and function. A quantitative diagnosis technique consists of localizing the CT slice passing through the middle of the third lumbar area (L3) and segmenting muscles at this level. In this paper, we propose a deep reinforcement learning method for accurate localization of the L3 CT slice. Our method trains a reinforcement learning agent by incentivizing it to discover the right position. Specifically, a Deep Q-Network is trained to find the best policy to follow for this problem. Visualizing the training process shows that the agent mimics the scrolling of an experienced radiologist. Extensive experiments against other state-of-the-art deep learning based methods for L3 localization prove the superiority of our technique which performs well even with a limited amount of data and annotations.

Keywords: L3 slice \cdot CT slice localization \cdot Deep Reinforcement Learning \cdot Sarcopenia

1 Introduction

Sarcopenia corresponds to muscle atrophy which may be due to ageing, inactivity, or disease. The decrease of skeletal muscle is a good indicator of the overall health state of a patient [11]. In oncology, it has been shown that sarcopenia is linked to outcome in patients treated by chemotherapy [14, 5], immunotherapy [20], or surgery [9]. There are multiple definitions of sarcopenia [7, 23] and consequently multiple ways of assessing it. On CT imaging, the method used is based on muscle mass quantification. Muscle mass is most commonly assessed at a level passing through the middle of the third lumbar vertebra area (L3), which has been found to be representative of the body composition [29]. After manual selection of the correct CT slice at the L3 level, segmentation of muscles is performed to calculate the skeletal muscle area [8]. In practice, the evaluation is tedious, time-consuming, and rarely done by radiologists, highlighting 2 O. Laousy et al.

the need for an automatic diagnosis tool that could be integrated into clinical practice. Such automated measurement of muscle mass could be of great help for introducing sarcopenia assessment in daily clinical practice.

Muscle segmentation and quantification on a single slice have been thoroughly addressed in multiple works using simple 2D U-Net like architectures [4, 6]. Few works, however, focus on L3 slice detection. The main challenges for solving this task rely on the inherent diversity in patient's anatomy, the strong resemblance between vertebrae, the variability of CT fields of view as well as their acquisition and reconstruction protocols.

The most straightforward approach to address L3 localization is by investigating methods for multiple vertebrae labeling in 3D images using detection [25] or even segmentation algorithms [22]. Such methods require a substantial volume of annotations and are computationally inefficient when dealing with the entire 3D CT scan. In fact, even if our input is 3D, a one-dimensional output as the z-coordinate of the slice is sufficient to solve the L3 localization problem.

In terms of L3 slice detection, the closest methods leverage deep learning [3, 13] and focus on training simple convolutional neural networks (CNN). These techniques use maximal intensity projection (MIP), where the objective is to project voxels with maximal intensity values into a 2D plane. Frontal view MIP projections contain enough information towards the body and vertebra's bone structure differentiation. On the sagittal view, restricted MIP projections are used to focus solely on the spinal area. In [3] the authors tackle this problem through regression, training the CNN with parts of the MIP that contain the L3 vertebra only. More recently, in [13] a UNet-like architecture (L3UNet-2D) is proposed to draw a 2D confidence map over the position of the L3 slice.

In this paper, we propose a reinforcement learning algorithm for accurate detection of the L3 slice in CT scans, automatizing the process of sarcopenia assessment. The main contribution of our paper is a novel formulation for the problem of L3 localization, exploiting different deep reinforcement learning (DRL) schemes that boost the state of the art for this challenging task, even on scarce data settings. Moreover, in this paper we demonstrate that the use of 2D approaches for vertebrae detection provides state of the art results compared to 3D landmark detection methods, simplifying the problem, reducing the search space and the amount of annotations needed. To the best of our knowledge, this is the first time a reinforcement learning algorithm is explored on vertebrae slice localization, reporting performances similar to medical experts and opening new directions for this challenging task.

2 Background

Reinforcement Learning is a fundamental tool of machine learning which allows dealing efficiently with the exploration/exploitation trade-off [24]. Given state-reward pairs, a reinforcement learning agent can pick actions to reach unexplored states or increase its accumulated future reward. Those principles are appealing for medical applications because they imitate a practitioner's behavior and self-learn from experience based on ground-truth. One of the main issues of this class of algorithm is its sample complexity: a large amount of interaction with its environment is needed before obtaining an agent close to an optimal state [26]. However, those techniques were recently combined with deep learning approaches, which efficiently addressed this issue [17] by incorporating priors based on neural networks. In the context of highly-dimensional computer vision applications, this approach allowed RL algorithms to obtain outstanding accuracy [12] in a variety of tasks and applications.

In medical imaging, model-free reinforcement learning algorithms are highly used for landmark detection [10] as well as localization tasks [16]. In [1], a Deep Q-Network (DQN) that automates the view planning process on brain and cardiac MRI was proposed. This framework takes as an input a single plane and updates its angle and position during the training process until convergence. Moreover, in [27] the authors again present a DQN framework for the localization of different anatomical landmarks introducing multiple agents that act and learn simultaneously. DRL has also been studied for object or lesion localization. More recently, in [19] the authors propose a DQN framework for the localization of 6 different organs from CT scans achieving a performance comparable to supervised CNNs. This framework uses a 3D volume as input with 11 different actions to generate bounding boxes for these organs. Our work is the first to explore and validate a RL scheme on MIP representations for a single slice detection using the discrete and 2D nature of the problem.

3 Reinforcement Learning Strategy

In this paper, we formulate the slice localization problem as a *Markov Decision* Process (MDP), which contains a set of states S, actions A, and rewards R.

States S: For our formulation, the environment \mathcal{E} that we explore and exploit is a 2D image representing the frontal MIP projection of the 3D CT scans. This projection allows us to reduce our problem's dimensionality from a volume of size $512 \times 512 \times N$ (N being the varying heights of the volumes) to an image of size $512 \times N$. The reinforcement learning agent is self-taught by interacting with this environment, executing a set of actions, and receiving a reward linked to the action taken. An input example is shown in Figure 1. We define a state $s \in S$ as an image of size 512×200 in \mathcal{E} . We consider the middle of the image to be the slice's current position on a z-axis. To highlight this, we assign a line of maximum intensity pixel value to the middle of each image provided as input to our DQN.

Actions A: We define a set of discrete actions $\mathcal{A} = \{t_z^+, t_z^-\} \in \mathbb{R}^2$. t_z^+ corresponds to a positive translation (going up by one slice) and t_z^- corresponds to a negative translation (going down by one slice). These two actions allow us to explore the entirety of our environment \mathcal{E} .

Rewards R: In reinforcement learning designing a good reward function is crucial in learning the goal to achieve. To measure the quality of taking an action $a \in A$ we use the distance over z between the current slice and the annotated

slice g. The reward for non-terminating states is computed with:

$$R_a(s,s') = sign(\mathcal{D}(p,g) - \mathcal{D}(p',g)) \tag{1}$$

where we denote as s and s' the current and next state and g the ground truth annotation. Moreover, our positions p and p' are the z-coordinates of the current and next state respectively. \mathcal{D} is the Euclidean distance between both coordinates over the z-axis. The reward is non-sparse and binary $r \in \{-1, +1\}$ and helps the agent differentiate between good and bad actions. A good action being when the agent gets closer to the correct slice. For a terminating state, we assign a reward of r = 0.5.

Starting States: An episode starts by randomly sampling a slice over the z-axis and ends when the agent has achieved its goal of finding the right slice. The agent then executes a set of actions and collects rewards until the episode terminates. When reaching the upper or lower borders of an image, the current state is assigned to the next state (i.e., the agent does not move), and a reward of r = -1 is appointed to this action.

Final States: During training, a terminal state is defined as a state in which the agent has reached the right slice. A reward of r = 0.5 is assigned in this case, and an episode is terminated. During testing, the termination of an episode happens when oscillations occur. We adopted the same approach as [1], and chose actions with the lowest *Q*-value, which have been found to be closest to the right slice since the DQN outputs higher *Q*-values to actions when the current slice is far from the ground truth.

3.1 Deep Q-Learning

To find the optimal policy π^* of the MDP, a state-action value function Q(s, a) can be learned. In Q-Learning, the expected value of the accumulated discounted future rewards can be estimated recursively using the Bellman optimality equation:

$$Q_{i+1}(s,a) = \mathbb{E}[r + \gamma \max_{a'} Q_i(s',a') \mid s,a]$$
(2)

In practice, since the state S is not easily exploitable, we can take advantage of neural networks as universal function approximators to approximate Q(s, a) [18]. We utilize an experience replay technique that consists in storing the agent's experience $e_t = (s, a, r, s', a')$ at each time step in a replay memory M. To break the correlation between consecutive samples, we will uniformly batch a set of experiences from M. The Deep Q-Network (DQN) will iteratively optimize its parameters θ by minimizing the following loss function:

$$L_i(\theta_i) = \mathbb{E}_{(s,a,r,s',a')} \left[\left(r + \gamma \max_{a'} Q_{target}(s',a';\theta_i^-) - Q_{policy}(s,a;\theta_i) \right)^2 \right]$$
(3)

with θ_i and θ_i^- being the parameters of the policy and the target network respectively. To stabilize rapid policy changes due to the distribution of the data and the variations in Q-values, the DQN uses $Q_{target}(\theta_i^-)$, a fixed version of $Q_{policy}(\theta_i)$ that is updated periodically. For our experiments, we update θ_i^- every 50 iterations.



Fig. 1. The implemented Deep Q-Network architecture for L3 slice localization. The network takes as input an image of size 512×200 with a single channel. The output is the q-values corresponding to each of the two actions.

3.2 Network Architecture

Our Deep Q-Network takes as input the state s and passes it through a convolutional network. The network contains four convolution layers separated by parametric ReLU in order to break the linearity of the network, and four linear layers with LeakyReLU. Contrary to [1], we chose not to add the history of previously visited states in our case. We opted for this approach since there is a single path that leads to the right slice. This approach allows us to simplify our problem even more. Ideally, our agent should learn, just by looking at the current state, whether to go up or down when the current slice is respectively below or above the L3 slice. An overview of our framework is presented in Figure 1.

We also explore dueling DQNs from [28]. Dueling DQNs rely on the concept of an advantage which calculates the benefit that each action can provide. The advantage is defined as A(s, a) = Q(s, a) - V(s) with V(s) being our state value function. This algorithm will use the advantage of the Q-values to distinguish between actions from the state's baseline values. Dueling DQNs were shown to provide more robust agents that are wiser in choosing the next best action. For our dueling DQN, we use the same architecture as the one in Figure 1 but change the second to last fully connected layer to compute state values on one side, and action values on the other.

3.3 Training

Since our agent is unaware of the possible states and rewards in \mathcal{E} , the exploration step is implemented first. After a few iterations, our agent can start exploiting what it has learned on \mathcal{E} . In order to balance between exploration and exploitation, we use an ϵ -greedy strategy. This strategy consists of defining an exploration rate ϵ , which is initialed to 1 with a decay of 0.1, allowing the agent to become greedy and exploit the environment. A batch size of 48 and an experience replay of 17×10^3 are used. The entire framework was developed in Pytorch [21] library using an NVIDIA GTX 1080Ti GPU. We trained our model for 10^5 episodes, requiring approximately 20-24 hours. Our source code is available on GitHub: https://git.io/JRyYw.

6 O. Laousy et al.

4 Experiments and Results

4.1 Dataset

A diverse dataset of 1000 CT scans has been retrospectively collected for this study. CT scans were acquired on four different CT models from three manufacturers (Revolution HD from GE Healthcare, Milwaukee, WI; Brillance 16 from Philips Healthcare, Best, Netherlands; and Somatom AS+ & Somatom Edge from Siemens Healthineer, Erlangen, Germany). Exams were either abdominal, thoracoabdominal, or thoraco-abdominopelvic CT scans acquired with or without contrast media injection. Images were reconstructed using abdominal kernel with either filtered back-projection or iterative reconstruction. Slice thickness ranged from 0.625 to 3mm, and the number of slices varied from 121 to 1407. The heterogeneity of our dataset highlights the challenges of the problem from a clinical perspective.

Experienced radiologists manually annotated the dataset, indicating the position of the middle of the L3 slice. Before computing the MIP, all of the CT scans are normalized to 1mm over the z-axis. This normalisation step harmonises our network's input, especially since the agent performs actions along the z-axis. After the MIP, we apply a threshold of 100 HU (Hounsfield Unit) to 1500 HU allowing us to eliminate artifacts and foreign metal bodies while keeping the skeleton structure. The MIP are finally normalized to [0,1]. From the entire dataset, we randomly selected 100 patients for testing and the rest 900 for training and validation. For the testing cohort, annotations of L3 from a second experienced radiologist have been provided to measure the interobserver performance.

4.2 Results and Discussion

Our method is compared with other techniques from the literature. The error is calculated as the distance in millimeters (mm) between the predicted L3 slice and the one annotated by the experts. In particular, we performed experiments with the L3UNet-2D [13] approach and the winning SC-Net [22] method of the Verse₂₀₂₀⁶ challenge. Even if SC-Net is trained on more accurate annotations with 3D landmarks as well as vertebrae segmentations, and addresses a different problem, we applied it to our testing cohort. The comparison of the different methods is summarised in Table 1. SC-Net reports 12 CT scans with an error higher than 10mm. Moreover, L3UNet-2D [13] reports a mean error of 4.24mm ± 6.97 mm when the method is trained on the entire training set, giving only 7 scans with an error higher than 10mm for the L3 detection. Our proposed method gives the lowest errors with a mean error of 3.77mm ± 4.71 mm, proving its superiority. Finally, we evaluated our technique's performance with a Duel DQN strategy, reporting higher errors than the proposed one. This observation could be linked to the small action space that is designed for this study. Duel DQNs were proven to be powerful in cases with higher action spaces and in which the computation of the advantage function makes a difference.

⁶ https://verse2020.grand-challenge.org/

Method	# of samples	Mean	Std	Median	Max	Error > 10mm
Interobserver	-	2.04	4.36	1.30	43.19	1
SC-Net [22]	-	6.78	13.96	1.77	46.98	12
L3UNet-2D [13]	900	4.24	6.97	2.19	40	7
Ours (Duel-DQN)	900	4.30	5.59	3	38	8
Ours	900	3.77	4.71	2.0	24	9
L3UNet-2D [13]	100	145.37	161.91	32.8	493	68
Ours	100	5.65	5.83	4	26	19
L3UNet-2D [13]	50	108.7	97.33	87.35	392.02	86
Ours	50	6.88	5.79	6.5	26	11
L3UNet-2D [13]	10	242.85	73.07	240.5	462	99
Ours	10	8.97	8.72	7	56	33

Table 1. Quantitative evaluation of the different methods using different number of training samples (metrics in mm).

For the proposed reinforcement learning framework, trained on the whole training set, 9 CTs had a detection error of more than 10mm. These scans were analysed by a medical expert who indicated that 2 of them have a lumbosacral transitional vertebrae (LSTV) anomaly [15]. Transitional vertebrae cases are common and observed in 15-35% of the population [2] highlighting once again the challenges of this task. For both cases, the localization of the L3 vertebra for sarcopenia assessment is ambiguous for radiologists and consequently for the reinforcement learning agent. In fact, the only error higher than 10mm in the interobserver comparison corresponds to an LSTV case where each radiologist chose a different vertebrae as a basis for sarcopenia assessment. Even if the interobserver performance is better than the one reported by the algorithms, our method reports the lowest errors, proving its potential.

Qualitative results are displayed in Figure 2. The yellow line represents the medical expert's annotation and the blue one the prediction of the different employed models. One can notice that all of the different methods converge to the correct L3 region with our method reporting great performance. It is important to note that for sarcopenia assessment, an automatic system does not need to be at the exact middle of the slice; a few millimeters around will not skew the end result since muscle mass in the L3 zone does not change significantly. Concerning prediction times for the RL agent, they depend on the initial slice that is randomly sampled on the MIP. Computed inference time for a single step is approximately 0.03 seconds.

To highlight the robustness of our network on a low number of annotated samples, we performed different experiments using 100, 50, and 10 CTs corresponding respectively to 10%, 5% and 1% of our dataset. We tested those 3 agents on the same 100 patients test set and report results in Table 1. Our experiments prove the robustness of reinforcement learning algorithms compared to traditional CNN based ones [13] in the case of small annotated datasets. One



Fig. 2. Qualitative comparison of different localization methods for two patients. First left to right represents: interobserver (4mm/2mm), SC-Net (4mm/2mm), L3UNet-2D (8mm/1mm), Ours (1mm/4mm). In the parenthesis we present the reported errors for the first and second row respectively. The yellow line represents the ground truth and the blue one the prediction.

can observe that the traditional methods fail to be trained properly with a small number of annotations, reporting errors higher than 100mm for all three experiments. In our case, decreasing the dataset size does not significantly affect the performance. In fact, trained on only 10 CTs with the same number of iterations and memory size, our agent was able to learn a correct policy and achieve a mean error of 8.97mm \pm 8.72mm. Learning a valid policy from a low number of annotations is one of the strengths of reinforcement learning. Traditional deep learning techniques rely on pairs of images and annotations in order to build a robust generalization. Thus, each pair is exploited only once by the learning algorithm. Reinforcement learning, however, relies on experiences, each experience $e_t = (s, a, r, s', a')$ being a tuple of state, action, reward, next state and next action. Therefore, a single CT scan can provide multiple experiences to the self-learning agent, making our method ideal for slice localization problems using datasets with a limited amount of annotations.

5 Conclusion

In this paper, we propose a novel direction to address the problem of CT slice localization. Our experiments empirically prove that reinforcement learning schemes work very well on small datasets and boost performance compared to classical convolutional architectures. One limitation of our work lies in the fact that our agent is always moving 1mm independently of the location, slowing down the process. In the future, we aim to explore different ways to adapt the action taken depending on the current location, with one possibility being to incentivize actions with higher increments. Future work also includes the use of reinforcement learning in multiple vertebrae detection with competitive or collaborative agents.

References

- Alansary, A., Le Folgoc, L., Vaillant, G., Oktay, O., Li, Y., Bai, W., Passerat-Palmbach, J., Guerrero, R., Kamnitsas, K., Hou, B., et al.: Automatic view planning with multi-scale deep reinforcement learning agents. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 277–285. Springer (2018)
- 2. Apazidis, A., Ricart, P.A., Diefenbach, C.M., Spivak, J.M.: The prevalence of transitional vertebrae in the lumbar spine. The Spine Journal **11**(9), 858–862 (2011)
- Belharbi, S., Chatelain, C., Hérault, R., Adam, S., Thureau, S., Chastan, M., Modzelewski, R.: Spotting l3 slice in ct scans using deep convolutional network and transfer learning. Computers in biology and medicine 87, 95–103 (2017)
- Blanc-Durand, P., Schiratti, J.B., Schutte, K., Jehanno, P., Herent, P., Pigneur, F., Lucidarme, O., Benaceur, Y., Sadate, A., Luciani, A., Ernst, O., Rouchaud, A., Creze, M., Dallongeville, A., et al.: Abdominal musculature segmentation and surface prediction from ct using deep learning for sarcopenia assessment. Diagnostic and Interventional Imaging **101**(12), 789–794 (2020)
- 5. Bozzetti, F.: Forcing the vicious circle: sarcopenia increases toxicity, decreases response to chemotherapy and worsens with chemotherapy. Annals of Oncology **28**(9), 2107–2118 (2017), a focus on esophageal squamous cell carinoma
- Castiglione, J., Somasundaram, E., Gilligan, L.A., Trout, A.T., Brady, S.: Automated segmentation of abdominal skeletal muscle in pediatric ct scans using deep learning. Radiology: Artificial Intelligence p. e200130 (2021)
- Cruz-Jentoft, A.J., Bahat, G., Bauer, J., Boirie, Y., Bruyère, O., Cederholm, T., Cooper, C., Landi, F., Rolland, Y., Sayer, A.A., et al.: Sarcopenia: revised european consensus on definition and diagnosis. Age and ageing 48(1), 16–31 (2019)
- Derstine, B.A., Holcombe, S.A., Goulson, R.L., Ross, B.E., Wang, N.C., Sullivan, J.A., Su, G.L., Wang, S.C.: Quantifying Sarcopenia Reference Values Using Lumbar and Thoracic Muscle Areas in a Healthy Population. J Nutr Health Aging 21(10), 180–185 (2017)
- Du, Y., Karvellas, C.J., Baracos, V., Williams, D.C., Khadaroo, R.G.: Sarcopenia is a predictor of outcomes in very elderly patients undergoing emergency surgery. Surgery 156(3), 521–527 (2014)
- Ghesu, F.C., Georgescu, B., Grbic, S., Maier, A., Hornegger, J., Comaniciu, D.: Towards intelligent robust detection of anatomical structures in incomplete volumetric data. Med Image Anal 48, 203–213 (08 2018)
- Gilligan, L.A., Towbin, A.J., Dillman, J.R., Somasundaram, E., Trout, A.T.: Quantification of skeletal muscle mass: sarcopenia as a marker of overall health in children and adults. Pediatric Radiology 50(4), 455–464 (Apr 2020)
- 12. Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Doersch, C., Pires, B.A., Guo, Z.D., Azar, M.G., et al.: Bootstrap your own latent: A new approach to self-supervised learning. arXiv preprint arXiv:2006.07733 (2020)
- Kanavati, F., Islam, S., Aboagye, E.O., Rockall, A.: Automatic 13 slice detection in 3d ct images using fully-convolutional networks (2018)
- Lee, J., Chang, C.L., Lin, J.B., Wu, M.H., Sun, F.J., Jan, Y.T., Hsu, S.M., Chen, Y.J.: Skeletal muscle loss is an imaging biomarker of outcome after definitive chemoradiotherapy for locally advanced cervical cancer. Clinical Cancer Research 24(20), 5028–5036 (2018)
- Lian, J., Levine, N., Cho, W.: A review of lumbosacral transitional vertebrae and associated vertebral numeration. European Spine Journal 27(5), 995–1004 (May 2018)

- 10 O. Laousy et al.
- Maicas, G., Carneiro, G., Bradley, A.P., Nascimento, J.C., Reid, I.: Deep reinforcement learning for active breast lesion detection from dce-mri. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) Medical Image Computing and Computer Assisted Intervention MICCAI 2017. pp. 665–673. Springer International Publishing, Cham (2017)
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., Riedmiller, M.: Playing atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602 (2013)
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., Hassabis, D.: Human-level control through deep reinforcement learning. Nature **518**(7540), 529–533 (Feb 2015)
- Navarro, F., Sekuboyina, A., Waldmannstetter, D., Peeken, J.C., Combs, S.E., Menze, B.H.: Deep reinforcement learning for organ localization in ct. In: Medical Imaging with Deep Learning. pp. 544–554. PMLR (2020)
- Nishioka, N., Uchino, J., Hirai, S., Katayama, Y., Yoshimura, A., Okura, N., Tanimura, K., Harita, S., Imabayashi, T., Chihara, Y., Tamiya, N., Kaneko, Y., Yamada, T., Takayama, K.: Association of sarcopenia with and efficacy of anti-pd-1/pd-l1 therapy in non-small-cell lung cancer. Journal of Clinical Medicine 8(4) (2019)
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in PyTorch. In: NIPS Autodiff Workshop (2017)
- 22. Payer, C., Stern, D., Bischof, H., Urschler, M.: Coarse to fine vertebrae localization and segmentation with spatialconfiguration-net and u-net. In: Proceedings of the 15th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications - Volume 5: VISAPP. vol. 5, pp. 124–133 (2020). https://doi.org/10.5220/0008975201240133
- Santilli, V., Bernetti, A., Mangone, M., Paoloni, M.: Clinical definition of sarcopenia. Clinical cases in mineral and bone metabolism 11(3), 177 (2014)
- Sutton, R.S., Barto, A.G.: Reinforcement learning: An introduction. MIT press (2018)
- Suzani, A., Seitel, A., Liu, Y., Fels, S., Rohling, R.N., Abolmaesumi, P.: Fast automatic vertebrae detection and localization in pathological ct scans - a deep learning approach. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015. pp. 678–686. Springer International Publishing (2015)
- Tarbouriech, J., Garcelon, E., Valko, M., Pirotta, M., Lazaric, A.: No-regret exploration in goal-oriented reinforcement learning. In: International Conference on Machine Learning. pp. 9428–9437. PMLR (2020)
- Vlontzos, A., Alansary, A., Kamnitsas, K., Rueckert, D., Kainz, B.: Multiple landmark detection using multi-agent reinforcement learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 262–270. Springer (2019)
- Wang, Z., de Freitas, N., Lanctot, M.: Dueling network architectures for deep reinforcement learning. CoRR abs/1511.06581 (2015), http://arxiv.org/abs/1511.06581
- Zopfs, D., Theurich, S., Große Hokamp, N., Naetlitz, J., Gerecht, L., Borggrefe, J., Schlaak, M., Pinto dos Santos, D.: Single-slice ct measurements allow for accurate assessment of sarcopenia and body composition. European radiology **30** (11 2019)