

Goal-conditioned reinforcement learning for ultrasound navigation guidance

Abdoul Aziz Amadou^{1,2}, Vivek Singh³, Florin C. Ghesu⁴, Young-Ho Kim³, Laura Stanciulescu^{3,5}, Harshitha P. Sai^{6,7}, Puneet Sharma³, Alistair Young¹, Ronak Rajani^{1,8}, and Kawal Rhode¹

¹ King's College London, School of Biomedical Engineering & Imaging Sciences, London, United Kingdom
`abdoul.a.amadou@kcl.ac.uk`

² Siemens Healthcare Limited, Camberley, United Kingdom

³ Siemens Healthineers, Digital Technology and Innovation, Princeton, NJ, USA

⁴ Siemens Healthineers AG, Digital Technology and Innovation, Erlangen, Germany

⁵ Carol Davila University of Medicine and Pharmacy Bucharest, Romania

⁶ Siemens Healthineers, Sonographer, Bangalore, Karnataka

⁷ Fortis Institute of Medical Sciences, Affiliated by Rajiv Gandhi University of Medical Sciences, Department of Cardiology, Bangalore, Karnataka

⁸ Guy's and St Thomas' NHS Foundation Trust, London, United Kingdom

Abstract. Transesophageal echocardiography (TEE) plays a pivotal role in cardiology for diagnostic and interventional procedures. However, using it effectively requires extensive training due to the intricate nature of image acquisition and interpretation. To enhance the efficiency of novice sonographers and reduce variability in scan acquisitions, we propose a novel ultrasound (US) navigation assistance method based on contrastive learning as goal-conditioned reinforcement learning (GCRL). We augment the previous framework using a novel contrastive patient batching method (CPB) and a data-augmented contrastive loss, both of which we demonstrate are essential to ensure generalization to anatomical variations across patients. The proposed framework enables navigation to both standard diagnostic as well as intricate interventional views with a single model. Our method was developed with a large dataset of 789 patients and obtained an average error of 6.56 mm in position and 9.36 degrees in angle on a testing dataset of 140 patients, which is competitive or superior to models trained on individual views. Furthermore, we quantitatively validate our method's ability to navigate to interventional views such as the Left Atrial Appendage (LAA) view used in LAA closure. Our approach holds promise in providing valuable guidance during transesophageal ultrasound examinations, contributing to the advancement of skill acquisition for cardiac ultrasound practitioners.

Keywords: Ultrasound · Echocardiography · Deep reinforcement learning · Goal-conditioned reinforcement learning

1 Introduction

Echocardiography is a key imaging modality in the diagnosis and treatment of cardiovascular diseases. While several US modalities are used in practice, in TEE, the transducer images the heart from the oesophagus, often yielding better scan quality and helping circumvent issues caused by acoustic windows in other modalities such as transthoracic echocardiography (TTE). Training operators for TEE is time-consuming due to complex controls and image interpretation, with an added risk of patient injury due to incorrect transducer manipulation. Additionally, in structural heart procedures where TEE is coupled with fluoroscopy, health issues arise for catheterization lab staff due to orthopaedic strain and radiation exposure [1].

AI-assisted image acquisition guidance for transducer manipulation has been proven to benefit operator training, lower the learning curve, and reduce intra and inter-user variability [2,3]. Additional advantages include shortening of TEE examinations, enhancing patient comfort and reducing radiation exposure during interventional procedures.

Various deep reinforcement learning (DRL) approaches for ultrasound autonomous navigation have been proposed, primarily focusing on extracorporeal scanning of anatomies like the spine [4,5] and neck [6]. However, Li et al. [4] suffer from a lack of generalization to unseen patient datasets, and [5] employs simplified state and action spaces that do not capture the real-world scanning conditions well. While previous works rely on simulation environments, both [7,8] use additional hardware attached to the transducers to acquire datasets for imitation learning. However, the scalability of the data acquisition (time and cost) is reported as one of the limitations of such approaches [8]. TEE imaging has been less explored, with Wang et al. [9] using a simulation environment based on segmented pre-operative scans to find robotic poses corresponding to desired views pre-operatively. However, this approach requires a manual intervention to define the views. Finally, authors in Li et al. [10] use a simulation environment to train models to navigate to standard TEE views. However, their approach involves training one model for each target view, which does not scale well to support additional views or supporting manoeuvres to visualize specific structures more clearly. Furthermore, they only control 3 out of 5 transducer degrees of freedom and test on a limited dataset of 5 patients.

This paper introduces a novel approach to training a navigation model using goal-conditioned reinforcement learning. We build upon Contrastive RL (CRL) [11], a state-of-the-art goal-conditioned method which showed promising results in image-based robotic tasks. We train our model using random goal views, enabling navigation to arbitrary views given a user-defined goal. We make use of a simulation environment [12], where we leverage a large dataset of chest and cardiac CTs to train our model and enable generalization to unseen patients. An overview of the proposed workflow is shown in Fig. 1.

The contributions of this work are the following: **1)** We propose a novel methodology for TEE imaging guidance to arbitrary views using goal-conditioned reinforcement learning. This not only enables navigation to standard views but




Fig. 1: *System overview of Goal-conditioned RL for Ultrasound Navigation.* We first segment CTs and generate ultrasound volume reconstructions for rapid sampling during training. The model is trained to reach randomly selected goal views by employing the contrastive patient batching (CPB) mechanism to create a contrastive batch from the collected experience. When deployed, the trained model can navigate to arbitrary views, including standard and interventional views.

also to alternative views showing specific structures. **2)** We enable the generalization of the CRL framework by introducing: *(i)* Contrastive patient batching, a simple yet effective method to sample hard contrastive pairs and improve performance; *(ii)* A novel contrastive data augmentation loss to improve both robustness and the quality of learnt representations. **3)** We demonstrate the effectiveness of our approach by performing two experiments on a dataset of 140 patients: *(i)* By navigating to standard views, including views that were not explicitly sampled during training. Our method achieves competitive performance to RL methods trained to reach individual views; *(ii)* By navigating to a non-standard view used to monitor the deployment of devices in LAA closure. This showcases the usability of our method both for diagnostic and interventional cases. To the best of our knowledge, this work is the first attempt to develop an ultrasound navigation model capable of navigating to arbitrary views given a goal.

2 Methodology

2.1 Simulation environment

Acquiring real datasets for navigation is a cumbersome, expensive and time-consuming task, as reported in [7,8]. Hence, as shown in Fig. 1, we employ a physics-based Computed Tomography (CT) to ultrasound simulation pipeline to train our model. [12]. The pipeline takes as input chest and cardiac CTs and automatically segments them to obtain masks of the organs of interest, namely the oesophagus, heart chambers, aorta, lungs and pulmonary artery. A Monte Carlo path tracing algorithm is then used to simulate ultrasound wave propagation in tissue. The pipeline was extensively validated with phantom experiments,




Fig. 2: *Contrastive critic training*: We build a contrastive batch using trajectories from two patients with CPB and pass (observation, action) pairs and goal images to the state-action and goal encoders respectively. Goal and observations are augmented K times and we build K^2 intermediate matrices (not shown) from the inner product between all the encoded representations, with Q_M^{Aug} as their average. The critic is trained to maximize the similarity between state-action and goal representations of the same trajectories, which corresponds to the diagonal of the matrices.

where US image properties were assessed, and a view classification experiment, in which we demonstrated the usefulness of the pipeline in generating data for model training. More details on the pre-processing and US simulation pipeline are provided in the anonymized submission in the supplementary material.

As simulating ultrasound images on the fly from the CT is computationally expensive and would significantly slow down the training, we followed Li et al. [10] and generated simulated US images by translating the transducer down the oesophagus, rotating it by 360 degrees at every position. Simulation and volume reconstruction were done offline on the GPU.

2.2 Goal-Conditioned Reinforcement Learning

The goal-conditioned navigation task is defined by: S , which represents the environment’s state, defined by the transducer’s pose in the CT coordinate system; A is the set of TEE transducer movements, i.e. translation along the oesophagus, transducer rotation, the electronic rotation of the scanning plane and the left/right and retro/ante flexions; $p(s_{t+1}|s_t, a_t)$ are the transition probabilities between s_{t+1} and s_t after taking action a_t ; $r_g(s, a) = (1 - \gamma)p(s_{t+1} = s_g|s_t, a_t)$ is the goal-conditioned reward function, defined as the probability density of reaching the goal s_g at the next step; Ω is the set of observations o , which correspond to ultrasound images acquired from the transducer in a given state s ; $\gamma \in [0, 1]$ the discount factor.

Our goal-conditioned framework follows the actor-critic architecture, where the critic takes as input a triplet of (observation, action, goal) (o_t, a_t, o_g) and returns the probability (density) of reaching goal o_g when taking action a_t when given an observation o_t . The actor takes as input a pair (o_t, o_g) and returns the action a_t to take to reach the goal. The critic is trained to correctly predict which actions lead to a goal, and the actor learns to output correct actions by maximizing the critic’s output.

Similarly to [11], contrastive learning is used to train a critic function by making use of two models ϕ and ψ which encode state-action (SA) pairs (o_t, a_t) and goals o_g respectively. The critic function measures the similarity of the latent representations of dimension H via inner-product $f(o_t, a_t, o_g) = \langle \phi(o_t, a_t), \psi(o_g) \rangle$, as illustrated in Fig. 2. When an action likely leads to a goal from a given pose, the inner product will have a high value, indicating the probability (density) of reaching the goal is high.

During training, the transducer is initialized at a given pose s_0 (yielding observation o_0) and is given a goal observation o_g . The sequence of observations/actions until the last timestep gives a trajectory $\tau_i = (o_0^i, a_0^i, o_1^i, \dots, o_n^i)$.

Critic loss: To train the critic, as illustrated in Fig. 2, we sampled the input triplet from a trajectory $(o_t^i, a_t^i, o_g^i) \sim \tau_i$, where the positive goal timestep T is a future timestep ($T > t$) sampled from a geometric distribution $T \sim \text{Geom}(1-\gamma)$. The negative goal $o_{g'}^j$ is sampled randomly from another trajectory τ_j . The critic loss is based on the infoNCE loss [13] and computed as:

$$\max_f \mathbb{E}_{\substack{(o_t^i, a_t^i, o_g^i) \sim \tau_i \\ (o_{g'}^j) \sim \tau_j}} \log \left[\frac{e^{f(o_t^i, a_t^i, o_g^i)^+}}{e^{f(o_t^i, a_t^i, o_g^i)^+} + \sum_j e^{f(o_t^i, a_t^i, o_{g'}^j)^-}} \right] \quad (1)$$

Where $f(o_t, a_t, o_g)^+$ and, $f(o_t, a_t, o_g)^-$ denote the critic output for positive and negative examples respectively. The inner product between all SA and goal representations in a batch of size N gives a matrix Q_M of size (N, N) on which we apply a cross-entropy loss row and column-wise, with the true labels being on the diagonal. In order to stabilize the critic training, we observed that the normalization of the goal representations was necessary. Furthermore, the use of a temperature scaling parameter of state-action representations, combined with L2 regularization was necessary. Hyperparameters are listed in the supplementary material.

Actor loss: The actor takes as input observation and goal pairs (o_t, o_g) and returns an action a_t to reach the goal. The actor simply aims at maximizing the critic output such that:

$$\max_{\pi} \mathbb{E}_{a_t \sim \pi(o_t, o_g)} f(o_t, a_t, o_g) \quad (2)$$

In practice, the actor outputs the mean and standard deviation of a multivariate Gaussian from which we sample and apply tanh squashing to obtain the bounded actions. During training, we noticed that using random goals sampled from other trajectories rather than goals from the same trajectory as o_t led to better actor performance, as also reported in [11].

Data augmented contrastive loss: Given a triplet (o_t, a_t, o_g) , the critic output should be similar to the data augmented triplet (o'_t, a_t, o'_g) , where o'_t is a randomly shifted version of o_t . We apply K random shifts to the observations and goal images, where the k -th augmentation is denoted as $o_{t,k}$. The critic loss is then computed on the average of the K^2 matrices Q_M^i resulting from the inner products of the augmented observations and goals.

$$Q_M^{aug} = \mathbb{E}_i[Q_M^i] = \frac{1}{K^2} \sum_{k=1}^K \sum_{k'=1}^K f(o_{t,k}, a_t, o_{g,k'}) \quad (3)$$

Contrastive Patient Batching (CPB): We observed empirically that the composition of the contrastive batch plays a significant role in the convergence of the critic. Following the strategy proposed in [11], where samples are randomly chosen from the replay buffer yields poor results in our setting. A closer investigation revealed that a randomly sampled batch contains samples from different patients with different intermediate states, and the critic ends up learning features associated with anatomical differences between patients, rather than general anatomical features necessary for the control task. To address this, we tag the trajectories by the corresponding patient identifier during training. While creating a batch of size N , we sample (o_t, a_t, o_g) triplets from two patients, with $\frac{N}{2}$ samples per patient. Having a significant number of samples coming from the same patient creates harder negatives for the critic, which improves its effectiveness in discriminating trajectories. Ablation studies in the supplementary material show performance for different numbers of patients per batch.

Training loop: We automatically find probe poses to obtain standard views using landmarks extracted from the automatic segmentations and by following clinical guidelines [15]. When applying actions, the transducer is translated along the oesophagus centerline and we constrain all motions to remain within its walls. At the start of each episode, we initialize the transducer at one of the standard view poses. Random perturbations are applied to obtain the starting pose s_0 . For CRL, we obtain a goal pose by applying additional random perturbations from s_0 , yielding a goal pose s_g . Hence our model is always trained with random goals and never explicitly trained to navigate to a standard view. Perturbation ranges are listed in the supplementary material.

3 Experiments and results

Standard view navigation: We first compare our approach with existing methods by examining their ability to reach standard views, a task essential in TEE examinations. We processed 929 patient CT datasets from the LIDC-IDRI dataset [16]: 653 were used for training, 136 for validation and the rest 140 for testing. Goal images in the test dataset were reviewed and confirmed by a cardiologist. We evaluated our method in two scenarios based on how the goal is specified: Using the (synthetic) US view generated from the same patient as the goal or using a US view from another "template" patient as a goal. The latter

Table 1: Quantitative results (mean $\pm$ std) for the standard view navigation experiment. Goal type *Patient/Template* indicates whether the input goal was generated from the same patient or from a template patient. Note that no perturbations were explicitly sampled around the ME 5CH view during training. (*) Results for RL-TEE and SAC are obtained from several models, each one trained separately on a view. CRL+B indicates CRL-D trained with CPB and CRL + BA is CRL+B with the data augmented contrastive loss ($K = 2$).

Views	Goal type	Method	Angle Error (deg)	Position error (mm)
ME AV SAX, 2CH, 4CH, LAX	N/A	RL-TEE [10]*	9.90 ± 8.04	9.17 ± 6.87
		SAC* [14]	9.77 ± 10.89	7.92 ± 9.35
	Patient	CRL-D [11]	18.47 ± 19.89	13.27 ± 17.96
		CRL+B	9.00 ± 12.37	5.93 ± 8.17
		CRL+BA	9.36 ± 9.52	6.56 ± 6.46
ME 5CH	Patient	CRL-D [11]	35.24 ± 25.67	15.23 ± 15.08
		CRL+B	19.88 ± 35.91	8.88 ± 11.02
		CRL+BA	11.40 ± 5.30	7.80 ± 4.11
ME 2CH,4CH	Template	CRL-D [11]	28.91 ± 21.86	24.59 ± 24.63
		CRL+B	13.39 ± 7.74	12.95 ± 7.32
		CRL+BA	12.19 ± 6.88	10.54 ± 6.34

corresponds to a scenario where a user may not have access to prior scans of the patient and hence uses a similar view from another reference to specify the goal. We report the position and angle error at the end of the episode w.r.t the ground truth pose. For each patient and goal pair, we ran 10 experiments with the transducer initialized at random positions. Results are reported in Table 1 and a breakdown of the results per view is included in the supplementary material, alongside videos showing the navigation process.

We compare the performance of our model with Li et al. [10] (RL-TEE), Soft Actor-Critic (SAC) [14] which is a state-of-the-art off-policy reinforcement learning algorithm and the default implementation of CRL (CRL-D)[11]. For RL-TEE and SAC, we train one model per view as the algorithms are not goal-conditioned. SAC models are trained using the same rewards as in [10]. We use four mid-oesophageal (ME) views for training: Two and four chambers (2CH, 4CH), long-axis (LAX) and aortic valve short-axis (AV SAX). Due to the similarity between ME AV SAX and ME Right Ventricle Inflow-Outflow views, samples resembling one or the other class were considered to be of the AV SAX class. We use templates from ME 2CH and 4CH views as they are better geometrically defined across patients.

Finally, we showcase the versatility of our method by inputting ME five chambers (5CH) views as goals to the model during testing. ME 5CH views were not used as a starting point for random perturbations during training. In Table 1, CRL+B model corresponds to CRL + CPB, and CRL+BA is CRL + CPB + data augmented contrastive loss.

Table 2: Quantitative results for the LAA view navigation experiment. The high-quality representations learnt by the model with the data augmented contrastive loss allow for better generalization and transfer.

Views	Goal type	Method	Angle Error (deg)	Position error (mm)
LAA	Patient	CRL-D [11]	37.70 ± 28.74	30.99 ± 28.61
		CRL+B	24.63 ± 23.69	17.54 ± 17.22
		CRL+BA	10.18 ± 5.58	9.02 ± 4.33

Interventional view navigation: In a second experiment, we showcase the usefulness of goal-conditioning by navigating to a non-standard view used in LAA closure procedures. We use the FUMPE dataset [17] (train: 21 / test: 5) for which we have additional LAA segmentations, hence finetuning is required as the LAA was missing in the previous datasets. Previously trained CRL models are finetuned for 250K steps, without changing the training procedure or sampling trajectories near the LAA explicitly. Quantitative results are reported in Table 2, where the performance is on par with standard view navigation. Qualitative results are shown in Fig. 3.

For all experiments and models, we use a ResNet-18 [18] as image an encoder, Adam optimizer [19] and train with A4500 GPUs. Detailed result tables and demonstration videos are included in the supplementary material.




Fig. 3: Example navigation to a view showing the LAA (orange box). The two rightmost pictures are projections showing the desired (green) and current (red) transducer positions.

4 Discussion and conclusion

Discussion: Our generalist model achieves competitive performance to specialist models trained to navigate to single views, whether it is given goal images from the same or a template patient. Additionally, the model robustly navigates to arbitrary views without explicit sampling during training, as shown by

the results on ME 5CH and LAA views, thus demonstrating the versatility of the goal-conditioned framework. Note that the performance in such scenarios is highly dependent on the agent’s exploration of the environment during training. A drawback of CRL is the longer training time, as the contrastive critic needs many samples to converge. We alleviate this with an efficient asynchronous implementation using RLLib [20], yielding a training time of two days for 200M steps. Finally, deployment in a real-world setting would potentially require fine-tuning using either real data and/or improved simulations with generative models to address any reality gap.

Conclusion: We have presented a novel approach for ultrasound navigation using goal-conditioned reinforcement learning. Given a goal image, our versatile model navigates robustly both to standard and arbitrary views showing specific structures. Using this method as an image acquisition guidance system could help train sonographers, improve the acquisition quality and reduce variability among experienced users.

Acknowledgements. The authors acknowledge the National Cancer Institute and the Foundation for the National Institutes of Health, and their critical role in the creation of the free publicly available LIDC/IDRI Database used in this study. For the purpose of open access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission

Disclaimer. The concepts and information presented in this paper are based on research results that are not commercially available. Future commercial availability cannot be guaranteed.

Disclosure of Interests. This research was funded in part, by the Wellcome Trust under Grant WT203148/Z/16/Z and in part by Siemens Healthineers.

References

1. Andreassi, M.G., Piccaluga, E., Guagliumi, G., Greco, M.D., Gaita, F., Picano, E.: Occupational health risks in cardiac catheterization laboratory workers. *Circulation: Cardiovascular Interventions* **9**, e003273 (2016)
2. Narang, A., Bae, R., Hong, H., Thomas, Y., Surette, S., Cadieu, C.F., Chaudhry, A.K., Martin, R.P., McCarthy, P.M., Rubenson, D., Goldstein, S.A., Little, S.H., Lang, R.M., Weissman, N., Thomas, J.D.: Utility of a deep-learning algorithm to guide novices to acquire echocardiograms for limited diagnostic use. *JAMA Cardiology* **6**, 1 – 9 (2021)
3. Sabo, S., Padeloup, D., Pettersen, H.N., Smistad, E., Østvik, A., Olaisen, S.H., Stølen, S.B., Grenne, B.L., Holte, E., Lovstakken, L., Dalen, H.: Real-time guidance by deep learning of experienced operators to improve the standardization of echocardiographic acquisitions. *European Heart Journal - Imaging Methods and Practice* **1**(2), qyad040 (2023)
4. Li, K., Wang, J., Xu, Y., Qin, H., Liu, D., Liu, L., Meng, M.Q.: Autonomous navigation of an ultrasound probe towards standard scan planes with deep reinforcement learning. 2021 IEEE International Conference on Robotics and Automation (ICRA) pp. 8302–8308 (2021)
5. Hase, H., Azampour, M.F., Tirindelli, M., Paschali, M., Simson, W., Fatemizadeh, E., Navab, N.: Ultrasound-guided robotic navigation with deep reinforcement learning. 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) pp. 5534–5541 (2020)
6. Bi, Y., Jiang, Z., Gao, Y., Wendler, T., Karlas, A., Navab, N.: Vesnet-rl: Simulation-based reinforcement learning for real-world us probe navigation. *IEEE Robotics and Automation Letters* **7**, 6638–6645 (2022)
7. Droste, R., Drukker, L., Papageorgiou, A.T., Noble, J.A.: Automatic probe movement guidance for freehand obstetric ultrasound. *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)* **12263**, 583–592 (2020)
8. Milletari, F., Birodar, V., Sofka, M.: Straight to the Point: Reinforcement Learning for User Guidance in Ultrasound. In: *Smart Ultrasound Imaging and Perinatal, Preterm and Paediatric Image Analysis*. pp. 3–10. Springer International Publishing, Cham (2019)
9. Wang, S., Housden, J., Bai, T., Liu, H., Back, J., Singh, D., Rhode, K.S., Hou, Z.G., Wang, F.Y.: Robotic intra-operative ultrasound: Virtual environments and parallel systems. *IEEE/CAA Journal of Automatica Sinica* **8**, 1095–1106 (2021)
10. Li, K., Li, A., Xu, Y., Xiong, H., Meng, M.Q.H.: Rl-tee: Autonomous probe guidance for transesophageal echocardiography based on attention-augmented deep reinforcement learning. *IEEE Transactions on Automation Science and Engineering* (2023)
11. Eysenbach, B., Zhang, T., Salakhutdinov, R., Levine, S.: Contrastive learning as goal-conditioned reinforcement learning. In: *Neural Information Processing Systems* (2022)
12. Amadou, A.A., Peralta, L., Dryburgh, P., Klein, P., Petkov, K., Housden, R.J., Singh, V., Liao, R., Kim, Y.H., Ghesu, F.C., Mansi, T., Rajani, R., Young, A., Rhode, K.: Cardiac ultrasound simulation for autonomous ultrasound navigation. *arXiv preprint arXiv:2402.06463* (2024)
13. van den Oord, A., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. *CoRR* **abs/1807.03748** (2018)

14. Haarnoja, T., Zhou, A., Abbeel, P., Levine, S.: Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. ICML (2018)
15. Hahn, R.T., Abraham, T., Adams, M.S., Bruce, C.J., Glas, K.E., Lang, R.M., Reeves, S.T., Shanewise, J.S., Siu, S.C., Stewart, W., Picard, M.H.: Guidelines for performing a comprehensive transesophageal echocardiographic examination: Recommendations from the american society of echocardiography and the society of cardiovascular anesthesiologists. *Journal of the American Society of Echocardiography* **26**(9), 921–964 (2013)
16. Armato, S.G., McNitt-Gray, M.F.: The lung image database consortium (lidc) and image database resource initiative (idri): a completed reference database of lung nodules on ct scans. *Medical physics* **38** **2**, 915–31 (2011). <https://doi.org/10.1118/1.3528204>
17. Masoudi, M., Pourreza, H.R., Saadatmand-Tarzjan, M., Eftekhari, N., Zargar, F.S., Rad, M.P.: A new dataset of computed-tomography angiography images for computer-aided detection of pulmonary embolism. *Scientific Data* **5** (2018). <https://doi.org/10.6084/m9.figshare.c.4107803.v1>
18. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) pp. 770–778 (2015), <https://api.semanticscholar.org/CorpusID:206594692>
19. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *CoRR abs/1412.6980* (2014)
20. Liang, E., Liaw, R., Nishihara, R., Moritz, P., Fox, R., Gonzalez, J., Goldberg, K., Stoica, I.: Ray rllib: A composable and scalable reinforcement learning library. *CoRR abs/1712.09381* (2017)