Towards More Sample Efficiency in Reinforcement Learning with Data Augmentation

Yijiong Lin¹, Jiancong Huang¹, Matthieu Zimmer², Yisheng Guan¹, Juan Rojas¹, and Paul Weng²

¹Guangdong University of Technology ²Shanghai Jiao Tong University

Abstract

Deep reinforcement learning (DRL) is a promising approach for adaptive robot control, but its current application to robotics is currently hindered by high sample requirements. We propose two novel data augmentation techniques for DRL in order to reuse more efficiently observed data. The first one called Kaleido-scope Experience Replay exploits reflectional symmetries, while the second called Goal-augmented Experience Replay takes advantage of lax goal definitions. Our preliminary experimental results show a large increase in learning speed.

1 Introduction

Deep reinforcement learning (DRL) has demonstrated great promise in recent years [Mnih *et al.*, 2015; Silver *et al.*, 2016]. However, despite being shown to be a viable approach in robotics [Kalashnikov *et al.*, 2018; OpenAI *et al.*, 2018], DRL still suffers from high sample complexity in practice—an acute issue in robot learning.

Given how critical this issue is, many diverse propositions have been presented. For brevity, we only recall the most related to our work. A first idea is to better utilize observed samples, e.g., memory replay [Lin, 1992] or hindsight experience replay (HER) [Andrychowicz *et al.*, 2017]. Although better observation reuse does not reduce sample requirements in a DRL algorithm, it decreases the number of actual interactions with the environment, which is the most important factor in robot learning. Another idea is to exploit any a priori domain knowledge one may have (e.g., symmetry [Kidziński *et al.*, 2018a]) to support learning. Besides, a robot is generally expected to solve not only one fixed task, but multiple related ones. Multi-task reinforcement learning [Plappert *et al.*, 2018] is considered beneficial as it would not be feasible to repeatedly solve each encountered task tabula rasa. Finally, in order to avoid or reduce learning on an actual robot, recent works have investigated how policies learned in a simulator can be transferred to a real robot [Tobin *et al.*, 2017].

In this work, in order to further reduce the requirement on the number of actual samples, we propose two novel data augmentation methods that better reuse the samples observed in the true environment by exploiting the symmetries (i.e., any invariant transformation) of the problem. For simplicity, we present them in the same set-up as HER [Andrychowicz *et al.*, 2017], although the techniques can be instantiated in other settings with any DRL algorithm based on memory replay.

The first technique, Kaleidoscope Experience Replay (KER), is based on reflectional symmetry, it posits that trajectories in robotic spaces can enjoy multiple reflections and remain in the valid workspace. Namely, for a given robotic problem, there is a set of reflections that can generate from any valid observed trajectory many new artificial valid ones for training. For concreteness, in this paper we focus on reflections with respect to hyperplanes.

Accepted after double-blind peer-review at NeurIPS 2019 Workshop on "Robot Learning: Control and Interaction in the Real World", Vancouver, Canada, 2019. Correspondence to: paul.weng@sjtu.edu.cn The second technique, Goal-augmented Experience Replay (GER), can be seen as a generalization of HER: any artificial goal g generated by HER can be instead replaced by a random goal sampled in a small ball around g. This idea takes advantage of tasks where success is defined as reaching a final pose within a distance of the goal set by a threshold (such tasks are common in robotics). Here, successful trajectories are augmented using the invariant transformation that consists in changing actual goals to random close goals.

In Sec. 2, we present this work's setup. Sec. 3 introduces related work. Sec. 4 details our two data augmentation techniques. Sec. 5 presents preliminary results; and Sec. 6 highlights key lessons.

2 Background

In this work, we consider robotic tasks that are modeled as multi-goal Markov decision processes [Schaul *et al.*, 2015] with continuous state and action spaces: $\langle S, A, G, T, R, p, \gamma \rangle$ where S is a continuous state space, A is a continuous action space, G is a set of goals, T is the unknown transition function that describes the effects of actions, R(s, a, s', g) is the immediate reward when reaching state $s' \in S$ after performing action $a \in A$ in state $s \in S$ if the goal were $g \in G$. Finally, p is a joint probability distribution over initial states and initial goals, and $\gamma \in (0, 1)$ is a discount factor. In this framework, the robot learning problem corresponds to an RL problem that aims at obtaining a policy $\pi : S \times G \to A$ such that the expected discounted sum of rewards is maximized for any given goal.

When the reward function is sparse, as assumed here, this RL problem is particularly hard to solve. In particular, we consider here reward functions that are described as follows: $R(s, a, s', g) = \mathbf{1}[d(s', g) \leq \epsilon_R] - 1$ where **1** is the indicator function, d is a distance, and $\epsilon_R > 0$ is a fixed threshold.

To tackle this issue, Andrychowicz *et al.* [2017] proposed HER, which is based on the following principle: Any trajectory that failed to reach its goal still carries useful information; it has at least reached the states of its trajectory path. Using this natural and powerful idea, memory replay can be augmented with the failed trajectories by changing their goals in *hindsight*.

3 Related Work

HER [Andrychowicz *et al.*, 2017; Plappert *et al.*, 2018] has been extended in various ways. Prioritized replay was incorporated in HER to learn from more valuable episodes with higher priority [Zhao and Tresp, 2018]. In [Fang *et al.*, 2019], HER was generalized to deal with dynamic goals. In [Gerken and Spranger, 2019], a variant of HER was also investigated where completely random goals replace achieved goals and in [Rauber *et al.*, 2017], it was adapted to work with on-policy RL algorithms.

Symmetry has been considered in MDPs [Zinkevich and Balch, 2001] and RL [Kamal and Murata, 2008; Agostini and Celaya, 2009; Mahajan and Tulabandhula, 2017; Kidziński *et al.*, 2018b; Amadio *et al.*, 2019]. It can be known a priori or learned [Mahajan and Tulabandhula, 2017]. In this work, we assume the former, which is reasonable in many robotics tasks. A natural approach to exploit symmetry in sequential decision-making is by aggregating states that satisfy an equivalence relation induced by a symmetry [Zinkevich and Balch, 2001; Kamal and Murata, 2008]. A dihedral group with finite invariant elements has been leveraged to implement symmetry on the state representation of board position in Go [Silver *et al.*, 2016]. Another related approach takes into account symmetry in the policy representation [Amadio *et al.*, 2019]. Doing so reduces representation may be difficult to recover, especially if many symmetries are considered simultaneously. Still another approach is to use symmetry during training instead. One simple idea is to learn the Q-function by performing an additional symmetrical update [Agostini and Celaya, 2009]. Another method is to augment the training points with their reflections [Kidziński *et al.*, 2018b].

In this paper, we generalize further this idea as a data augmentation technique where many symmetries can be considered and pairs of symmetrical updates do not need to be simultaneously applied.

While, to the best of our knowledge, data augmentation has not been considered much to accelerate learning in RL, it has been used extensively and with great success in machine learning [Baird,



Figure 1: Framework overview: real and symmetrically transformed transitions are stored in the replay buffer. Sampled minibatches are then augmented with GER before updating the policy.

1992] and more so in deep learning [Krizhevsky *et al.*, 2012]. Interestingly, symmetries can also be exploited in neural network architecture design [Gens and Domingos, 2014]. However, in our case, the integration of symmetry in deep networks will be left as future work.

4 Data Augmentation for RL

To reduce the number of interactions with the real environment our goal is to generate artificial training data based on actual trajectories collected during the robot's learning.

Our architecture leverages our two proposed techniques, Kaleidoscope experience replay (KER) and Goal-Augmented Experience Replay (GER). While the two methods are combined here, only one of them could be used instead. An overview of our architecture is illustrated in Fig. 1.

4.1 Kaleidoscope Experience Replay (KER)

KER uses reflectional symmetry¹. Consider a 3D workspace with a bisecting plane xoz as shown in Fig. 2. If a valid trajectory is generated in the workspace (blue in Fig. 2), natural symmetry would then yield a new valid trajectory reflected on this plane. More generally, the xoz plane may be rotated by some angle θ_z along axis \vec{z} and still define an invariant symmetry for the robotic task. We



Figure 2: Kaleidoscope Experience Replay leverages natural symmetry. Valid trajectories are reflected with respect to plane xoz, where the latter can itself be rotated by some θ_z along axis \vec{z} .

can now precisely define KER, which amounts to augmenting any original trajectory with a certain number of random symmetries. A random symmetry is a reflectional symmetry with respect to the xoz plane after it has been rotated by a random angle about the \vec{z} -axis.

Note that instead of storing the reflected trajectories in the replay buffer, the random symmetries could be instead applied to sampled minibatches. This approach was tried previously for single-symmetry scenarios [Kidziński *et al.*, 2018a]. Doing so, however, is more computationally taxing as transitions are reflected every time they are sampled and more significantly leads to lower performance². Our conjecture is that such approach leads to a lower diversity in the minibatches.

¹Though more general invariant transformations could also be used in place of reflectional symmetry.

²Please visit our project page for this and more supplementary information (not referenced until after review).

4.2 Goal-Augmented Experience Replay (GER)

GER exploits the formulation of any reward function that defines a successful trajectory as one whose end position is within a small radial threshold (a ball) centered around the goal. When the robot obtains a successful trajectory, we therefore know that it can in fact be considered successful for any goal within a ball centered around its end position. Based on this observation, GER augments successful trajectories by replacing the original goal with a random goal sampled within that ball. This ball can be formally described as $\{g \in \mathcal{G} | d(s_f, g) \leq \epsilon\}$ where s_f is the final state reached in the original trajectory and $\epsilon < \epsilon_R$ is a threshold, which does not have to be constant for each application of GER. Therefore, GER can be seen as a generalization of HER and can be implemented in the same fashion. This is why in our architecture, GER is applied on minibatches, like HER.

5 Preliminary Experimental Results

Our experimental evaluation is performed according to the HER formulation [Andrychowicz *et al.*, 2017]. Namely, we use a simulated 7-DOF Fetch Robotics arm trained with DDPG on the pushing, sliding, and pick-and-place tasks.

We design our experiments to demonstrate the effectiveness of our propositions and final combination, which uses n = 8 random symmetries for KER and 4 applications of GER (where one of them uses a threshold equal to zero in order to also take full advantage of realized goals). We now present some initial experimental results. As shown in Fig. 3, our method vastly improves the learning speed compared to vanilla HER.

In our experiments, we have observed that performance is monotonic with respect to the number n of random symmetries, although the gain diminishes for larger n. Similar observations can be made for the number of applications of GER.



Figure 3: Comparison of vanilla HER with the combination of 8 KER symmetries and 4 GERs.

6 Conclusion

We proposed two novel data augmentation techniques KER and GER to amplify the efficiency of observed samples in a memory replay mechanism. KER exploited reflectional symmetry in the valid workspace (though in general it could be employed with other types of symmetries). GER, as an extension of HER, is specific to goal-oriented tasks where success is defined in terms of a thresholded distance. The combination of these techniques greatly accelerated learning as demonstrated in our experiments.

Our next step is to use our method to solve the same tasks with a simulated Baxter robot, and then transfer to a real Baxter using a sim2real methodology. Furthermore, we aim at extending our proposition to other types of symmetries and to other robotic tasks as well.

References

- Alejandro Agostini and Enric Celaya. Exploiting Domain Symmetries in Reinforcement Learning with Continuous State and Action Spaces. In *ICMLA*, 2009.
- Fabio Amadio, Adria Colome, and Carme Torras. Exploiting Symmetries in Reinforcement Learning of Bimanual Robotic Tasks. *IEEE Robotics and Automation Letters*, 4(2):1838–1845, 4 2019.
- Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. In Advances in Neural Information Processing Systems, pages 5049–5059, 2017.
- Henry S. Baird. Document Image Defect Models. In *Structured Document Image Analysis*, pages 546–556. Springer, 1992.
- Meng Fang, Cheng Zhou, Bei Shi, Boqing Gong, Jia Xu, and Tong Zhang. DHER: Hindsight Experience Replay. In *ICLR*, 2019.
- Robert Gens and Pedro Domingos. Deep symmetry networks. In Advances in Neural Information Processing Systems, pages 2537–2545, 2014.
- Andreas Gerken and Michael Spranger. Continuous Value Iteration (CVI) Reinforcement Learning and Imaginary Experience Replay (IER) For Learning Multi-Goal, Continuous Action and State Space Controllers. In 2019 International Conference on Robotics and Automation (ICRA), pages 7173–7179. IEEE, 5 2019.
- Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, and Sergey Levine. Scalable deep reinforcement learning for vision-based robotic manipulation. In Aude Billard, Anca Dragan, Jan Peters, and Jun Morimoto, editors, *Proceedings of The 2nd Conference on Robot Learning*, volume 87 of *Proceedings of Machine Learning Research*, pages 651–673. PMLR, 29–31 Oct 2018.
- M.A.S. Kamal and Junichi Murata. Reinforcement learning for problems with symmetrical restricted states. *Robotics and Autonomous Systems*, 56(9):717–727, 9 2008.
- Łukasz Kidziński, Sharada Prasanna Mohanty, Carmichael F Ong, Zhewei Huang, Shuchang Zhou, Anton Pechenko, Adam Stelmaszczyk, Piotr Jarosik, Mikhail Pavlov, Sergey Kolesnikov, et al. Learning to run challenge solutions: Adapting reinforcement learning methods for neuromusculoskeletal environments. In *The NIPS'17 Competition: Building Intelligent Systems*, pages 121–153. Springer, 2018.
- Łukasz Kidziński, Sharada Prasanna Mohanty, Carmichael F. Ong, Zhewei Huang, Shuchang Zhou, Anton Pechenko, Adam Stelmaszczyk, Piotr Jarosik, Mikhail Pavlov, Sergey Kolesnikov, Sergey Plis, Zhibo Chen, Zhizheng Zhang, Jiale Chen, Jun Shi, Zhuobin Zheng, Chun Yuan, Zhihui Lin, Henryk Michalewski, Piotr Milos, Blazej Osinski, Andrew Melnik, Malte Schilling, Helge Ritter, Sean F. Carroll, Jennifer Hicks, Sergey Levine, Marcel Salathé, and Scott Delp. Learning to Run Challenge Solutions: Adapting Reinforcement Learning Methods for Neuromusculoskeletal Environments. In *The NIPS '17 Competition: Building Intelligent Systems*, pages 121–153. Springer, 2018.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems, pages 1097–1105, 2012.
- Long-Ji Lin. Self-improving reactive agents based on reinforcement learning, planning and teaching. Machine learning, 8(3-4):293–321, 1992.
- Anuj Mahajan and Theja Tulabandhula. Symmetry learning for function approximation in reinforcement learning. arXiv preprint arXiv:1706.02999, 2017.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.

- OpenAI, Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafał Józefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, Jonas Schneider, Szymon Sidor, Josh Tobin, Peter Welinder, Lilian Weng, and Wojciech Zaremba. Learning dexterous in-hand manipulation. *CoRR*, 2018.
- Matthias Plappert, Marcin Andrychowicz, Alex Ray, Bob McGrew, Bowen Baker, Glenn Powell, Jonas Schneider, Josh Tobin, Maciek Chociej, Peter Welinder, et al. Multi-goal reinforcement learning: Challenging robotics environments and request for research. *arXiv preprint arXiv:1802.09464*, 2018.
- Paulo Rauber, Avinash Ummadisingu, Filipe Mutz, and Juergen Schmidhuber. Hindsight policy gradients. arXiv preprint arXiv:1711.06006, 2017.
- Tom Schaul, Daniel Horgan, Karol Gregor, and David Silver. Universal value function approximators. In *International Conference on Machine Learning*, pages 1312–1320, 2015.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484, 2016.
- Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 23–30. IEEE, 2017.

Rui Zhao and Volker Tresp. Energy-Based Hindsight Experience Prioritization. In coRL, 2018.

Martin Zinkevich and Tucker Balch. Symmetry in markov decision processes and its implications for single agent and multi agent learning. In *In Proceedings of the 18th International Conference on Machine Learning*. Citeseer, 2001.