

Hierarchical Deep Q-Network from imperfect demonstrations in Minecraft

Alexey Skrynnik^{a,*}, Aleksey Staroverov^b, Ermek Aitygulov^b,
Kirill Aksenov^c, Vasilii Davydov^d, Aleksandr I. Panov^{a,b}

^a Artificial Intelligence Research Institute FRC CSC RAS, Russia

^b Moscow Institute of Physics and Technology, Russia

^c Higher School of Economics, Russia

^d Moscow Aviation Institute, Russia

Received 29 July 2020; accepted 26 August 2020

Available online 22 October 2020

Abstract

We present Hierarchical Deep Q-Network (HDQfD) that won first place in the MineRL competition. The HDQfD works on imperfect demonstrations and utilizes the hierarchical structure of expert trajectories. We introduce the procedure of extracting an effective sequence of meta-actions and subgoals from the demonstration data. We present a structured task-dependent replay buffer and an adaptive prioritizing technique that allow the HDQfD agent to gradually erase poor-quality expert data from the buffer. In this paper, we present the details of the HDQfD algorithm and give the experimental results in the Minecraft domain.

© 2020 Elsevier B.V. All rights reserved.

Keywords: Reinforcement learning; Minecraft; Demonstrations

1. Introduction

Deep reinforcement learning (RL) has achieved compelling success on many complex sequential decision-making problems, especially in simple domains. In such examples as AlphaStar (Vinyals et al., 2019), AlphaZero (Silver et al., 2017), OpenAI Five human or superhuman performance was attained. However, RL algorithms usually require a huge amount of environment-samples required for training to reach good performance (Kakade et al., 2003). Learning from demonstration is a well-known alternative, but until now, this approach has not achieved any considerable success in complex non-single-task environments. This was largely due to the fact that

obtaining high-quality expert demonstrations in a sufficient quantity in sample-limited, real-world domains is a separate non-trivial problem.

Minecraft as a compelling domain for the development of reinforcement and imitation learning based methods was recently introduced by Guss, Houghton, et al. (2019). Minecraft presents unique challenges because it is a 3D, first-person, open-world game where the agent should gather resources and create structures and items to achieve a goal. Due to its popularity as a video game it turned out to be possible to collect a large number of expert trajectories in which individual subtasks are solved. This allowed the appealing MineRL competition to run. The organizers have released the largest-ever dataset of human demonstrations on a Minecraft domain. The primary goal of the competition is to foster the development of algorithms that can efficiently leverage human priors to drastically reduce the

* Corresponding author.

E-mail address: skrynnik@isa.ru (A. Skrynnik).

number of samples needed to solve complex, hierarchical, and sparse environments.

The main difficulty in solving the MineRL problem was the imperfection of demonstrations and the presence of hierarchical relationships of subtasks. In this paper we present hierarchical Deep Q-Network from Demonstrations (HDQfD) that allowed us to win first place in the MineRL competition proposed by Guss, Codel, et al. (2019). The HDQfD works on imperfect demonstrations and utilize a hierarchical structure of expert trajectories extracting effective sequence of meta-actions and subgoals. Each subtask is solved by its own simple strategy, which extends the DQfD approach proposed by Gao, Lin, Yu, Levine, and Darrell (2018) and relies on a structured buffer and gradually decrease the ratio of poor-quality expert data. In this paper, we present the details of our algorithm and provide the results that allow the HDQfD agent to play Minecraft at the human level.

2. Background

One way to explore the domain with the use of expert data is to do behavioral cloning (BC). Pure supervised learning BC methods suffer from a distribution shift, because the agent greedily imitates the demonstrated actions, it can drift away from the demonstrated states due to error accumulation. Another way to use expert data in search for the exploration policy is to use conventional RL methods like PPO, DDDQN, etc. and guide exploration through enforcing occupancy measure matching between the learned policy and current demonstrations.

The main approach is to use demonstration trajectories sampled from an expert policy to guide the learning procedure by either putting the demonstrations into a replay buffer or using them to pre-train the policy in a supervised manner.

The organizers of the MineRL competition provided a few baselines. Standard DQfD presented by Hester et al. (2018) get the max score of 64 after 1000 episodes, the PPO gets the max of 55 after 800 episodes, the rainbow also gets the max of 55 after 800 episodes of training.

Our best solution exploits the method of injecting expert data into the agent replay buffer. The DQfD, which our method is based on, is an advanced approach to reinforcement learning with additional expert demonstrations. The main idea of the DQfD is to use an algorithm called Deep Q-Network (DQN) and combine loss function $J(Q)$, with the main component $J_E(Q)$:

$$J(Q) = J_{DQ}(Q) + \lambda_1 J_n(Q) + \lambda_2 J_E(Q) + \lambda_3 J_{L2}(Q). \quad (1)$$

The loss function $J_{DQ}(Q)$ is a standard TD-error:

$$J_{DQ}(Q) = (R(s, a) + \gamma Q(s_{t+1}, a_{t+1}^{\max}; \theta') - Q(s, a; \theta))^2. \quad (2)$$

The loss function $J_n(Q)$ is the so-called N-step return, which allows the agent to extend the utility of trajectories to several steps, which lead to a better strategy:

$$J_n(Q) = r_t + \gamma r_{t+1} + \dots + \gamma^{n-1} r_{t+n-1} + \max_a \gamma^n Q(s_{t+n}, a). \quad (3)$$

The main part $J_E(Q)$ is a margin loss function. It is responsible for copying expert behavior and gives a penalty to the agent for performing actions other than those of the experts:

$$J_E(Q) = \max_{a \in \mathcal{A}} [Q(s, a) + l(a_E, a)] - Q(s, a_E). \quad (4)$$

Finally, $J_{L2}(Q)$ is L2 regularization added to prevent overfitting.

The problem of suboptimal demonstrations is considered by Nair, McGrew, Andrychowicz, Zaremba, and Abbeel (2018), who attempted to develop an approach closely related to ours. The authors use the BC loss as a margin loss only for expert data. The key idea of the suggested approach is Q-Filter, which applies the BC loss only to the states where the critic determines that the action of the demonstrator is better than that of the actor. The use of expert demonstrations depends on the quality of the critic, which can lead to difficulties in the case of sparse rewards. For such tasks, the authors of Nair et al. (2018) propose resets to the demonstration states technique for Hindsight Experience Replay (Andrychowicz et al., 2017). In contrast to the Q-filter, the HDQfD is suitable for solving a wide range of problems related to the quality of demonstrations. In this paper, our approach shows an ability to learn from expert data to which a change in the action space is applied.

3. Hierarchical Deep Q-Network from demonstrations

3.1. Action and state space

To make the demonstration data convenient for the RL agent we used action discretization and some techniques for state space preparation, in particular a frameskip and a framestack. The frameskip technique repeats the action selected by the agent over several steps. The framestack technique replaces the current state with the concatenation of the current and several previous states of the environment. In the MineRL simulator, the agent could choose between 10 actions (see Table 1). The expert action is mapped to the agent's action in the order shown in Table 1. For example, “turn the camera right 10 degrees, turn the camera up 5 degrees, run forward” will be mapped with the first action, i.e., turn the camera right 5 degrees and attack. All “move”-actions (back, forward, left, right) were allowed because the experts used mostly them to direct the camera to the tree block.

3.2. Adaptive ration of expert data

Despite this action, the space discretization allowed making good behavior cloning as there was some noise in the demonstrations due to which the agent could not

Table 1

Discretization of actions used for all subtasks with frameskip 4. The expert action is mapped to the agent's action in the order shown in this table. The rotation angle is determined using the sum of 4 frames. For other actions, the most frequent was selected.

actions	a_0	a_1	a_2	a_3	a_4	a_5	a_6	a_7	a_8	a_9
pitch + 5	+									
pitch - 5		+								
yaw + 5			+							
yaw - 5				+						
Forward					+	+				
Left							+			
Right								+		
Back									+	
Jump						+				+
Attack	+	+	+	+	+	+	+	+	+	+

improve its strategy above a certain threshold. We solved this problem by adding an ability to change the ratio of the expert data. The demonstrations and the agents' trajectories were stored separately in Aggregating Buffer, which controls the proportion of demonstrations in mini-batches. The proportion decreases linearly depending on the number of episodes (see Fig. 1).

3.3. Extracting hierarchical subtask structure

We examined each expert's trajectory separately and considered the time of appearance of the items in the inventory in a chronological order. An example of a possible order of obtaining items is shown in Fig. 2. In addition, this sequence can be considered as a semantic network with two

types of nodes: certain agent's actions and subtasks defined for the agent's inventory. We consider each subtask node in this network as a mandatory sub-goal that the agent must complete to move on. We train a separate strategy for the agent to achieve each sub-goal, and it can be considered as a set of individual agents. The task of such agents is to obtain the necessary number of the items in the inventory.

The agent that solves the subtask is divided into two agents that take actions at the same time (see Fig. 3): the agent performing basic actions in the environment (POV or *item* agent) and the agent interacting with semantic actions – sequentially perform the action denoted in the corresponding node of the semantic network. The training scheme for the *item* agents is presented in Fig. 4. During the training process, all the expert data from the ObtainIronPickaxe environment of the MineRL simulator is used.

The frames of a mini-batch that correspond to the currently trained *item* agent are considered as expert data. All the other frames are considered as additional data and their rewards are nullified. This approach allows both training the agent to move from solving one subtask to another in addition to the effective use of available data.

4. Experiments

Here we report our results from Round 2 of the MineRL Competition. All the agents except for the *log* agent were trained on the expert data gathered from the ObtainIronPickaxeDense dataset. A summary of all the submissions is presented in Table 2.

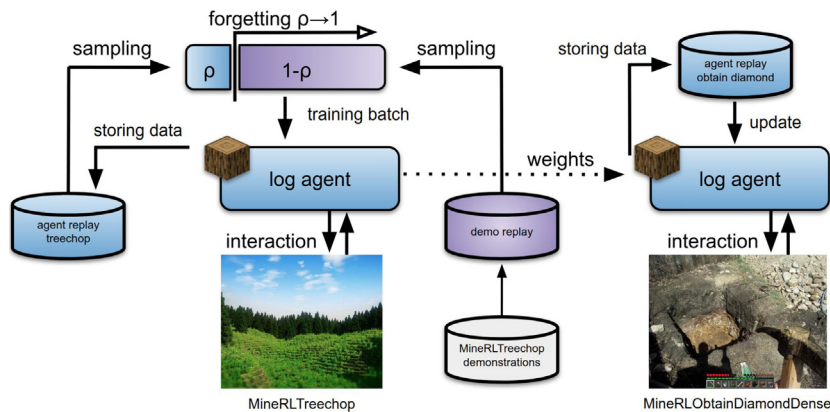


Fig. 1. Training *log* agent. The aggregated buffer is used to store expert and agent trajectories. The amount of data in the mini-batch sampled from the demo replay buffer is gradually decreasing.




Fig. 2. Example of expert's trajectory from demonstrations. Action nodes highlighted in red. Subtask nodes highlighted in blue. Additionally, the number of items for each subtask indicated. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

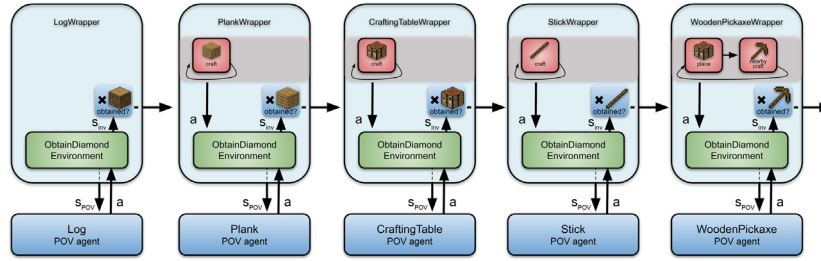


Fig. 3. Separation of the agent types (policies) for the subtasks and specific actions to craft and obtain the items.

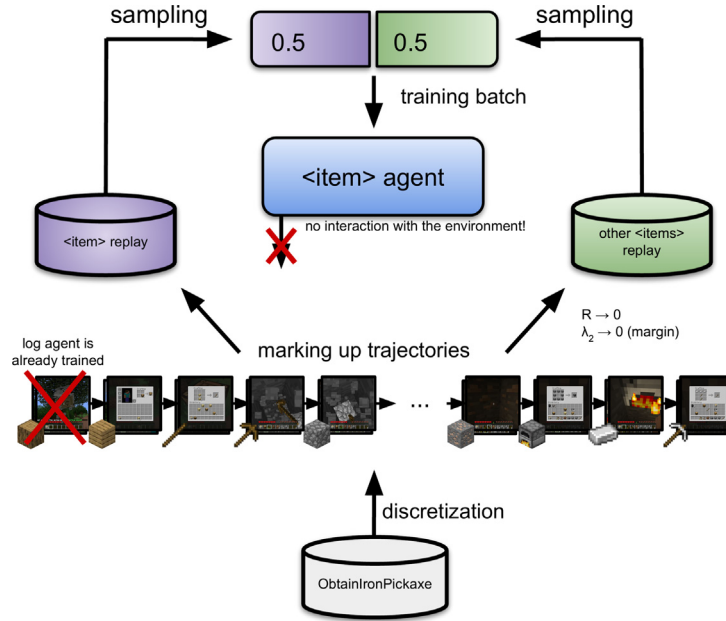


Fig. 4. Training item agents. For the item agents each trajectory is divided into expert and non-expert segments. The item agent learns to solve one subtask using the data from the other subtasks that it considers as non-expert.

Table 2
Round 2 submissions.

Log agent: Treechop \Rightarrow ObtainDiamondDense						
	Treechop episodes	Reward	Episodes	Reward	Pre-training	Evaluation
Submit 1	–	–	–	–	10^4 steps	20.72
Submit 2	200	53.22	300	16.31	10^4 steps	55.08
Submit 3	200	53.83	300	19.19	5×10^4 steps	61.61




Fig. 5. Log agent results for the Treechop (left) and ObtainDiamondDense environments (right).

Table 3
Comparison of the algorithms implemented for the *Treechop* environment.

	Demonstrations	Discretization	Embeddings	Episodes	Reward
SAC				300	5
GAIL	+	+	+	150	30
RnD		+	+	1000	35
PPO		+	+	1000	35
Pretrained PPO	+	+	+	150	50
HDQfD	+	+		200	60

In the first submit, the HDQfD agent was trained using only expert data. Each of the *item* agents was pre-trained using 10^4 steps. The *Log* agent learned on the *Treechop* environment data. The final result was **20.72**.

In the 2nd and the 3rd submissions we used interaction with the environment to train the *log* agent. The *log* agent was trained 200 episodes on the *Treechop* environment data, and then on 300 episodes of the *ObtainDiamondDense* environment data (see dynamics in the Fig. 5). The difference was in the number of pre-training steps. The final results were **55.08** and **61.61**, respectively. In addition, in Table 3 we provide a comparison of the algorithms that we tested for the *Treechop* environment.

5. Conclusion

In this paper we introduce a novel approach to learn from imperfect demonstrations. This hierarchical Deep Q-Network from Demonstrations won first place in the MineRL competition and received **61.61** score. In our further work, we plan to train all *item* agents for full hierarchical end-to-end architecture. Besides, for these agents we plan to ensure access to all demonstrations from all subtasks with respect to the agent’s inventory for additional performance.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by the Russian Science Foundation, project No. 18-71-00143. We thank the AIM Tech Company for its organizational and computing support.

References

- Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., et al. (2017). Hindsight experience replay. *Advances in Neural Information Processing Systems*, 5048–5058.
- Gao, Y., Lin, J., Yu, F., Levine, S., Darrell, T., et al. (2018). *Reinforcement learning from imperfect demonstrations*. arXiv preprint arXiv:1802.05313.
- Guss, W. H., Codel, C., Hofmann, K., Houghton, B., Kuno, N., Milani, S., et al. (2019). *The minerl competition on sample efficient reinforcement learning using human priors*. arXiv preprint arXiv:1904.10079.
- Guss, W. H., Houghton, B., Topin, N., Wang, P., Codel, C., Veloso, M., et al. (2019). *Minerl: A large-scale dataset of minecraft demonstrations*. arXiv preprint arXiv:1907.13440.
- Hester, T., Vecerik, M., Pietquin, O., Lanctot, M., Schaul, T., Piot, B., et al. (2018). Deep q-learning from demonstrations. In *Thirty-second AAAI conference on artificial intelligence*.
- Kakade, S. M. et al. (2003). *On the sample complexity of reinforcement learning* Ph.D. thesis. England: University of London London.
- Nair, A., McGrew, B., Andrychowicz, M., Zaremba, W., & Abbeel, P. (2018). Overcoming exploration in reinforcement learning with demonstrations. In *2018 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 6292–6299). IEEE.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., et al. (2017). Mastering the game of go without human knowledge. *Nature*, 550(7676), 354.
- Vinyals, O., Babuschkin, I., Chung, J., Mathieu, M., Jaderberg, M., Czarnecki, W. M., et al. (2019). Alphastar: Mastering the real-time strategy game starcraft ii. *DeepMind Blog*.