

# EasyRL: A Simple and Extensible Reinforcement Learning Framework

Neil Hulbert, Sam Spillers, Brandon Francis, James Haines-Temons, Ken Gil Romero, Benjamin De Jager, Sam Wong, Kevin Flora, Bowei Huang, Athirai A. Irissappane

University of Washington, Tacoma

{nhulbert, samds4, brandf3, jamesht, kgmr, benjad2, sam543, flora01, boweih, athirai}@uw.edu

## Abstract

In recent years, Reinforcement Learning (RL), has become a popular field of study as well as a tool for enterprises working on cutting-edge artificial intelligence research. To this end, many researchers have built RL frameworks such as openAI Gym, and KerasRL for ease of use. While these works have made great strides towards bringing down the barrier of entry for those new to RL, we propose a much simpler framework called EasyRL, by providing an interactive graphical user interface for users to train and evaluate RL agents. As it is entirely graphical, EasyRL does not require programming knowledge for training and testing simple built-in RL agents. EasyRL also supports custom RL agents and environments, which can be highly beneficial for RL researchers in evaluating and comparing their RL models.

## Introduction

Reinforcement Learning (RL) is a popular formalism for automated decision-making. It is a growing field with impressive advances. However, existing RL techniques have been mostly applied and evaluated on games (Mnih et al. 2015; Silver et al. 2016). Apart from a small number of real-world applications that use RL solutions (Zeng et al. 2016; Moriyama et al. 2018), RL is not a widely popular decision-making model across all domains. RL training, in general, is a cumbersome process. Several issues including (but not limited to) hyper-parameter tuning, sample efficiency, and training stability (Schulman et al. 2017) need to be carefully addressed during the training process. It has therefore become indispensable to have a solid background knowledge in RL and sufficient software development skills to successfully develop and train RL agents. The above requirements restrict the RL audience to only a handful of researchers who are experts in RL. Although RL can solve decision-making problems in other domains such as healthcare, transportation, networking, etc., people in these fields may find it cumbersome to train RL agents without sufficient expertise.

Our main goal is to build a RL framework that can be used by diverse audiences from different domains. To this end, we propose the EasyRL framework for both native as well as non-native RL users to easily develop, train, and evaluate RL agents. The existing RL frameworks Keras-RL, Ten-

sorforce, Horizon, HuskaRL, SimpleRL, AI-ToolBox, and Coach provide a range of built-in deepRL (deep reinforcement learning) techniques and some of them, such as Coach, create visualization of the training process for debugging purposes (Winder 2019). However, they do not support a user-friendly Graphical User Interface (GUI). Even the popular OpenAI gym only supports a range of environments to evaluate RL agents. Further, all existing frameworks require some amount of RL and programming knowledge.

Our EasyRL framework offers an interactive GUI to build, train, and evaluate RL agents. It hosts a number of built-in RL agents (algorithms) and environments. Additionally, a user may develop his own custom RL agent or environment and add it to the framework. The EasyRL framework requires no or minimal programming skills for training in-built RL agents. This differs from many previously existing RL frameworks by greatly simplifying access and reducing the technical barrier of entry to training RL agents. As the amount of personal computing resources has tremendously increased, the applicability of RL techniques to well-defined environments can be better leveraged to allow non-native RL users to themselves train RL agents. By introducing the user to RL and giving them the tools to create agents and environments themselves, our framework will improve the visibility as well as applicability of RL across different domains.




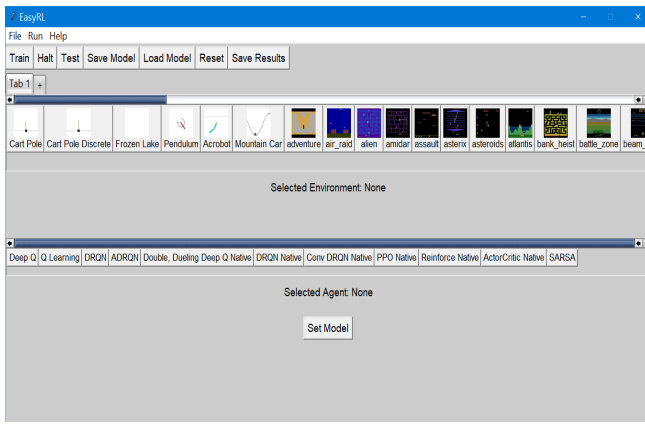
Figure 1: Structure of the EasyRL Framework

## The EasyRL Framework

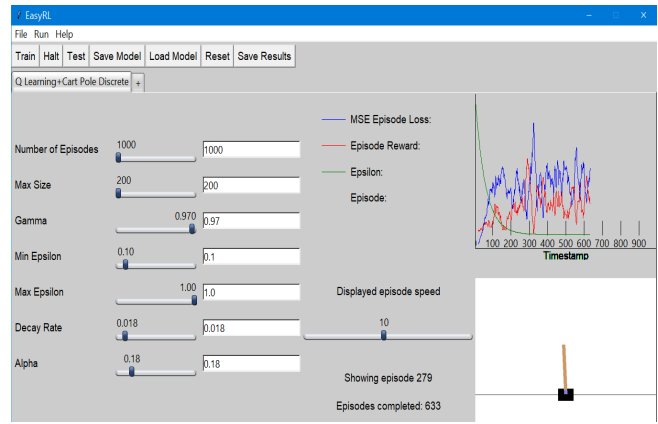
The proposed EasyRL framework<sup>1,2</sup> allows the training and evaluation of RL agents on a variety of openAI gym as

<sup>1</sup>Source Code: <https://github.com/easyRL/easyRL-v0>

<sup>2</sup>Demo Video: [https://youtu.be/Sc5phA.TR\\_o](https://youtu.be/Sc5phA.TR_o)



(a) Choosing Environment & RL Agent



(b) Setting Hyper-Parameters & Visualizing Results

Figure 2: EasyRL Framework GUI

well as custom real-world environments. EasyRL follows a highly modularized implementation with abstractions such as *Agent* and *Environment*. The sequence diagram for navigating through the framework is shown in Fig. 1. The GUI for the framework is shown in Fig. 2. The user can select from a variety of RL agents and environments (see Fig. 2a). The user can then set the hyper-parameters for training (see Fig. 2b). The training results are plotted using metrics such as mean rewards and training loss. The graphs also show the epsilon annealing process. The training environment is dynamically rendered on the screen. The rendering speed can also be changed (see *Display Episode Speed* in Fig. 2b).

The trained RL model as well as results can be saved for future use. The user can load a previously trained RL agent using *Load Model* and run test cases as well as visualize the results. The framework provides options to create custom RL agents and environments using *Load Agent* and *Load Environment* options in the File menu. The menu also has a detailed help guide to assist the user with these commands.

```
class CustomEnv(environment.Environment):
    displayName = 'CustomEnv'

    def __init__(self):
    def step(self, action):
    def reset(self):
    def sample_action(self):
    def render(self):
```

Figure 3: API for Custom Environment

## RL Algorithms & Environments

EasyRL currently hosts a list of model-free algorithms that can handle both fully-observable environments such as Q-learning (Watkins and Dayan 1992), SARSA (Rummery and Niranjan 1994), DQN (Mnih et al. 2015), DDQN (Van Hasselt, Guez, and Silver 2016), PPO (Schulman et al. 2017), REINFORCE (Williams 1992) and partially-observable environments such as DRQN (Hausknecht and Stone 2015)

and ADRQN (Zhu et al. 2017). The off-policy deepRL techniques mentioned above are implemented using standard experience replay for sampling experiences. It should be noted that our framework also supports model-based RL agents. The user is also allowed to create custom RL agents and import them to the EasyRL framework (as a Python file).

The framework hosts a variety of OpenAI Gym environments (classic control and atari). The user can also create a custom environment by following the API shown in Fig. 3. We have implemented some custom environments, e.g., seller selection in e-markets (Irissappane et al. 2014).

The EasyRL framework is highly modularized and extensible (MVC design pattern). It is predominately written in python and supports both Tensorflow as well as Pytorch deep learning libraries. EasyRL also supports C++ native implementations (see DRQNNative, DDQNNative) via CFFI, which speeds up the training atleast by 5 times. The framework, by default, uses local CPU/GPU during training, however, it can be easily configured to use resources remotely. Further, EasyRL supports training multiple RL agents in parallel via the Python Threading library. The EasyRL framework is easy to install and is supported by Windows, Linux as well as iOS. We also provide a command-line interface, offering the same functionality as the GUI.

## Demo

Our demonstration<sup>2</sup> will show how a GUI can greatly simplify the process of developing, training, and testing a RL agent. We will demonstrate our simple installation procedure and show how a user with minimal knowledge of RL and programming can successfully train a RL agent. In addition to training and testing different combinations of agents and environments, we will show how to save and load pre-trained RL agents along with the results from a training or test run. We will demonstrate how to create custom environments and RL agents and show the training results for one such custom environment including the visualization graphs. Furthermore, we will show how multiple agents can be trained simultaneously and the improvement in training speed when native C++ implementation is used.

## References

- Hausknecht, M.; and Stone, P. 2015. Deep Recurrent Q-learning for Partially Observable MDPs. In *2015 AAAI Fall Symposium Series*.
- Irissappane, A. A.; Oliehoek, F. A.; Zhang, J.; et al. 2014. A POMDP based Approach to Optimally Select Sellers in Electronic Marketplaces. In *International Conference on Autonomous Agents and Multi-Agent Systems*, 1329–1336.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level Control through Deep Reinforcement Learning. *Nature* 518(7540): 529–533.
- Moriyama, T.; De Magistris, G.; Tatsubori, M.; Pham, T.-H.; Munawar, A.; and Tachibana, R. 2018. Reinforcement learning testbed for power-consumption optimization. In *Asian Simulation Conference*, 45–59. Springer.
- Rummery, G. A.; and Niranjana, M. 1994. *On-line Q-learning using Connectionist Systems*, volume 37. University of Cambridge, Department of Engineering Cambridge, UK.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. *arXiv preprint arXiv:1707.06347*.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature* 529(7587): 484–489.
- Van Hasselt, H.; Guez, A.; and Silver, D. 2016. Deep Reinforcement Learning with Double Q-learning. In *AAAI Conference on Artificial Intelligence*.
- Watkins, C. J.; and Dayan, P. 1992. Q-learning. *Machine Learning* 8(3-4): 279–292.
- Williams, R. J. 1992. Simple Statistical Gradient-following Algorithms for Connectionist Reinforcement Learning. *Machine Learning* 8(3-4): 229–256.
- Winder. 2019. Comparison of Reinforcement Learning Frameworks: Dopamine, RLlib, KerasRL, Coach, Tensorforce and More. <https://winderresearch.com/a-comparison-of-reinforcement-learning-frameworks-dopamine-rl-lib-keras-rl-coach-trfl-tensorforce-coach-and-more/>. Accessed: 2020-11-18.
- Zeng, C.; Wang, Q.; Mokhtari, S.; and Li, T. 2016. Online context-aware recommendation with time varying multi-armed bandit. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2025–2034.
- Zhu, P.; Li, X.; Poupart, P.; and Miao, G. 2017. On Improving Deep Reinforcement Learning for POMDPs. *arXiv preprint arXiv:1704.07978*.