RLzoo

Release 0.0.1

May 02, 2020

User Guide

RLzoo is a collection of the most practical reinforcement learning algorithms, frameworks and applications. It is implemented with Tensorflow 2.0 and API of neural network layers in TensorLayer 2, to provide a hands-on fast-developing approach for reinforcement learning practices and benchmarks. It supports basic toy-tests like OpenAI Gym and DeepMind Control Suite with very simple configurations. Moreover, RLzoo supports robot learning benchmark environment RLBench based on Vrep/Pyrep simulator. Other large-scale distributed training framework for more realistic scenarios with Unity 3D, Mujoco, Bullet Physics, etc, will be supported in the future. A Springer textbook is also provided, you can get the free PDF if your institute has Springer license.

CHAPTER 1

Installation

Direct installation:

```
pip install rlzoo
```

Install from the source code on github:

```
git clone https://github.com/tensorlayer/RLzoo.git
cd RLzoo
pip install .
```

CHAPTER 2

Quick Start

Open `./run_rlzoo.py`:

```
from rlzoo.common.env_wrappers import build_env
from rlzoo.common.utils import call_default_params
from rlzoo.algorithms import TD3
# choose an algorithm
AlgName = 'TD3'
# chose an environment
EnvName = 'Pendulum-v0'
# select a corresponding environment type
EnvType = 'classic_control'
# build an environment with wrappers
env = build_env(EnvName, EnvType)
# call default parameters for the algorithm and learning process
alg_params, learn_params = call_default_params(env, EnvType, AlgName)
# instantiate the algorithm
alg = eval(AlgName+'(**alg_params)')
# start the training
alg.learn(env=env, mode='train', render=False, **learn_params)
# test after training
alg.learn(env=env, mode='test', render=True, **learn_params)
```

Run the example:

```
python run_rlzoo.py
```

CHAPTER 3

DQN and Variants

3.1 Deep Q-Networks

3.2 Default Hyper-parameters

CHAPTER 4

VPG

4.1 Vanilla Policy Gradient

4.2 Default Hyper-parameters

CHAPTER 5

AC

5.1 Actor-Critic

5.2 Default Hyper-parameters

CHAPTER 6

A3C

6.1 Asychronous Advantage Actor-Critic

6.2 Default Hyper-parameters

CHAPTER 7

DDPG

7.1 Deep Deterministic Policy Gradient

7.2 Default Hyper-parameters

CHAPTER 8

TD3

8.1 Twin Delayed DDPG

8.2 Default Hyper-parameters

CHAPTER 9

SAC

9.1 Soft Actor-Critic

9.2 Default Hyper-parameters

CHAPTER 10

TRPO

10.1 Trust Region Policy Optimization

10.2 Default Hyper-parameters

CHAPTER 11

PPO

11.1 Proximal Policy Optimization (Penalty)

11.2 Proximal Policy Optimization (Clip)

11.3 Default Hyper-parameters

CHAPTER 12

DPPO

12.1 Distributed Proximal Policy Optimization (Penalty)

12.2 Distributed Proximal Policy Optimization (Clip)

12.3 Default Hyper-parameters

CHAPTER 13

Common

CHAPTER 14

DRL Book

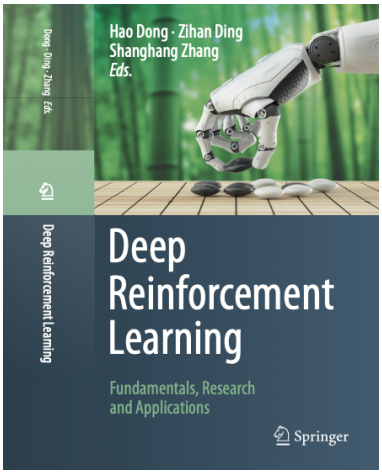

- You can get the free PDF if your institute has Springer license.

Deep reinforcement learning (DRL) relies on the intersection of reinforcement learning (RL) and deep learning (DL). It has been able to solve a wide range of complex decision-making tasks that were previously out of reach for a machine and famously contributed to the success of AlphaGo. Furthermore, it opens up numerous new applications in domains such as healthcare, robotics, smart grids, and finance.

Divided into three main parts, this book provides a comprehensive and self-contained introduction to DRL. The first part introduces the foundations of DL, RL and widely used DRL methods and discusses their implementation. The second part covers selected DRL research topics, which are useful for those wanting to specialize in DRL research. To help readers gain a deep understanding of DRL and quickly apply the techniques in practice, the third part presents mass applications, such as the intelligent transportation system and learning to run, with detailed explanations.

The book is intended for computer science students, both undergraduate and postgraduate, who would like to learn DRL from scratch, practice its implementation, and explore the research topics. This book also appeals to engineers and practitioners who do not have strong machine learning background, but want to quickly understand how DRL works and use the techniques in their applications.

14.1 Editors

- Hao Dong - Peking University
- Zihan Ding - Princeton University
- Shanghang Zhang - University of California, Berkeley

14.2 Authors

- Hao Dong - Peking University
- Zihan Ding - Princeton University
- Shanghang Zhang - University of California, Berkeley
- Hang Yuan - Oxford University
- Hongming Zhang - Peking University
- Jingqing Zhang - Imperial College London
- Yanhua Huang - Xiaohongshu Technology Co.
- Tianyang Yu - Nanchang University
- Huaqing Zhang - Google
- Ruitong Huang - Borealis AI

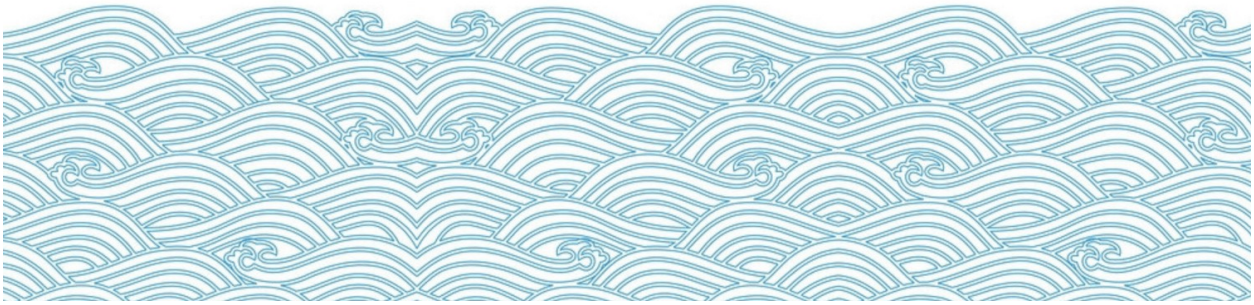

CHAPTER 15

DRL Tutorial

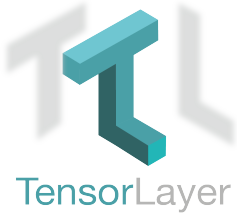

Different from RLzoo for simple usage with **high-level APIs**, the RL tutorial aims to make the reinforcement learning tutorial simple, transparent and straight-forward with **low-level APIs**, as this would not only benefits new learners of reinforcement learning, but also provide convenience for senior researchers to testify their new ideas quickly.

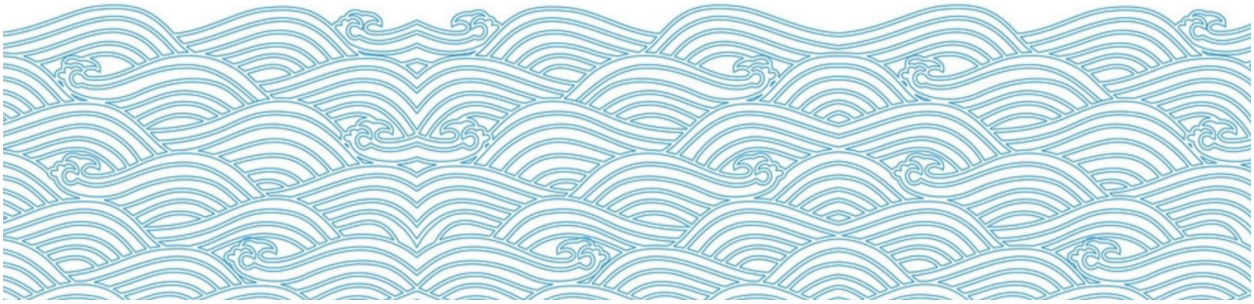

CHAPTER 16

Contributing

This project is under active development, if you want to join the core team, feel free to contact Zihan Ding at ....

CHAPTER 17

Citation

- genindex
- modindex
- search