
RLLTE: Long-Term Evolution Project of Reinforcement Learning

¹Mingqi Yuan, ²Zequn Zhang, ³Yang Xu, ⁴Shihao Luo, ¹Bo Li, ²Xin Jin, and
²Wenjun Zeng

¹The Hong Kong Polytechnic University

²Eastern Institute for Advanced Study

³Purdue University

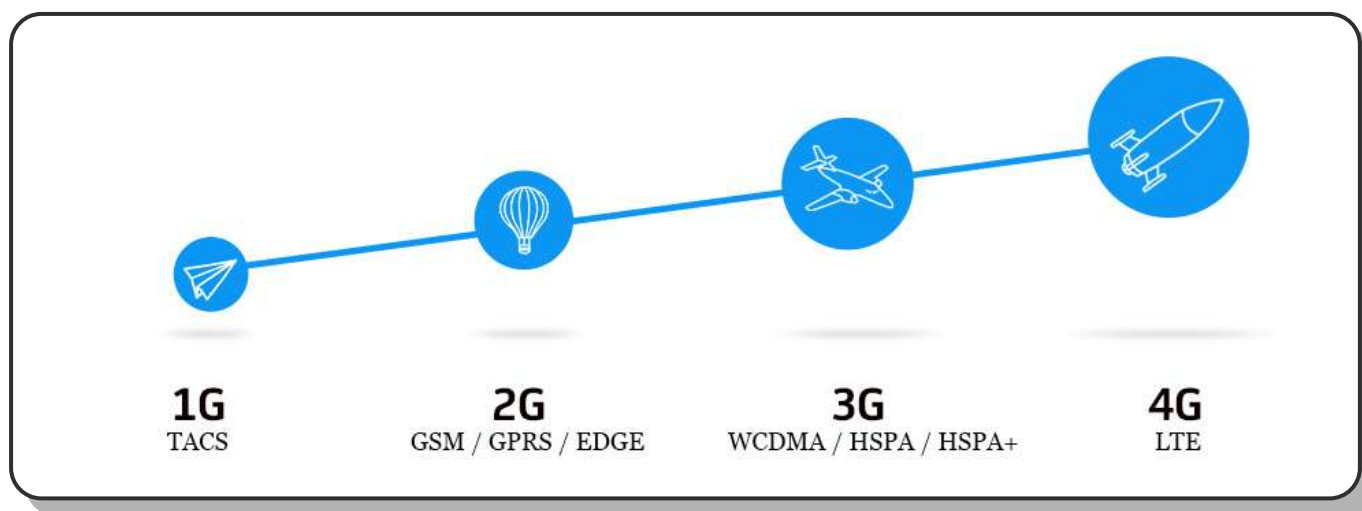
⁴Dajiang Innovation Technology Co., Ltd.

- ❑ **Volatile performance** of different implementations;
- ❑ Algorithm **updates** are very **complex** and **miscellaneous**;
- ❑ **Unfriendly** support for the latest **tricks**;
- ❑ **Incomplete** benchmark testing;
- ❑ **Expensive computational cost** of algorithm reproduction;
- ❑ Few **active** repositories;
- ❑ **High learning costs** for developers.



What is RLLTE?

- A novel **reinforcement learning** (RL) framework inspired by the **long-term evolution (LTE)** standard project in telecommunications.



- GitHub Link: <https://github.com/RLE-Foundation/rllte>



What is RLLTE for?

For Academia:









- ❑ Accelerating algorithm development;
- ❑ Tracking the latest research progress;
- ❑ Reusable and reliable baselines;



For Industry:

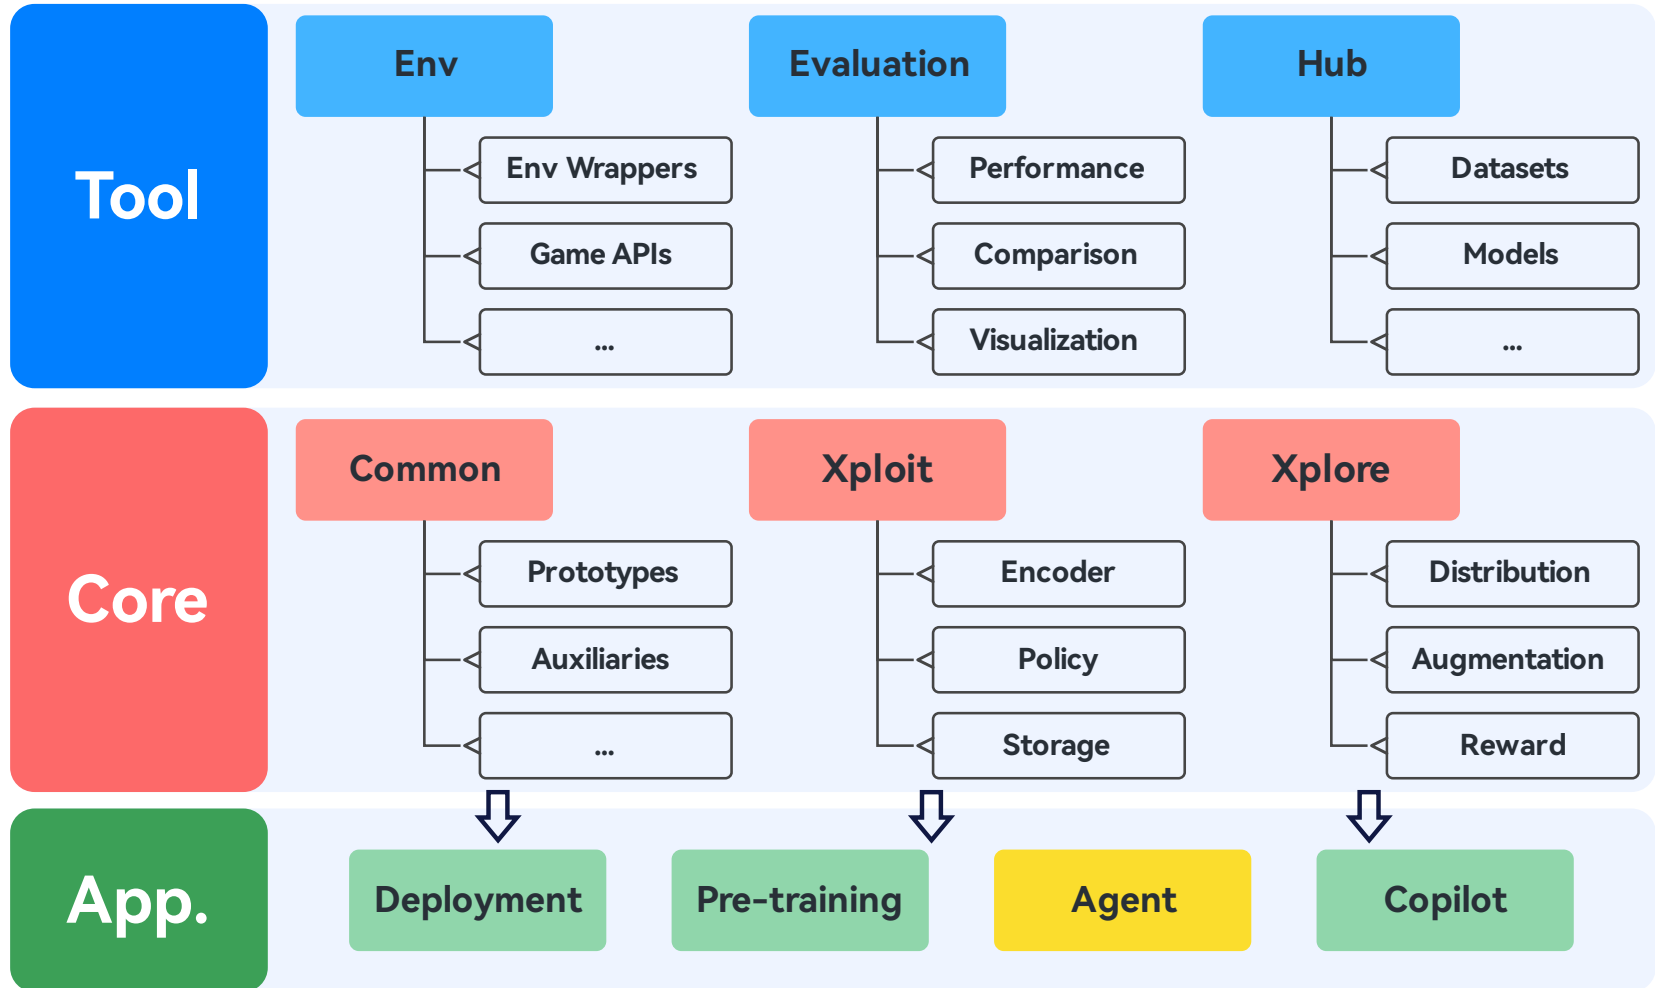
- ❑ Ultrafast application construction;
- ❑ High scalability and friendly interface;
- ❑ Convenient model deployment.



- ❑  Long-term evolution for providing latest **algorithms** and **tricks**;
- ❑  Complete **ecosystem** for task **design**, model **training**, **evaluation**, and **deployment** (TensorRT, CANN, ...);
- ❑  **Module-oriented** design for complete **decoupling** of RL algorithms;
- ❑  Optimized workflow for full hardware **acceleration**;
- ❑  Support **custom environments** and **modules**;
- ❑  Support multiple computing devices like **GPU** and **NPU**;
- ❑  Large number of reusable **benchmarks** ([rllte-hub](#));
- ❑  Large language model-empowered **copilot**.



Architecture (Overview)



- ❑ **Common:** **Prototypes** and auxiliary modules.

- ❑ **Xploit:** Modules that focus on **exploitation** in RL.
 - **Encoder:** **feature** extraction;
 - **Policy:** **interaction** and learning;
 - **Storage:** experience storage and sampling.

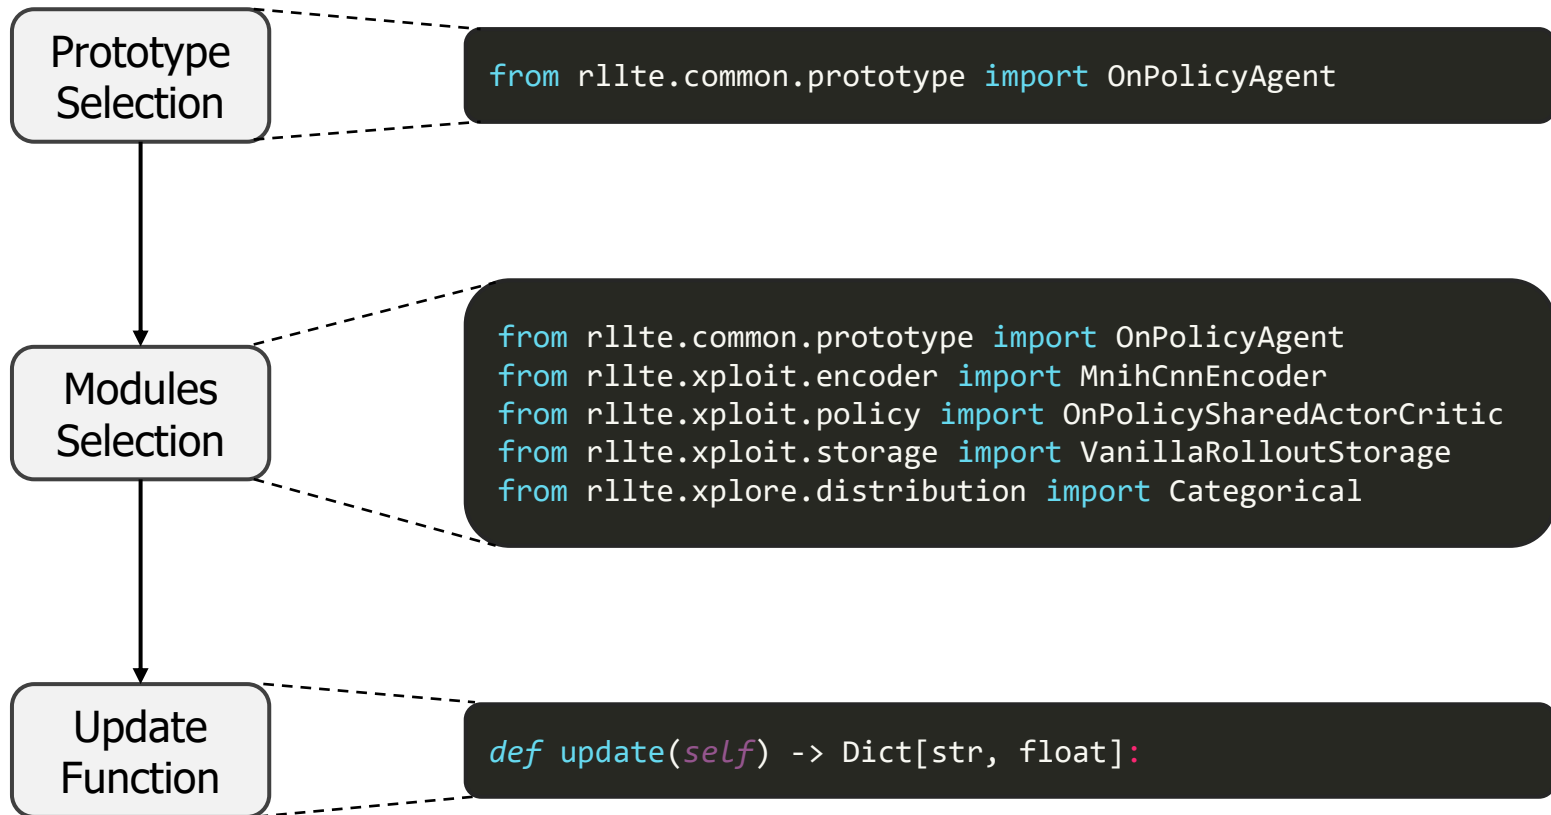
- ❑ **Xplore:** Modules that focus on **exploration** in RL.
 - **Distribution:** **action** sampling;
 - **Augmentation:** observation data augmentation;
 - **Reward:** **intrinsic reward** modules.



- ❑ **Agent:** Implemented RL Agents using RLLTE building blocks.
- ❑ **Pre-Training:** Methods of **pre-training** in RL.
- ❑ **Deployment:** Methods of model **deployment** in RL.
- ❑ **Copilot:** **LLM-based copilot** that helps developer build RL applications;
- ❑ **Hub:** Fast **training API** and reusable **benchmarks**.
- ❑ **Evaluation:** **Reasonable** and **reliable** metrics for algorithm evaluation.
- ❑ **Env:** **Packaged** environments (e.g., Atari games) for fast invocation.



- ❑ **Three steps** to implement an agent:



- ❑ RLLTE provides implementations for **well-recognized** RL algorithms and **simple interface** for building applications:

```
# import `env` and `agent` api
from rllte.env import make_dmc_env
from rllte.agent import DrQv2

if __name__ == "__main__":
    device = "cuda:0"
    # create env, `eval_env` is optional
    env = make_dmc_env(env_id="cartpole_balance", device=device)
    # create agent
    agent = DrQv2(env=env, device=device, tag="drqv2_dmc_pixel")
    # start training
    agent.train(num_train_steps=500000)
```



❑ Training Example:

```
PROBLEMS OUTPUT DEBUG CONSOLE TERMINAL PORTS  
○ (rllte) yuanmingqi@YUAN-W5:/export/yuanmingqi/code/rllte$ python test.py
```



- ❑ The **module-oriented** design allows developers to perform module replacement to make model **comparison** and **improvement**:

```
# compare the performance of different encoders
from rllte.agent import DrQv2
from rllte.xploit.encoder import MnihCnnEncoder, TassaCnnEncoder

if __name__ == "__main__":
    agent = DrQv2(...)

    encoder1 = MnihCnnEncoder(...)
    encoder2 = TassaCnnEncoder(...)

    agent.set(encoder=encoder1)
    agent.train(...)

    agent.set(encoder=encoder2)
    agent.train(...)
```



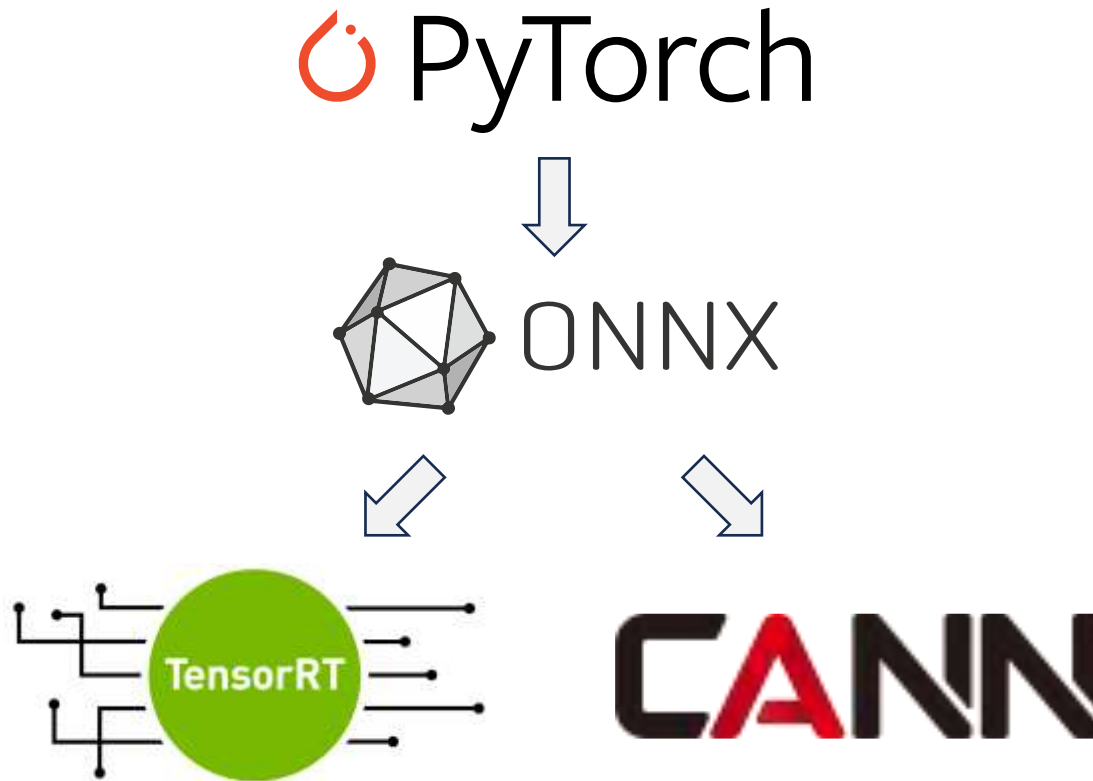
□ Pre-training Based on Intrinsic Rewards

```
from rllte.agent import PPO
from rllte.env import make_atari_env
from rllte.xplore.reward import RE3

if __name__ == "__main__":
    # env setup
    device = "cuda:0"
    env = make_atari_env(device=device)
    # create agent and turn on pre-training mode
    agent = PPO(env=env,
                device=device,
                tag="ppo_atari",
                pretraining=True)
    # create intrinsic reward
    re3 = RE3(observation_space=env.observation_space,
              action_space=env.action_space,
              device=device)
    # set the intrinsic reward module
    agent.set(reward=re3)
    # start training
    agent.train(num_train_steps=25000000)
```



- ❑ Model Deployment Based-on **TensorRT** and **CANN**



❑ LLM-Based Copilot: An attempt



👉 Ask anything you want to know about RL!

Chatbot

hi



Hi! How may I assist you today?

I want to train an PPO agent on the Procgen benchmark.

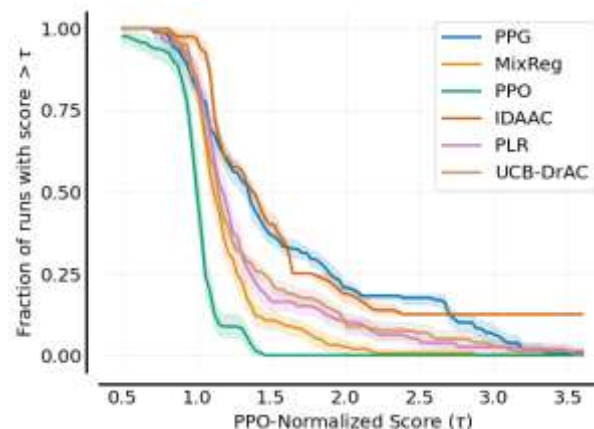



To train a PPO agent on the Procgen benchmark using RLLTE, you can use the following command: `python -m rllte.hub.apps.ppo_procgen --env-id bigfish.`



□ RLLTE provides evaluation methods based on:

Agarwal R, Schwarzer M, Castro P S, et al. Deep reinforcement learning at the edge of the statistical precipice[J]. Advances in neural information processing systems, 2021, 34: 29304-29320.



□ **Hub:** Fast training **API** and reusable **benchmarks**.

- **Datasets:** **test scores** and **learning curves** of various RL algorithms on different benchmarks.

```
from rllte.hub.datasets import Procgen
```

- **Models:** **trained models** of various RL algorithms on different benchmarks.

```
from rllte.hub.models import Procgen
```

- **Applications:** **fast-API** for training RL agents with one-line command.

```
python -m rllte.hub.apps.ppo_procgen --env_id bigfish
```



❑ Packaged environments (Part)

Function	Name	Remark
<code>make_atari_env</code>	Atari Games	Discrete control
<code>make_bullet_env</code>	PyBullet Robotics Environments	Continuous control
<code>make_dmc_env</code>	DeepMind Control Suite	Continuous control
<code>make_minigrid_env</code>	MiniGrid Games	Discrete control
<code>make_procgen_env</code>	Procgen Games	Discrete control
<code>make_robosuite_env</code>	Robosuite Robotics Environments	Continuous control



□ RLLTE Project Update **Tenet**

- **General**;
- Improvements in **sample efficiency or generalization ability**;
- Excellent **performance** on recognized benchmarks;
- Promising **tools** for RL.



- ❑ Advanced **LLM**-Based Copilot;
- ❑ Support **Multi-Agent** Reinforcement Learning;
- ❑ Support **Offline** Reinforcement Learning;
- ❑ **Hardware**-Level Code Acceleration;
- ❑ More Convenient **Interface** for Everyone;
- ❑ **General** Reinforcement Learning Model.



- ❑  GitHub Link: <https://github.com/RLE-Foundation/rllte>
- ❑  E-mail: friedrichyuan19990827@gmail.com
- ❑  Documentation: <https://docs.rllte.dev/>
- ❑  Benchmarks: <https://hub.rllte.dev/>
- ❑  Discussions: <https://github.com/RLE-Foundation/rllte/discussions>



Thanks!

