

---

# **Tianshou**

***Release 0.4.8***

**Tianshou contributors**

**May 05, 2022**



## TUTORIALS

<b>1</b>	<b>Installation</b>	<b>3</b>
<b>2</b>	<b>Indices and tables</b>	<b>143</b>
	<b>Bibliography</b>	<b>145</b>
	<b>Python Module Index</b>	<b>147</b>
	<b>Index</b>	<b>149</b>



**Tianshou** () is a reinforcement learning platform based on pure PyTorch. Unlike existing reinforcement learning libraries, which are mainly based on TensorFlow, have many nested classes, unfriendly API, or slow-speed, Tianshou provides a fast-speed framework and pythonic API for building the deep reinforcement learning agent. The supported interface algorithms include:

- *DQNPOLICY* Deep Q-Network
- *DQNPOLICY* Double DQN
- *DQNPOLICY* Dueling DQN
- *C51POLICY* Categorical DQN
- *RainbowPOLICY* Rainbow DQN
- *QRDQNPOLICY* Quantile Regression DQN
- *IQNPOLICY* Implicit Quantile Network
- *FQFPOLICY* Fully-parameterized Quantile Function
- *PGPOLICY* Policy Gradient
- *NPGPOLICY* Natural Policy Gradient
- *A2CPOLICY* Advantage Actor-Critic
- *TRPOPOLICY* Trust Region Policy Optimization
- *PPOPOLICY* Proximal Policy Optimization
- *DDPGPOLICY* Deep Deterministic Policy Gradient
- *TD3POLICY* Twin Delayed DDPG
- *SACPOLICY* Soft Actor-Critic
- *REDQPOLICY* Randomized Ensembled Double Q-Learning
- *DiscreteSACPOLICY* Discrete Soft Actor-Critic
- *ImitationPOLICY* Imitation Learning
- *BCQPOLICY* Batch-Constrained deep Q-Learning
- *CQLPOLICY* Conservative Q-Learning
- *DiscreteBCQPOLICY* Discrete Batch-Constrained deep Q-Learning
- *DiscreteCQLPOLICY* Discrete Conservative Q-Learning
- *DiscreteCRRPOLICY* Critic Regularized Regression
- *GAILPOLICY* Generative Adversarial Imitation Learning
- *PSRLPOLICY* Posterior Sampling Reinforcement Learning
- *ICMPOLICY* Intrinsic Curiosity Module
- *PrioritizedReplayBuffer* Prioritized Experience Replay
- *compute\_episodic\_return()* Generalized Advantage Estimator

Here is Tianshou's other features:

- Elegant framework, using only ~3000 lines of code
- State-of-the-art MuJoCo benchmark
- Support vectorized environment (synchronous or asynchronous) for all algorithms: *Parallel Sampling*

- Support super-fast vectorized environment `EnvPool` for all algorithms: *EnvPool Integration*
- Support recurrent state representation in actor network and critic network (RNN-style training for POMDP): *RNN-style Training*
- Support any type of environment state/action (e.g. a dict, a self-defined class, ...): *User-defined Environment and Different State Representation*
- Support *Customize Training Process*
- Support n-step returns estimation `compute_nstep_return()` and prioritized experience replay `PrioritizedReplayBuffer` for all Q-learning based algorithms; GAE, nstep and PER are very fast thanks to numba jit function and vectorized numpy operation
- Support *Multi-Agent RL*
- Support both `TensorBoard` and `W&B` log tools
- Support multi-GPU training *Multi-GPU Training*
- Comprehensive `unit tests`, including functional checking, RL pipeline checking, documentation checking, PEP8 code-style checking, and type checking

<https://tianshou.readthedocs.io/zh/master/>

## INSTALLATION

Tianshou is currently hosted on [PyPI](#) and [conda-forge](#). It requires Python  $\geq 3.6$ .

You can simply install Tianshou from PyPI with the following command:

```
$ pip install tianshou
```

If you use Anaconda or Miniconda, you can install Tianshou from conda-forge through the following command:

```
$ conda -c conda-forge install tianshou
```

You can also install with the newest version through GitHub:

```
$ pip install git+https://github.com/thu-ml/tianshou.git@master --upgrade
```

After installation, open your python console and type

```
import tianshou
print(tianshou.__version__)
```

If no error occurs, you have successfully installed Tianshou.

Tianshou is still under development, you can also check out the documents in stable version through [tianshou.readthedocs.io/en/stable/](https://tianshou.readthedocs.io/en/stable/).

### 1.1 Get Started with Jupyter Notebook

In this tutorial, we will use Google Colaboratory to show you the most basic usages of common building blocks in Tianshou. You will be guided step by step to see how different modules in Tianshou collaborate with each other to conduct a classic DRL experiment (PPO algorithm for CartPole-v0 environment).

- [L0: Overview](#)
- [L1: Batch](#)
- [L2: Replay Buffer](#)
- [L3: Vectorized Environment](#)
- [L4: Policy](#)
- [L5: Collector](#)
- [L6: Trainer](#)
- [L7: Experiment](#)

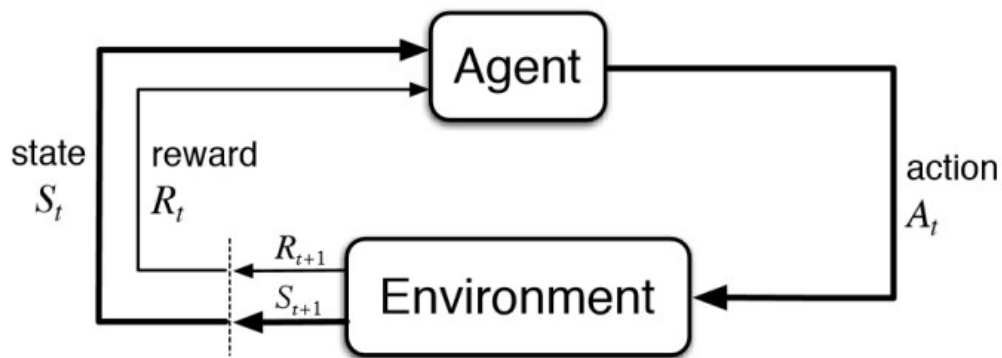
## 1.2 Deep Q Network

Deep reinforcement learning has achieved significant successes in various applications. **Deep Q Network** (DQN) [MKS+15] is the pioneer one. In this tutorial, we will show how to train a DQN agent on CartPole with Tianshou step by step. The full script is at [test/discrete/test\\_dqn.py](#).

Contrary to existing Deep RL libraries such as **RLlib**, which could only accept a config specification of hyperparameters, network, and others, Tianshou provides an easy way of construction through the code-level.

### 1.2.1 Overview

In reinforcement learning, the agent interacts with environments to improve itself.

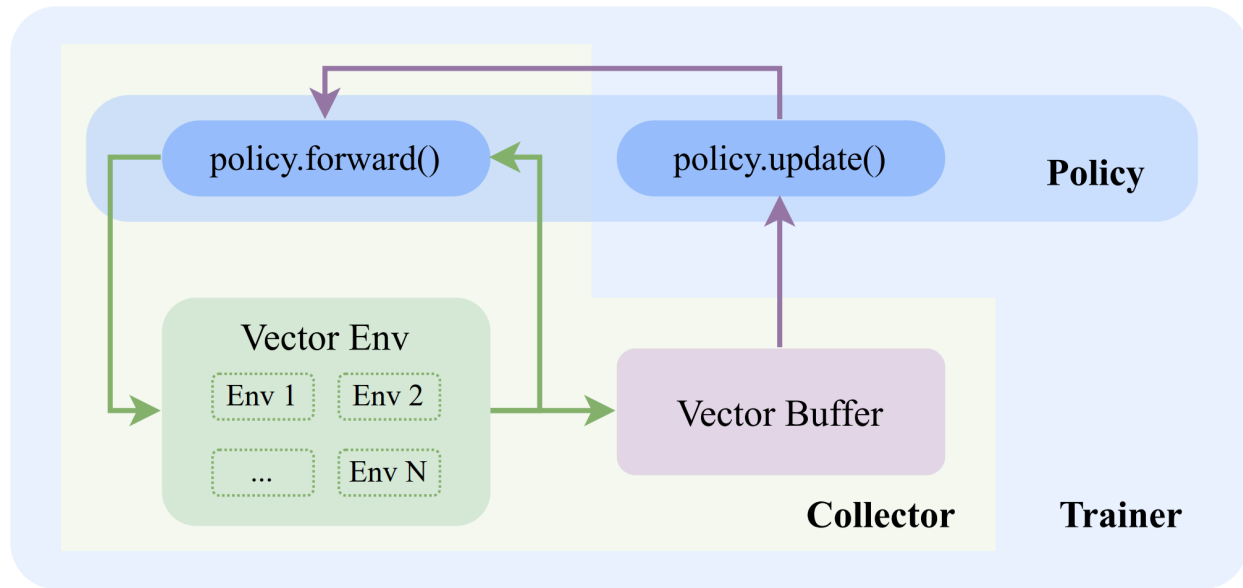


There are three types of data flow in RL training pipeline:

1. Agent to environment: `action` will be generated by agent and sent to environment;
2. Environment to agent: `env.step` takes action, and returns a tuple of (`observation`, `reward`, `done`, `info`);
3. Agent-environment interaction to agent training: the data generated by interaction will be stored and sent to the learner of agent.

In the following sections, we will set up (vectorized) environments, policy (with neural network), collector (with buffer), and trainer to successfully run the RL training and evaluation pipeline. Here is the overall system:





### 1.2.2 Make an Environment

First of all, you have to make an environment for your agent to interact with. You can use `gym.make(environment_name)` to make an environment for your agent. For environment interfaces, we follow the convention of [OpenAI Gym](#). In your Python code, simply import Tianshou and make the environment:

```
import gym
import tianshou as ts

env = gym.make('CartPole-v0')
```

`CartPole-v0` includes a cart carrying a pole moving on a track. This is a simple environment with a discrete action space, for which DQN applies. You have to identify whether the action space is continuous or discrete and apply eligible algorithms. DDPG [LHP+16], for example, could only be applied to continuous action spaces, while almost all other policy gradient methods could be applied to both.

Here is the detail of useful fields of `CartPole-v0`:

- `state`: the position of the cart, the velocity of the cart, the angle of the pole and the velocity of the tip of the pole;
- `action`: can only be one of `[0, 1, 2]`, for moving the cart left, no move, and right;
- `reward`: each timestep you last, you will receive a +1 reward;
- `done`: if `CartPole` is out-of-range or timeout (the pole is more than 15 degrees from vertical, or the cart moves more than 2.4 units from the center, or you last over 200 timesteps);
- `info`: extra info from environment simulation.

The goal is to train a good policy that can get the highest reward in this environment.

### 1.2.3 Setup Vectorized Environment

If you want to use the original `gym.Env`:

```
train_envs = gym.make('CartPole-v0')
test_envs = gym.make('CartPole-v0')
```

Tianshou supports vectorized environment for all algorithms. It provides four types of vectorized environment wrapper:

- *DummyVectorEnv*: the sequential version, using a single-thread for-loop;
- *SubprocVectorEnv*: use python multiprocessing and pipe for concurrent execution;
- *ShmemVectorEnv*: use share memory instead of pipe based on *SubprocVectorEnv*;
- *RayVectorEnv*: use Ray for concurrent activities and is currently the only choice for parallel simulation in a cluster with multiple machines. It can be used as follows: (more explanation can be found at [Parallel Sampling](#))

```
train_envs = ts.env.DummyVectorEnv([lambda: gym.make('CartPole-v0') for _ in
    range(10)])
test_envs = ts.env.DummyVectorEnv([lambda: gym.make('CartPole-v0') for _ in
    range(100)])
```

Here, we set up 10 environments in `train_envs` and 100 environments in `test_envs`.

You can also try the super-fast vectorized environment *EnvPool* by

```
import envpool
train_envs = envpool.make_gym("CartPole-v0", num_envs=10)
test_envs = envpool.make_gym("CartPole-v0", num_envs=100)
```

For the demonstration, here we use the second code-block.

**Warning:** If you use your own environment, please make sure the `seed` method is set up properly, e.g.,

```
def seed(self, seed):
    np.random.seed(seed)
```

Otherwise, the outputs of these envs may be the same with each other.

### 1.2.4 Build the Network

Tianshou supports any user-defined PyTorch networks and optimizers. Yet, of course, the inputs and outputs must comply with Tianshou's API. Here is an example:

```
import torch, numpy as np
from torch import nn

class Net(nn.Module):
    def __init__(self, state_shape, action_shape):
        super().__init__()
        self.model = nn.Sequential(
            nn.Linear(np.prod(state_shape), 128), nn.ReLU(inplace=True),
            nn.Linear(128, 128), nn.ReLU(inplace=True),
            nn.Linear(128, 128), nn.ReLU(inplace=True),
            nn.Linear(128, np.prod(action_shape)),
```

(continues on next page)

(continued from previous page)

```

    )

    def forward(self, obs, state=None, info={}):
        if not isinstance(obs, torch.Tensor):
            obs = torch.tensor(obs, dtype=torch.float)
        batch = obs.shape[0]
        logits = self.model(obs.view(batch, -1))
        return logits, state

state_shape = env.observation_space.shape or env.observation_space.n
action_shape = env.action_space.shape or env.action_space.n
net = Net(state_shape, action_shape)
optim = torch.optim.Adam(net.parameters(), lr=1e-3)

```

You can also use pre-defined MLP networks in *common*, *discrete*, and *continuous*. The rules of self-defined networks are:

1. Input: observation `obs` (may be a `numpy.ndarray`, `torch.Tensor`, dict, or self-defined class), hidden state `state` (for RNN usage), and other information `info` provided by the environment.
2. Output: some logits, the next hidden state `state`. The logits could be a tuple instead of a `torch.Tensor`, or some other useful variables or results during the policy forwarding procedure. It depends on how the policy class process the network output. For example, in PPO [SWD+17], the return of the network might be `(mu, sigma)`, `state` for Gaussian policy.

**Note:** The logits here indicates the raw output of the network. In supervised learning, the raw output of prediction/classification model is called logits, and here we extend this definition to any raw output of the neural network.

## 1.2.5 Setup Policy

We use the defined `net` and `optim` above, with extra policy hyper-parameters, to define a policy. Here we define a DQN policy with a target network:

```

policy = ts.policy.DQNPolicy(net, optim, discount_factor=0.9, estimation_step=3,
    ↪target_update_freq=320)

```

## 1.2.6 Setup Collector

The collector is a key concept in Tianshou. It allows the policy to interact with different types of environments conveniently. In each step, the collector will let the policy perform (at least) a specified number of steps or episodes and store the data in a replay buffer.

The following code shows how to set up a collector in practice. It is worth noticing that `VectorReplayBuffer` is to be used in vectorized environment scenarios, and the number of buffers, in the following case 10, is preferred to be set as the number of environments.

```

train_collector = ts.data.Collector(policy, train_envs, ts.data.
    ↪VectorReplayBuffer(20000, 10), exploration_noise=True)
test_collector = ts.data.Collector(policy, test_envs, exploration_noise=True)

```

The main function of collector is the `collect` function, which can be summarized in the following lines:

```

result = self.policy(self.data, last_state) # the agent
↳ predicts the batch action from batch observation
act = to_numpy(result.act)
self.data.update(act=act) # update the data
↳ with new action/policy
result = self.env.step(act, ready_env_ids) # apply action to
↳ environment
obs_next, rew, done, info = result
self.data.update(obs_next=obs_next, rew=rew, done=done, info=info) # update the data
↳ with new state/reward/done/info

```

## 1.2.7 Train Policy with a Trainer

Tianshou provides `onpolicy_trainer()`, `offpolicy_trainer()`, and `offline_trainer()`. The trainer will automatically stop training when the policy reach the stop condition `stop_fn` on test collector. Since DQN is an off-policy algorithm, we use the `offpolicy_trainer()` as follows:

```

result = ts.trainer.offpolicy_trainer(
    policy, train_collector, test_collector,
    max_epoch=10, step_per_epoch=10000, step_per_collect=10,
    update_per_step=0.1, episode_per_test=100, batch_size=64,
    train_fn=lambda epoch, env_step: policy.set_eps(0.1),
    test_fn=lambda epoch, env_step: policy.set_eps(0.05),
    stop_fn=lambda mean_rewards: mean_rewards >= env.spec.reward_threshold)
print(f'Finished training! Use {result["duration"]}')

```

The meaning of each parameter is as follows (full description can be found at `offpolicy_trainer()`):

- `max_epoch`: The maximum of epochs for training. The training process might be finished before reaching the `max_epoch`;
- `step_per_epoch`: The number of environment step (a.k.a. transition) collected per epoch;
- `step_per_collect`: The number of transition the collector would collect before the network update. For example, the code above means “collect 10 transitions and do one policy network update”;
- `episode_per_test`: The number of episodes for one policy evaluation.
- `batch_size`: The batch size of sample data, which is going to feed in the policy network.
- `train_fn`: A function receives the current number of epoch and step index, and performs some operations at the beginning of training in this epoch. For example, the code above means “reset the epsilon to 0.1 in DQN before training”.
- `test_fn`: A function receives the current number of epoch and step index, and performs some operations at the beginning of testing in this epoch. For example, the code above means “reset the epsilon to 0.05 in DQN before testing”.
- `stop_fn`: A function receives the average undiscounted returns of the testing result, return a boolean which indicates whether reaching the goal.
- `logger`: See below.

The trainer supports `TensorBoard` for logging. It can be used as:

```

from torch.utils.tensorboard import SummaryWriter
from tianshou.utils import TensorboardLogger
writer = SummaryWriter('log/dqn')
logger = TensorboardLogger(writer)

```

Pass the logger into the trainer, and the training result will be recorded into the TensorBoard.

The returned result is a dictionary as follows:

```
{
  'train_step': 9246,
  'train_episode': 504.0,
  'train_time/collector': '0.65s',
  'train_time/model': '1.97s',
  'train_speed': '3518.79 step/s',
  'test_step': 49112,
  'test_episode': 400.0,
  'test_time': '1.38s',
  'test_speed': '35600.52 step/s',
  'best_reward': 199.03,
  'duration': '4.01s'
}
```

It shows that within approximately 4 seconds, we finished training a DQN agent on CartPole. The mean returns over 100 consecutive episodes is 199.03.

## 1.2.8 Save/Load Policy

Since the policy inherits the class `torch.nn.Module`, saving and loading the policy are exactly the same as a torch module:

```
torch.save(policy.state_dict(), 'dqn.pth')
policy.load_state_dict(torch.load('dqn.pth'))
```

## 1.2.9 Watch the Agent's Performance

`Collector` supports rendering. Here is the example of watching the agent's performance in 35 FPS:

```
policy.eval()
policy.set_eps(0.05)
collector = ts.data.Collector(policy, env, exploration_noise=True)
collector.collect(n_episode=1, render=1 / 35)
```

## 1.2.10 Train a Policy with Customized Codes

“I don't want to use your provided trainer. I want to customize it!”

Tianshou supports user-defined training code. Here is the code snippet:

```
# pre-collect at least 5000 transitions with random action before training
train_collector.collect(n_step=5000, random=True)

policy.set_eps(0.1)
for i in range(int(1e6)): # total step
    collect_result = train_collector.collect(n_step=10)

    # once if the collected episodes' mean returns reach the threshold,
    # or every 1000 steps, we test it on test_collector
    if collect_result['rews'].mean() >= env.spec.reward_threshold or i % 1000 == 0:
```

(continues on next page)

(continued from previous page)

```

policy.set_eps(0.05)
result = test_collector.collect(n_episode=100)
if result['rews'].mean() >= env.spec.reward_threshold:
    print(f'Finished training! Test mean returns: {result["rews"].mean()}')
    break
else:
    # back to training eps
    policy.set_eps(0.1)

# train policy with a sampled batch data from buffer
losses = policy.update(64, train_collector.buffer)

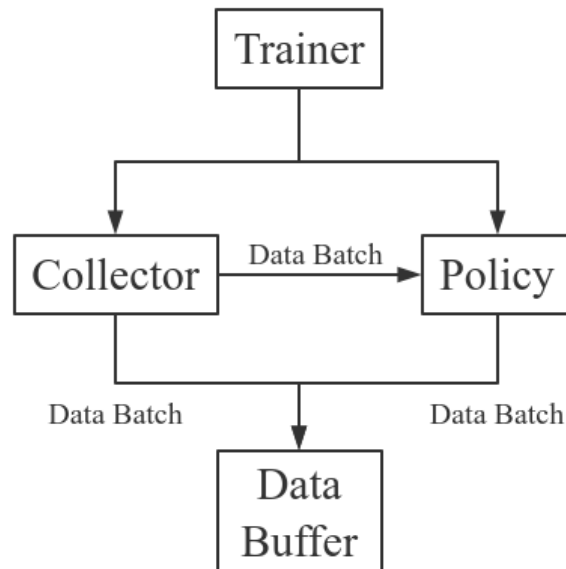
```

For further usage, you can refer to the [Cheat Sheet](#).

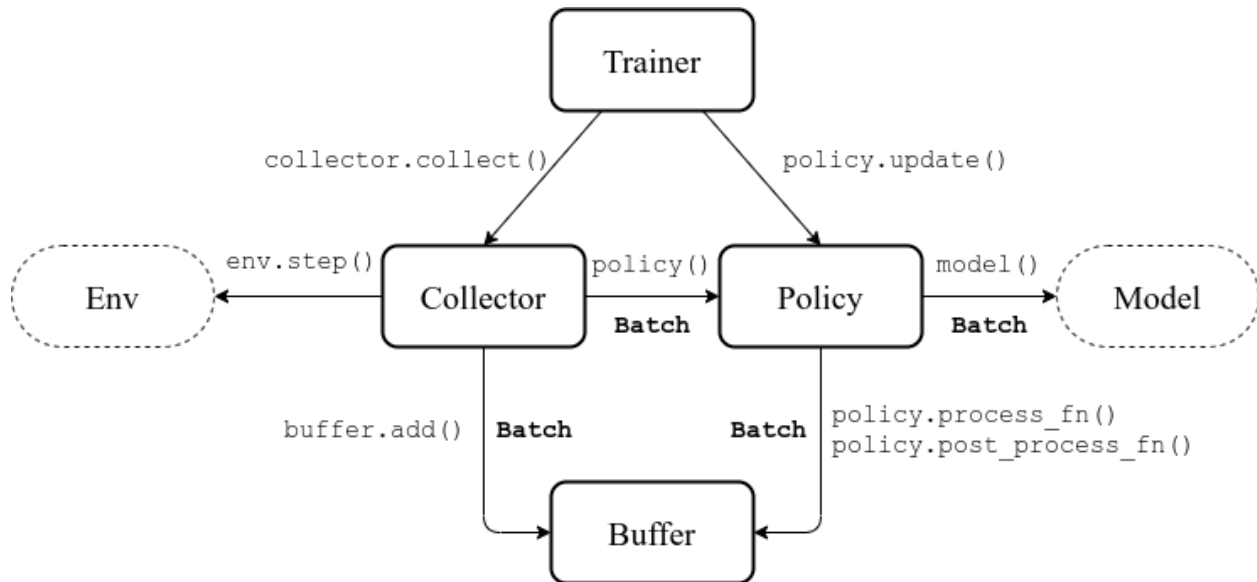
## References

### 1.3 Basic concepts in Tianshou

Tianshou splits a Reinforcement Learning agent training procedure into these parts: trainer, collector, policy, and data buffer. The general control flow can be described as:



Here is a more detailed description, where `Env` is the environment and `Model` is the neural network:



### 1.3.1 Batch

Tianshou provides `Batch` as the internal data structure to pass any kind of data to other methods, for example, a collector gives a `Batch` to policy for learning. Let's take a look at this script:

```

>>> import torch, numpy as np
>>> from tianshou.data import Batch
>>> data = Batch(a=4, b=[5, 5], c='2312312', d=('a', -2, -3))
>>> # the list will automatically be converted to numpy array
>>> data.b
array([5, 5])
>>> data.b = np.array([3, 4, 5])
>>> print(data)
Batch(
  a: 4,
  b: array([3, 4, 5]),
  c: '2312312',
  d: array(['a', '-2', '-3'], dtype=object),
)
>>> data = Batch(obs={'index': np.zeros((2, 3))}, act=torch.zeros((2, 2)))
>>> data[:, 1] += 6
>>> print(data[-1])
Batch(
  obs: Batch(
    index: array([0., 6., 0.]),
  ),
  act: tensor([0., 6.]),
)

```

In short, you can define a `Batch` with any key-value pair, and perform some common operations over it.

[Understand Batch](#) is a dedicated tutorial for `Batch`. We strongly recommend every user to read it so as to correctly understand and use `Batch`.

### 1.3.2 Buffer

*ReplayBuffer* stores data generated from interaction between the policy and environment. *ReplayBuffer* can be considered as a specialized form (or management) of *Batch*. It stores all the data in a batch with circular-queue style.

The current implementation of Tianshou typically use 7 reserved keys in *Batch*:

- `obs` the observation of step  $t$  ;
- `act` the action of step  $t$  ;
- `rew` the reward of step  $t$  ;
- `done` the done flag of step  $t$  ;
- `obs_next` the observation of step  $t + 1$  ;
- `info` the info of step  $t$  (in `gym.Env`, the `env.step()` function returns 4 arguments, and the last one is `info`);
- `policy` the data computed by policy in step  $t$ ;

The following code snippet illustrates its usage, including:

- the basic data storage: `add()` ;
- get attribute, get slicing data, ... ;
- sample from buffer: `sample_indices(batch_size)` and `sample(batch_size)` ;
- get previous/next transition index within episodes: `prev(index)` and `next(index)` ;
- save/load data from buffer: `pickle` and `HDF5` ;

```
>>> import pickle, numpy as np
>>> from tianshou.data import Batch, ReplayBuffer
>>> buf = ReplayBuffer(size=20)
>>> for i in range(3):
...     buf.add(Batch(obs=i, act=i, rew=i, done=0, obs_next=i + 1, info={}))

>>> buf.obs
# since we set size = 20, len(buf.obs) == 20.
array([0, 1, 2, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0])
>>> # but there are only three valid items, so len(buf) == 3.
>>> len(buf)
3
>>> # save to file "buf.pkl"
>>> pickle.dump(buf, open('buf.pkl', 'wb'))
>>> # save to HDF5 file
>>> buf.save_hdf5('buf.hdf5')

>>> buf2 = ReplayBuffer(size=10)
>>> for i in range(15):
...     done = i % 4 == 0
...     buf2.add(Batch(obs=i, act=i, rew=i, done=done, obs_next=i + 1, info={}))
>>> len(buf2)
10
>>> buf2.obs
# since its size = 10, it only stores the last 10 steps' result.
array([10, 11, 12, 13, 14, 5, 6, 7, 8, 9])

>>> # move buf2's result into buf (meanwhile keep it chronologically)
>>> buf.update(buf2)
```

(continues on next page)



(continued from previous page)

```

>>> buf.obs
array([ 0,  1,  2,  5,  6,  7,  8,  9, 10, 11, 12, 13, 14,  0,  0,  0,
        0,  0,  0,  0])

>>> # get all available index by using batch_size = 0
>>> indices = buf.sample_indices(0)
>>> indices
array([ 0,  1,  2,  3,  4,  5,  6,  7,  8,  9, 10, 11, 12])
>>> # get one step previous/next transition
>>> buf.prev(indices)
array([ 0,  0,  1,  2,  3,  4,  5,  7,  7,  8,  9, 11, 11])
>>> buf.next(indices)
array([ 1,  2,  3,  4,  5,  6,  6,  8,  9, 10, 10, 12, 12])

>>> # get a random sample from buffer
>>> # the batch_data is equal to buf[indices].
>>> batch_data, indices = buf.sample(batch_size=4)
>>> batch_data.obs == buf[indices].obs
array([ True,  True,  True,  True])
>>> len(buf)
13

>>> buf = pickle.load(open('buf.pkl', 'rb')) # load from "buf.pkl"
>>> len(buf)
3
>>> # load complete buffer from HDF5 file
>>> buf = ReplayBuffer.load_hdf5('buf.hdf5')
>>> len(buf)
3

```

`ReplayBuffer` also supports `frame_stack` sampling (typically for RNN usage, see issue#19), ignoring storing the next observation (save memory in Atari tasks), and multi-modal observation (see issue#38):

```

>>> buf = ReplayBuffer(size=9, stack_num=4, ignore_obs_next=True)
>>> for i in range(16):
...     done = i % 5 == 0
...     ptr, ep_rew, ep_len, ep_idx = buf.add(
...         Batch(obs={'id': i}, act=i, rew=i,
...               done=done, obs_next={'id': i + 1}))
...     print(i, ep_len, ep_rew)
0 [1] [0.]
1 [0] [0.]
2 [0] [0.]
3 [0] [0.]
4 [0] [0.]
5 [5] [15.]
6 [0] [0.]
7 [0] [0.]
8 [0] [0.]
9 [0] [0.]
10 [5] [40.]
11 [0] [0.]
12 [0] [0.]
13 [0] [0.]
14 [0] [0.]
15 [5] [65.]

```

(continues on next page)

(continued from previous page)

```

>>> print(buf) # you can see obs_next is not saved in buf
ReplayBuffer(
  obs: Batch(
    id: array([ 9, 10, 11, 12, 13, 14, 15,  7,  8]),
  ),
  act: array([ 9, 10, 11, 12, 13, 14, 15,  7,  8]),
  rew: array([ 9., 10., 11., 12., 13., 14., 15.,  7.,  8.]),
  done: array([False, True, False, False, False, False, True, False,
               False]),
)
>>> index = np.arange(len(buf))
>>> print(buf.get(index, 'obs').id)
[[ 7  7  8  9]
 [ 7  8  9 10]
 [11 11 11 11]
 [11 11 11 12]
 [11 11 12 13]
 [11 12 13 14]
 [12 13 14 15]
 [ 7  7  7  7]
 [ 7  7  7  8]]
>>> # here is another way to get the stacked data
>>> # (stack only for obs and obs_next)
>>> abs(buf.get(index, 'obs')['id'] - buf[index].obs.id).sum().sum()
0
>>> # we can get obs_next through __getitem__, even if it doesn't exist
>>> # however, [:] will select the item according to timestamp,
>>> # that equals to index == [7, 8, 0, 1, 2, 3, 4, 5, 6]
>>> print(buf[:].obs_next.id)
[[ 7  7  7  8]
 [ 7  7  8  9]
 [ 7  8  9 10]
 [ 7  8  9 10]
 [11 11 11 12]
 [11 11 12 13]
 [11 12 13 14]
 [12 13 14 15]
 [12 13 14 15]]
>>> full_index = np.array([7, 8, 0, 1, 2, 3, 4, 5, 6])
>>> np.allclose(buf[:].obs_next.id, buf[full_index].obs_next.id)
True

```

Tianshou provides other type of data buffer such as *PrioritizedReplayBuffer* (based on Segment Tree and `numpy.ndarray`) and *VectorReplayBuffer* (add different episodes' data but without losing chronological order). Check out *ReplayBuffer* for more detail.

### 1.3.3 Policy

Tianshou aims to modularize RL algorithms. It comes into several classes of policies in Tianshou. All of the policy classes must inherit `BasePolicy`.

A policy class typically has the following parts:

- `__init__()`: initialize the policy, including copying the target network and so on;
- `forward()`: compute action with given observation;
- `process_fn()`: pre-process data from the replay buffer;
- `learn()`: update policy with a given batch of data.
- `post_process_fn()`: update the buffer with a given batch of data.
- `update()`: the main interface for training. This function samples data from buffer, pre-process data (such as computing n-step return), learn with the data, and finally post-process the data (such as updating prioritized replay buffer); in short, `process_fn` -> `learn` -> `post_process_fn`.

#### States for policy

During the training process, the policy has two main states: training state and testing state. The training state can be further divided into the collecting state and updating state.

The meaning of training and testing state is obvious: the agent interacts with environment, collects training data and performs update, that's training state; the testing state is to evaluate the performance of the current policy during training process.

As for the collecting state, it is defined as interacting with environments and collecting training data into the buffer; we define the updating state as performing a model update by `update()` during training process.

In order to distinguish these states, you can check the policy state by `policy.training` and `policy.updating`. The state setting is as follows:

State for policy		<code>policy.training</code>	<code>policy.updating</code>
Training state	Collecting state	True	False
	Updating state	True	True
Testing state		False	False

`policy.updating` is helpful to distinguish the different exploration state, for example, in DQN we don't have to use epsilon-greedy in a pure network update, so `policy.updating` is helpful for setting epsilon in this case.

#### `policy.forward`

The `forward` function computes the action over given observations. The input and output is algorithm-specific but generally, the function is a mapping of `(batch, state, ...)` -> `batch`.

The input batch is the environment data (e.g., observation, reward, done flag and info). It comes from either `collect()` or `sample()`. The first dimension of all variables in the input batch should be equal to the batch-size.

The output is also a Batch which must contain "act" (action) and may contain "state" (hidden state of policy), "policy" (the intermediate result of policy which needs to save into the buffer, see `forward()`), and some other algorithm-specific keys.

For example, if you try to use your policy to evaluate one episode (and don't want to use `collect()`), use the following code-snippet:

```
# assume env is a gym.Env
obs, done = env.reset(), False
while not done:
    batch = Batch(obs=[obs]) # the first dimension is batch-size
    act = policy(batch).act[0] # policy.forward return a batch, use ".act" to
    ↪extract the action
    obs, rew, done, info = env.step(act)
```

Here, `Batch(obs=[obs])` will automatically create the 0-dimension to be the batch-size. Otherwise, the network cannot determine the batch-size.

## policy.process\_fn

The `process_fn` function computes some variables that depends on time-series. For example, compute the N-step or GAE returns.

Take 2-step return DQN as an example. The 2-step return DQN compute each transition's return as:

$$G_t = r_t + \gamma r_{t+1} + \gamma^2 \max_a Q(s_{t+2}, a)$$

where  $\gamma$  is the discount factor,  $\gamma \in [0, 1]$ . Here is the pseudocode showing the training process **without Tianshou framework**:

```
# pseudocode, cannot work
obs = env.reset()
buffer = Buffer(size=10000)
agent = DQN()
for i in range(int(1e6)):
    act = agent.compute_action(obs)
    obs_next, rew, done, _ = env.step(act)
    buffer.store(obs, act, obs_next, rew, done)
    obs = obs_next
    if i % 1000 == 0:
        b_s, b_a, b_s_, b_r, b_d = buffer.get(size=64)
        # compute 2-step returns. How?
        b_ret = compute_2_step_return(buffer, b_r, b_d, ...)
        # update DQN policy
        agent.update(b_s, b_a, b_s_, b_r, b_d, b_ret)
```

Thus, we need a time-related interface for calculating the 2-step return. `process_fn()` finishes this work by providing the replay buffer, the sample index, and the sample batch data. Since we store all the data in the order of time, you can simply compute the 2-step return as:

```
class DQN_2step(BasePolicy):
    """some code"""

    def process_fn(self, batch, buffer, indices):
        buffer_len = len(buffer)
        batch_2 = buffer[(indices + 2) % buffer_len]
        # this will return a batch data where batch_2.obs is s_t+2
        # we can also get s_t+2 through:
        # batch_2_obs = buffer.obs[(indices + 2) % buffer_len]
        # in short, buffer.obs[i] is equal to buffer[i].obs, but the former is more
        ↪effecient.
        Q = self(batch_2, eps=0) # shape: [batchsize, action_shape]
        maxQ = Q.max(dim=-1)
```

(continues on next page)

(continued from previous page)

```

batch.returns = batch.rew \
    + self._gamma * buffer.rew[(indices + 1) % buffer_len] \
    + self._gamma ** 2 * maxQ
return batch

```

This code does not consider the done flag, so it may not work very well. It shows two ways to get  $s_{t+2}$  from the replay buffer easily in `process_fn()`.

For other method, you can check out [tianshou.policy](#). We give the usage of policy class a high-level explanation in [A High-level Explanation](#).

### 1.3.4 Collector

The `Collector` enables the policy to interact with different types of environments conveniently.

`collect()` is the main method of `Collector`: it let the policy perform a specified number of step `n_step` or episode `n_episode` and store the data in the replay buffer, then return the statistics of the collected data such as episode's total reward.

The general explanation is listed in [A High-level Explanation](#). Other usages of collector are listed in [Collector](#) documentation. Here are some example usages:

```

policy = PGPolicy(...) # or other policies if you wish
env = gym.make("CartPole-v0")

replay_buffer = ReplayBuffer(size=10000)

# here we set up a collector with a single environment
collector = Collector(policy, env, buffer=replay_buffer)

# the collector supports vectorized environments as well
vec_buffer = VectorReplayBuffer(total_size=10000, buffer_num=3)
# buffer_num should be equal to (suggested) or larger than #envs
envs = DummyVectorEnv([lambda: gym.make("CartPole-v0") for _ in range(3)])
collector = Collector(policy, envs, buffer=vec_buffer)

# collect 3 episodes
collector.collect(n_episode=3)
# collect at least 2 steps
collector.collect(n_step=2)
# collect episodes with visual rendering ("render" is the sleep time between
# rendering consecutive frames)
collector.collect(n_episode=1, render=0.03)

```

There is also another type of collector `AsyncCollector` which supports asynchronous environment setting (for those taking a long time to step). However, `AsyncCollector` only supports **at least** `n_step` or `n_episode` collection due to the property of asynchronous environments.

### 1.3.5 Trainer

Once you have a collector and a policy, you can start writing the training method for your RL agent. Trainer, to be honest, is a simple wrapper. It helps you save energy for writing the training loop. You can also construct your own trainer: *Train a Policy with Customized Codes*.

Tianshou has three types of trainer: `onpolicy_trainer()` for on-policy algorithms such as Policy Gradient, `offpolicy_trainer()` for off-policy algorithms such as DQN, and `offline_trainer()` for offline algorithms such as BCQ. Please check out `tianshou.trainer` for the usage.

We also provide the corresponding iterator-based trainer classes `OnpolicyTrainer`, `OffpolicyTrainer`, `OfflineTrainer` to facilitate users writing more flexible training logic:

```
trainer = OnpolicyTrainer(...)
for epoch, epoch_stat, info in trainer:
    print(f"Epoch: {epoch}")
    print(epoch_stat)
    print(info)
    do_something_with_policy()
    query_something_about_policy()
    make_a_plot_with(epoch_stat)
    display(info)

# or even iterate on several trainers at the same time

trainer1 = OnpolicyTrainer(...)
trainer2 = OnpolicyTrainer(...)
for result1, result2, ... in zip(trainer1, trainer2, ...):
    compare_results(result1, result2, ...)
```

### 1.3.6 A High-level Explanation

We give a high-level explanation through the pseudocode used in section `policy.process_fn`:

<code># pseudocode, cannot work</code>	<code># methods in tianshou</code>
<code>obs = env.reset()</code>	
<code>buffer = Buffer(size=10000)</code>	<code># buffer = tianshou.</code>
<code>↪data.ReplayBuffer(size=10000)</code>	
<code>agent = DQN()</code>	<code># policy.__init__(...)</code>
<code>for i in range(int(1e6)):</code>	<code># done in trainer</code>
<code>act = agent.compute_action(obs)</code>	<code># act = policy(batch, _</code>
<code>↪...).act</code>	
<code>obs_next, rew, done, _ = env.step(act)</code>	<code># collector.collect(..</code>
<code>↪.)</code>	
<code>buffer.store(obs, act, obs_next, rew, done)</code>	<code># collector.collect(..</code>
<code>↪.)</code>	
<code>obs = obs_next</code>	<code># collector.collect(..</code>
<code>↪.)</code>	
<code>if i % 1000 == 0:</code>	<code># done in trainer</code>
	<code># the following is _</code>
<code>↪done in policy.update(batch_size, buffer)</code>	
<code>b_s, b_a, b_s_, b_r, b_d = buffer.get(size=64)</code>	<code># batch, indices = _</code>
<code>↪buffer.sample(batch_size)</code>	
<code># compute 2-step returns. How?</code>	
<code>b_ret = compute_2_step_return(buffer, b_r, b_d, ...)</code>	<code># policy.process_</code>
<code>↪fn(batch, buffer, indices)</code>	

(continues on next page)

(continued from previous page)

```
# update DQN policy
agent.update(b_s, b_a, b_s_, b_r, b_d, b_ret)      # policy.learn(batch, u
↪ ...)
```

### 1.3.7 Conclusion

So far, we go through the overall framework of Tianshou. Really simple, isn't it?

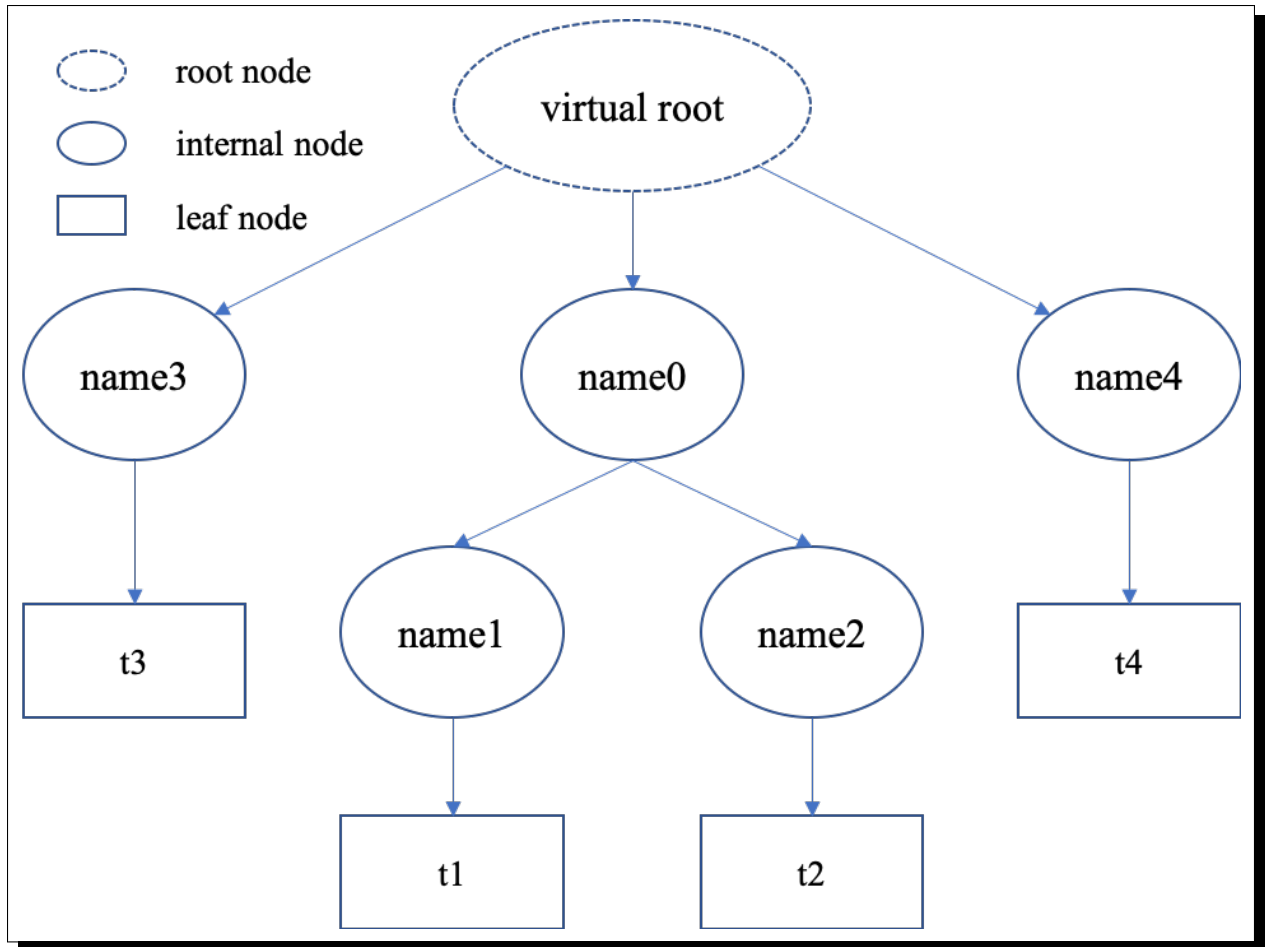
## 1.4 Understand Batch

*Batch* is the internal data structure extensively used in Tianshou. It is designed to store and manipulate hierarchical named tensors. This tutorial aims to help users correctly understand the concept and the behavior of *Batch* so that users can make the best of Tianshou.

The tutorial has three parts. We first explain the concept of hierarchical named tensors, and introduce basic usage of *Batch*, followed by advanced topics of *Batch*.

### 1.4.1 Hierarchical Named Tensors

The structure of a Batch shown by a tree



“Hierarchical named tensors” refers to a set of tensors where their names form a hierarchy. Suppose there are four tensors  $[t1, t2, t3, t4]$  with names  $[name1, name2, name3, name4]$ , where  $name1$  and  $name2$  belong to the same namespace  $name0$ , then the full name of tensor  $t1$  is  $name0.name1$ . That is, the hierarchy lies in the names of tensors.

We can describe the structure of hierarchical named tensors using a tree in the right. There is always a “virtual root” node to represent the whole object; internal nodes are keys (names), and leaf nodes are values (scalars or tensors).

Hierarchical named tensors are needed because we have to deal with the heterogeneity of reinforcement learning problems. The abstraction of RL is very simple, just:

```
state, reward, done = env.step(action)
```

`reward` and `done` are simple, they are mostly scalar values. However, the `state` and `action` vary with environments. For example, `state` can be simply a vector, a tensor, or a camera input combined with sensory input. In the last case, it is natural to store them as hierarchical named tensors. This hierarchy can go beyond `state` and `action`: we can store `state`, `action`, `reward`, and `done` together as hierarchical named tensors.

Note that, storing hierarchical named tensors is as easy as creating nested dictionary objects:

```
{
  'done': done,
  'reward': reward,
  'state': {
    'camera': camera,
```

(continues on next page)



(continued from previous page)

```

        'sensory': sensory
    }
    'action': {
        'direct': direct,
        'point_3d': point_3d,
        'force': force,
    }
}

```

The real problem is how to **manipulate them**, such as adding new transition tuples into replay buffer and dealing with their heterogeneity. `Batch` is designed to easily create, store, and manipulate these hierarchical named tensors.

## 1.4.2 Basic Usages

Here we cover some basic usages of `Batch`, describing what `Batch` contains, how to construct `Batch` objects and how to manipulate them.

### What Does Batch Contain

The content of `Batch` objects can be defined by the following rules.

1. A `Batch` object can be an empty `Batch()`, or have at least one key-value pairs. `Batch()` can be used to reserve keys, too. See [Key Reservations](#) for this advanced usage.
2. The keys are always strings (they are names of corresponding values).
3. The values can be scalars, tensors, or `Batch` objects. The recursive definition makes it possible to form a hierarchy of batches.
4. Tensors are the most important values. In short, tensors are n-dimensional arrays of the same data type. We support two types of tensors: `PyTorch` tensor type `torch.Tensor` and `NumPy` tensor type `np.ndarray`.
5. Scalars are also valid values. A scalar is a single boolean, number, or object. They can be python scalar (`False`, `1`, `2.3`, `None`, `'hello'`) or NumPy scalar (`np.bool_(True)`, `np.int32(1)`, `np.float64(2.3)`). They just shouldn't be mixed up with `Batch`/dict/tensors.

---

**Note:** `Batch` cannot store `dict` objects, because internally `Batch` uses `dict` to store data. During construction, `dict` objects will be automatically converted to `Batch` objects.

The data types of tensors are `bool` and numbers (any size of `int` and `float` as long as they are supported by `NumPy` or `PyTorch`). Besides, `NumPy` supports `ndarray` of objects and we take advantage of this feature to store non-number objects in `Batch`. If one wants to store data that are neither boolean nor numbers (such as strings and sets), they can store the data in `np.ndarray` with the `np.object` data type. This way, `Batch` can store any type of python objects.

---

## Construction of Batch

There are two ways to construct a Batch object: from a dict, or using kwargs. Below are some code snippets.

```
>>> # directly passing a dict object (possibly nested) is ok
>>> data = Batch({'a': 4, 'b': [5, 5], 'c': '2312312'})
>>> # the list will automatically be converted to numpy array
>>> data.b
array([5, 5])
>>> data.b = np.array([3, 4, 5])
>>> print(data)
Batch(
  a: 4,
  b: array([3, 4, 5]),
  c: '2312312',
)
>>> # a list of dict objects (possibly nested) will be automatically stacked
>>> data = Batch([{'a': 0.0, 'b': "hello"}, {'a': 1.0, 'b': "world"}])
>>> print(data)
Batch(
  a: array([0., 1.]),
  b: array(['hello', 'world'], dtype=object),
)

>>> # construct a Batch with keyword arguments
>>> data = Batch(a=[4, 4], b=[5, 5], c=[None, None])
>>> print(data)
Batch(
  a: array([4, 4]),
  b: array([5, 5]),
  c: array([None, None], dtype=object),
)
>>> # combining keyword arguments and batch_dict works fine
>>> data = Batch({'a': [4, 4], 'b': [5, 5]}, c=[None, None]) # the first argument is a_
↳ dict, and 'c' is a keyword argument
>>> print(data)
Batch(
  a: array([4, 4]),
  b: array([5, 5]),
  c: array([None, None], dtype=object),
)
>>> arr = np.zeros((3, 4))
>>> # By default, Batch only keeps the reference to the data, but it also supports_
↳ data copying
>>> data = Batch(arr=arr, copy=True) # data.arr now is a copy of 'arr'
```

## Data Manipulation With Batch

Users can access the internal data by `b.key` or `b[key]`, where `b.key` finds the sub-tree with `key` as the root node. If the result is a sub-tree with non-empty keys, the key-reference can be chained, i.e. `b.key.key1.key2.key3`. When it reaches a leaf node, users get the data (scalars/tensors) stored in that Batch object.

```
>>> data = Batch(a=4, b=[5, 5])
>>> print(data.b)
[5 5]
>>> # obj.key is equivalent to obj["key"]
```

(continues on next page)

(continued from previous page)

```

>>> print(data["a"])
4
>>> # iterating over data items like a dict is supported
>>> for key, value in data.items():
>>>     print(f"{key}: {value}")
a: 4
b: [5, 5]
>>> # obj.keys() and obj.values() work just like dict.keys() and dict.values()
>>> for key in data.keys():
>>>     print(f"{key}")
a
b
>>> # obj.update() behaves like dict.update()
>>> # this is the same as data.c = 1; data.c = 2; data.e = 3;
>>> data.update(c=1, d=2, e=3)
>>> print(data)
Batch(
  a: 4,
  b: array([5, 5]),
  c: 1,
  d: 2,
  e: 3,
)

```

**Note:** If data is a dict object, for `x in data` iterates over keys in the dict. However, it has a different meaning for Batch objects: for `x in data` iterates over `data[0]`, `data[1]`, ..., `data[-1]`. An example is given below.

Batch also partially reproduces the NumPy ndarray APIs. It supports advanced slicing, such as `batch[:, i]` so long as the slice is valid. Broadcast mechanism of NumPy works for Batch, too.

```

>>> # initialize Batch with tensors
>>> data = Batch(a=np.array([[0.0, 2.0], [1.0, 3.0]]), b=[[5, -5], [1, -2]])
>>> # if values have the same length/shape, that length/shape is used for this Batch
>>> # else, check the advanced topic for details
>>> print(len(data))
2
>>> print(data.shape)
[2, 2]
>>> # access the first item of all the stored tensors, while keeping the structure of
↳Batch
>>> print(data[0])
Batch(
  a: array([0., 2.])
  b: array([ 5, -5]),
)
>>> # iterates over ``data[0], data[1], ..., data[-1]``
>>> for sample in data:
>>>     print(sample.a)
[0. 2.]
[1. 3.]

>>> # Advanced slicing works just fine
>>> # Arithmetic operations are passed to each value in the Batch, with broadcast
↳enabled

```

(continues on next page)

(continued from previous page)

```

>>> data[:, 1] += 1
>>> print(data)
Batch(
  a: array([[0., 3.],
            [1., 4.]]) ,
  b: array([[ 5, -4]]) ,
)

>>> # amazingly, you can directly apply np.mean to a Batch object
>>> print(np.mean(data))
Batch(
  a: 1.5,
  b: -0.25,
)

>>> # directly converted to a list is also available
>>> list(data)
[Batch(
  a: array([0., 3.]),
  b: array([ 5, -4]),
),
Batch(
  a: array([1., 4.]),
  b: array([ 1, -1]),
)]

```

Stacking and concatenating multiple Batch instances, or split an instance into multiple batches, they are all easy and intuitive in Tianshou. For now, we stick to the aggregation (stack/concatenate) of homogeneous (same structure) batches. Stack/Concatenation of heterogeneous batches are discussed in [Aggregation of Heterogeneous Batches](#).

```

>>> data_1 = Batch(a=np.array([0.0, 2.0]), b=5)
>>> data_2 = Batch(a=np.array([1.0, 3.0]), b=-5)
>>> data = Batch.stack((data_1, data_2))
>>> print(data)
Batch(
  b: array([ 5, -5]),
  a: array([[0., 2.],
            [1., 3.]]) ,
)

>>> # split supports random shuffling
>>> data_split = list(data.split(1, shuffle=False))
>>> print(list(data.split(1, shuffle=False)))
[Batch(
  b: array([5]),
  a: array([[0., 2.]]) ,
), Batch(
  b: array([-5]),
  a: array([[1., 3.]]) ,
)]

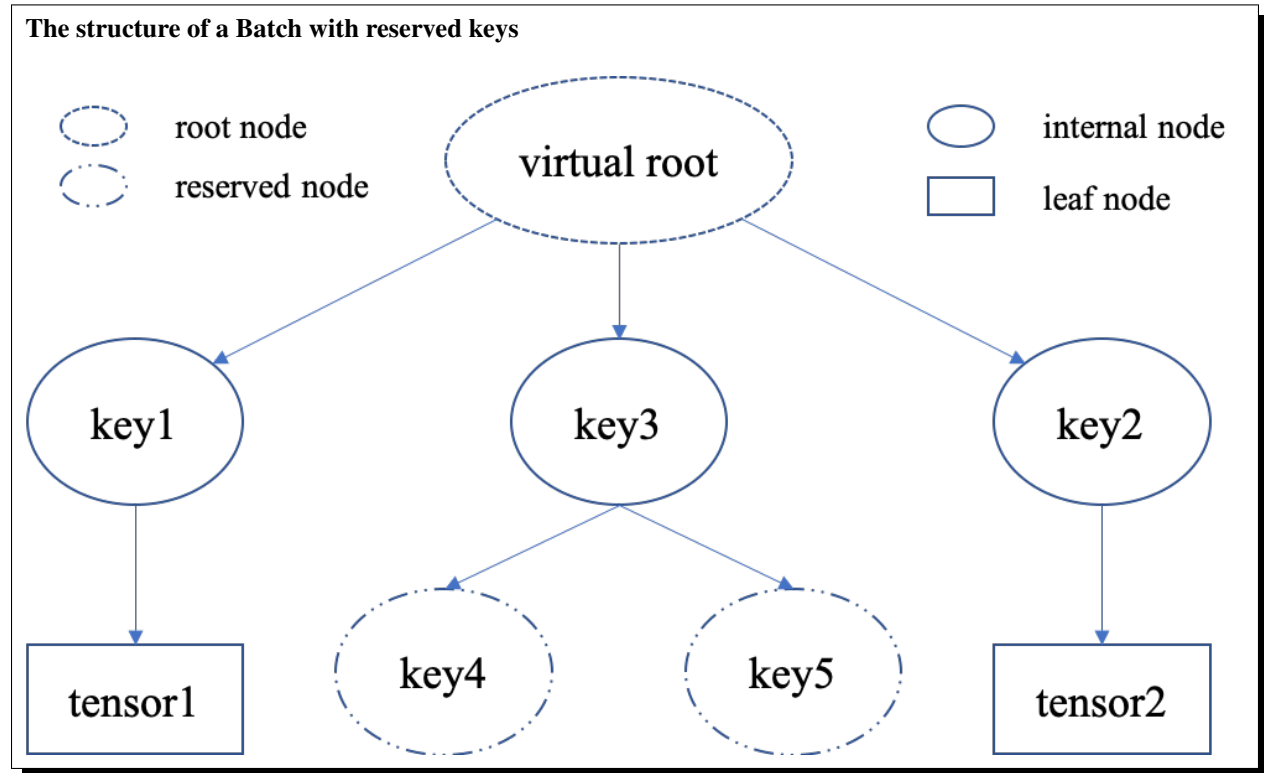
>>> data_cat = Batch.cat(data_split)
>>> print(data_cat)
Batch(
  b: array([ 5, -5]),
  a: array([[0., 2.],
            [1., 3.]]) ,
)

```

### 1.4.3 Advanced Topics

From here on, this tutorial focuses on advanced topics of `Batch`, including key reservation, length/shape, and aggregation of heterogeneous batches.

#### Key Reservations



In many cases, we know in the first place what keys we have, but we do not know the shape of values until we run the environment. To deal with this, Tianshou supports key reservations: **reserve a key and use a placeholder value**.

The usage is easy: just use `Batch()` to be the value of reserved keys.

```

a = Batch(b=Batch()) # 'b' is a reserved key
# this is called hierarchical key reservation
a = Batch(b=Batch(c=Batch()), d=Batch()) # 'c' and 'd' are reserved key
# the structure of this last Batch is shown in the right figure
a = Batch(key1=tensor1, key2=tensor2, key3=Batch(key4=Batch(), key5=Batch()))
  
```

Still, we can use a tree (in the right) to show the structure of `Batch` objects with reserved keys, where reserved keys are special internal nodes that do not have attached leaf nodes.

**Note:** Reserved keys mean that in the future there will eventually be values attached to them. The values can be scalars, tensors, or even **Batch** objects. Understanding this is critical to understand the behavior of `Batch` when dealing with heterogeneous Batches.

The introduction of reserved keys gives rise to the need to check if a key is reserved. Tianshou provides `Batch.is_empty` to achieve this.

```
>>> Batch().is_empty()
True
>>> Batch(a=Batch(), b=Batch(c=Batch())).is_empty()
False
>>> Batch(a=Batch(), b=Batch(c=Batch())).is_empty(recurse=True)
True
>>> Batch(d=1).is_empty()
False
>>> Batch(a=np.float64(1.0)).is_empty()
False
```

The `Batch.is_empty` function has an option to decide whether to identify direct emptiness (just a `Batch()`) or to identify recursive emptiness (a `Batch` object without any scalar/tensor leaf nodes).

---

**Note:** Do not get confused with `Batch.is_empty` and `Batch.empty`. `Batch.empty` and its in-place variant `Batch.empty_` are used to set some values to zeros or None. Check the API documentation for further details.

---

## Length and Shape

The most common usage of `Batch` is to store a Batch of data. The term “Batch” comes from the deep learning community to denote a mini-batch of sampled data from the whole dataset. In this regard, “Batch” typically means a collection of tensors whose first dimensions are the same. Then the length of a `Batch` object is simply the batch-size.

If all the leaf nodes in a `Batch` object are tensors, but they have different lengths, they can be readily stored in `Batch`. However, for `Batch` of this kind, the `len(obj)` seems a bit ambiguous. Currently, Tianshou returns the length of the shortest tensor, but we strongly recommend that users do not use the `len(obj)` operator on `Batch` objects with tensors of different lengths.

```
>>> data = Batch(a=[5., 4.], b=np.zeros((2, 3, 4)))
>>> data.shape
[2]
>>> len(data)
2
>>> data[0].shape
[]
>>> len(data[0])
TypeError: Object of type 'Batch' has no len()
```

---

**Note:** Following the convention of scientific computation, scalars have no length. If there is any scalar leaf node in a `Batch` object, an exception will occur when users call `len(obj)`.

Besides, values of reserved keys are undetermined, so they have no length, neither. Or, to be specific, values of reserved keys have lengths of **any**. When there is a mix of tensors and reserved keys, the latter will be ignored in `len(obj)` and the minimum length of tensors is returned. When there is not any tensor in the `Batch` object, Tianshou raises an exception, too.

---

The `obj.shape` attribute of `Batch` behaves somewhat similar to `len(obj)`:

1. If all the leaf nodes in a `Batch` object are tensors with the same shape, that shape is returned.
2. If all the leaf nodes in a `Batch` object are tensors but they have different shapes, the minimum length of each dimension is returned.
3. If there is any scalar value in a `Batch` object, `obj.shape` returns `[]`.

4. The shape of reserved keys is undetermined, too. We treat their shape as `[]`.

## Aggregation of Heterogeneous Batches

In this section, we talk about aggregation operators (stack/concatenate) on heterogeneous `Batch` objects. The following picture will give you an intuitive understanding of this behavior. It shows two examples of aggregation operators with heterogeneous `Batch`. The shapes of tensors are annotated in the leaf nodes.



We only consider the heterogeneity in the structure of `Batch` objects. The aggregation operators are eventually done by NumPy/PyTorch operators (`np.stack`, `np.concatenate`, `torch.stack`, `torch.cat`). Heterogeneity in values can fail these operators (such as stacking `np.ndarray` with `torch.Tensor`, or stacking tensors with different shapes) and an exception will be raised.

The behavior is natural: for keys that are not shared across all batches, batches that do not have these keys will be padded by zeros (or `None` if the data type is `np.object`). It can be written in the following scripts:

```
>>> # examples of stack: a is missing key `b`, and b is missing key `a`
>>> a = Batch(a=np.zeros([4, 4]), common=Batch(c=np.zeros([4, 5])))
>>> b = Batch(b=np.zeros([4, 6]), common=Batch(c=np.zeros([4, 5])))
>>> c = Batch.stack([a, b])
>>> c.a.shape
(2, 4, 4)
>>> c.b.shape
(2, 4, 6)
>>> c.common.c.shape
(2, 4, 5)
>>> # None or 0 is padded with appropriate shape
>>> data_1 = Batch(a=np.array([0.0, 2.0]))
>>> data_2 = Batch(a=np.array([1.0, 3.0]), b='done')
>>> data = Batch.stack((data_1, data_2))
>>> print(data)
Batch(
  a: array([[0., 2.],
            [1., 3.]])
  b: array([None, 'done'], dtype=object),
```

(continues on next page)

(continued from previous page)

```

)
>>> # examples of cat: a is missing key `b`, and b is missing key `a`
>>> a = Batch(a=np.zeros([3, 4]), common=Batch(c=np.zeros([3, 5])))
>>> b = Batch(b=np.zeros([4, 3]), common=Batch(c=np.zeros([4, 5])))
>>> c = Batch.cat([a, b])
>>> c.a.shape
(7, 4)
>>> c.b.shape
(7, 3)
>>> c.common.c.shape
(7, 5)

```

However, there are some cases when batches are too heterogeneous that they cannot be aggregated:

```

>>> a = Batch(a=np.zeros([4, 4]))
>>> b = Batch(a=Batch(b=Batch()))
>>> # this will raise an exception
>>> c = Batch.stack([a, b])

```

Then how to determine if batches can be aggregated? Let's rethink the purpose of reserved keys. What is the advantage of `a1=Batch(b=Batch())` over `a2=Batch()`? The only difference is that `a1.b` returns `Batch()` but `a2.b` raises an exception. That's to say, **we reserve keys for attribute reference**.

We say a key chain `k=[key1, key2, ..., keyn]` applies to `b` if the expression `b.key1.key2.{...}.keyn` is valid, and the result is `b[k]`.

For a set of `Batch` objects denoted as  $S$ , they can be aggregated if there exists a `Batch` object `b` satisfying the following rules:

1. Key chain applicability: For any object `bi` in  $S$ , and any key chain `k`, if `bi[k]` is valid, then `b[k]` is valid.
2. Type consistency: If `bi[k]` is not `Batch()` (the last key in the key chain is not a reserved key), then the type of `b[k]` should be the same as `bi[k]` (both should be scalar/tensor/non-empty `Batch` values).

The `Batch` object `b` satisfying these rules with the minimum number of keys determines the structure of aggregating  $S$ . The values are relatively easy to define: for any key chain `k` that applies to `b`, `b[k]` is the stack/concatenation of `[bi[k] for bi in S]` (if `k` does not apply to `bi`, the appropriate size of zeros or `None` are filled automatically). If `bi[k]` are all `Batch()`, then the aggregation result is also an empty `Batch()`.

## Miscellaneous Notes

1. `Batch` is serializable and therefore `Pickle` compatible. `Batch` objects can be saved to disk and later restored by the python `pickle` module. This pickle compatibility is especially important for distributed sampling from environments.

```

>>> data = Batch(a=np.zeros((3, 4)))
>>> data.to_torch(dtype=torch.float32, device='cpu')
>>> print(data.a)
tensor([[0., 0., 0., 0.],
        [0., 0., 0., 0.],
        [0., 0., 0., 0.]])
>>> # data.to_numpy is also available
>>> data.to_numpy()

```

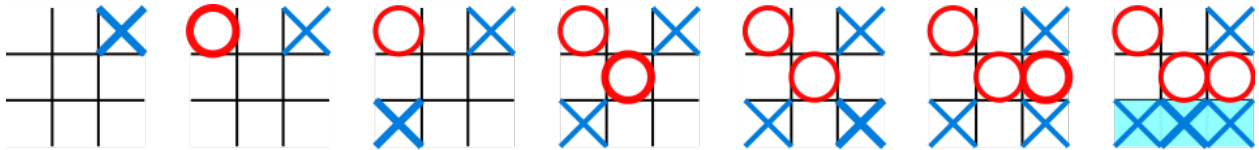
2. It is often the case that the observations returned from the environment are all NumPy ndarray but the policy requires `torch.Tensor` for prediction and learning. In this regard, Tianshou provides helper functions to convert the stored data in-place into Numpy arrays or Torch tensors.



3. `obj.stack_([a, b])` is the same as `Batch.stack([obj, a, b])`, and `obj.cat_([a, b])` is the same as `Batch.cat([obj, a, b])`. Considering the frequent requirement of concatenating two Batch objects, Tianshou also supports `obj.cat_(a)` to be an alias of `obj.cat_([a])`.
4. `Batch.cat` and `Batch.cat_` does not support `axis` argument as `np.concatenate` and `torch.cat` currently.
5. `Batch.stack` and `Batch.stack_` support the `axis` argument so that one can stack batches besides the first dimension. But be cautious, if there are keys that are not shared across all batches, `stack` with `axis != 0` is undefined, and will cause an exception currently.

## 1.5 Multi-Agent RL

In this section, we describe how to use Tianshou to implement multi-agent reinforcement learning. Specifically, we will design an algorithm to learn how to play **Tic Tac Toe** (see the image below) against a random opponent.



### 1.5.1 Tic-Tac-Toe Environment

The scripts are located at `test/pettingzoo/`. We have implemented `PettingZooEnv` which can wrap any `[PettingZoo]`(<https://www.pettingzoo.ml/>) environment. `PettingZoo` offers a 3x3 Tic-Tac-Toe environment, let's first explore it.

```
>>> from tianshou.env import PettingZooEnv          # wrapper for PettingZoo environments
>>> from pettingzoo.classic import tictactoe_v3     # the Tic-Tac-Toe environment to be_
↳wrapped
>>> # This board has 3 rows and 3 cols (9 places in total)
>>> # Players place 'x' and 'o' in turn on the board
>>> # The player who first gets 3 consecutive 'x's or 'o's wins
>>>
>>> env = PettingZooEnv(tictactoe_v3.env())
>>> obs = env.reset()
>>> env.render()                                     # render the empty board
board (step 0):
  | | |
- | - | -
_|_|_|
  | | |
- | - | -
_|_|_|
  | | |
- | - | -
  | | |
>>> print(obs)                                       # let's see the shape of the_
↳observation
{'agent_id': 'player_1', 'obs': array([[0, 0],
      [0, 0],
      [0, 0],
      [0, 0],
```

(continues on next page)

(continued from previous page)

```

        [0, 0],
        [0, 0]],
        [[0, 0],
         [0, 0],
         [0, 0]]], dtype=int8), 'mask': [True, True, True, True, True, True, True,
↪True, True]}

```

The observation variable `obs` returned from the environment is a dict, with three keys `agent_id`, `obs`, `mask`. This is a general structure in multi-agent RL where agents take turns. The meaning of these keys are:

- `agent_id`: the id of the current acting agent. In our Tic-Tac-Toe case, the `agent_id` can be `player_1` or `player_2`.
- `obs`: the actual observation of the environment. In the Tic-Tac-Toe game above, the observation variable `obs` is a `np.ndarray` with the shape of (3, 3, 2). For `player_1`, the first 3x3 plane represents the placement of Xs, and the second plane shows the placement of Os. The possible values for each cell are 0 or 1; in the first plane, 1 indicates that an X has been placed in that cell, and 0 indicates that X is not in that cell. Similarly, in the second plane, 1 indicates that an O has been placed in that cell, while 0 indicates that an O has not been placed. For `player_2`, the observation is the same, but Xs and Os swap positions, so Os are encoded in plane 1 and Xs in plane 2.
- `mask`: the action mask in the current timestep. In board games or card games, the legal action set varies with time. The mask is a boolean array. For Tic-Tac-Toe, index `i` means the place of `i/N`th row and `i%N`th column. If `mask[i] == True`, the player can place an x or o at that position. Now the board is empty, so the mask is all the true, contains all the positions on the board.

**Note:** There is no special formulation of `mask` either in discrete action space or in continuous action space. You can also use some action spaces like `gym.spaces.Discrete` or `gym.spaces.Box` to represent the available action space. Currently, we use a boolean array.

Let's play two steps to have an intuitive understanding of the environment.

```

>>> import numpy as np
>>> action = 0                                     # action is either an integer, or an
↪np.ndarray with one element
>>> obs, reward, done, info = env.step(action)      # the env.step follows the api of
↪OpenAI Gym
>>> print(obs)                                     # notice the change in the observation
{'agent_id': 'player_2', 'obs': array([[0, 1],
        [0, 0],
        [0, 0]],
        [[0, 0],
         [0, 0],
         [0, 0]],
        [[0, 0],
         [0, 0],
         [0, 0]]], dtype=int8), 'mask': [False, True, True, True, True, True, True,
↪True, True]}
>>> # reward has two items, one for each player: 1 for win, -1 for lose, and 0
↪otherwise
>>> print(reward)
[0. 0.]

```

(continues on next page)

(continued from previous page)

```
>>> print(done)                # done indicates whether the game is_
↪over
False
>>> # info is always an empty dict in Tic-Tac-Toe, but may contain some useful_
↪information in environments other than Tic-Tac-Toe.
>>> print(info)
{}
```

One worth-noting case is that the game is over when there is only one empty position, rather than when there is no position. This is because the player just has one choice (literally no choice) in this game.

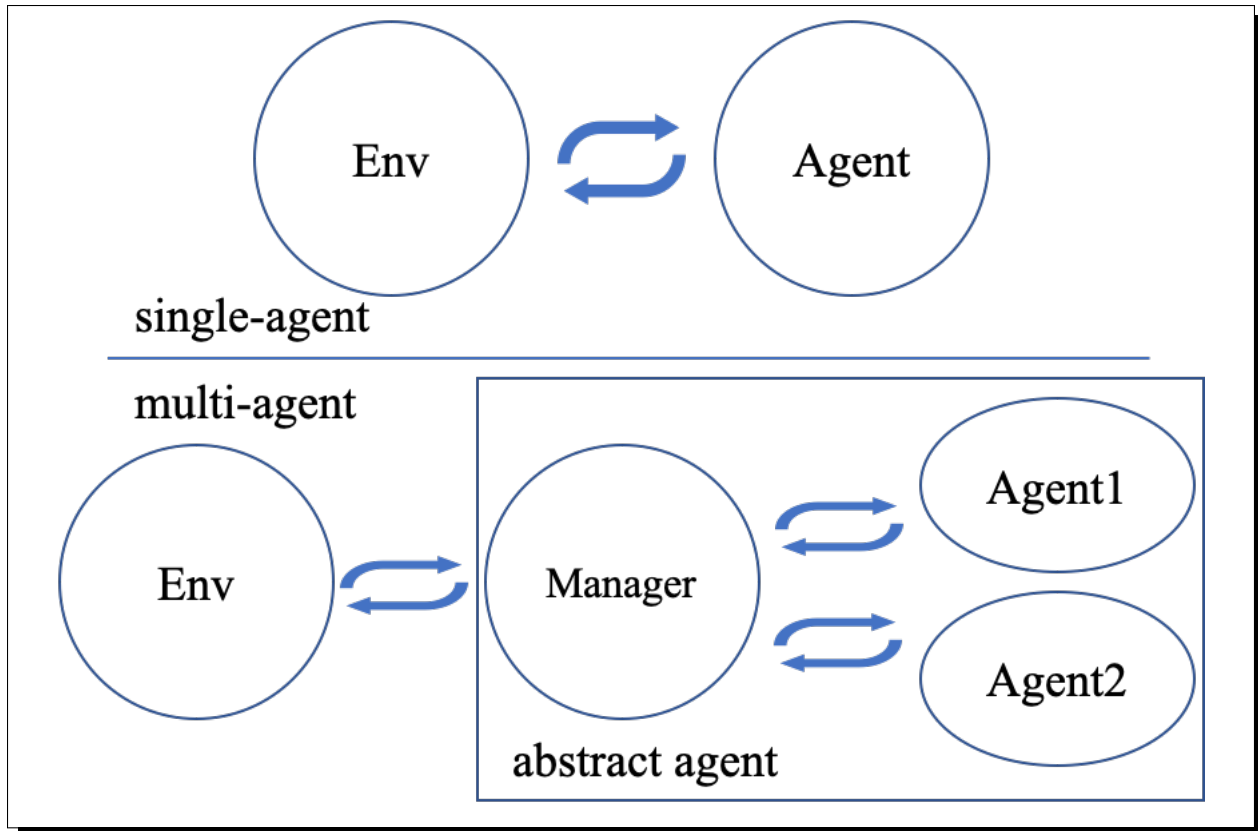
```
>>> # omitted actions: 3, 1, 4
>>> obs, reward, done, info = env.step(2) # player_1 wins
>>> print((reward, done))
([1, -1], True)
>>> env.render()

| | |
X | O | -
|_|_|_|
| | |
X | O | -
|_|_|_|
| | |
X | - | -
|_|_|_|
```

After being familiar with the environment, let's try to play with random agents first!

## 1.5.2 Two Random Agents

**The relationship between MultiAgentPolicyManager (Manager) and BasePolicy (Agent)**



Tianshou already provides some builtin classes for multi-agent learning. You can check out the API documentation for details. Here we use *RandomPolicy* and *MultiAgentPolicyManager*. The figure on the right gives an intuitive explanation.

```
>>> from tianshou.data import Collector
>>> from tianshou.env import DummyVectorEnv
>>> from tianshou.policy import RandomPolicy, MultiAgentPolicyManager
>>>
>>> # agents should be wrapped into one policy,
>>> # which is responsible for calling the acting agent correctly
>>> # here we use two random agents
>>> policy = MultiAgentPolicyManager([RandomPolicy(), RandomPolicy()], env)
>>>
>>> # need to vectorize the environment for the collector
>>> env = DummyVectorEnv([lambda: env])
>>>
>>> # use collectors to collect a episode of trajectories
>>> # the reward is a vector, so we need a scalar metric to monitor the training
>>> collector = Collector(policy, env)
>>>
>>> # you will see a long trajectory showing the board status at each timestep
>>> result = collector.collect(n_episode=1, render=.1)
(only show the last 3 steps)
```

X	X	-
X	O	-

(continues on next page)

(continued from previous page)

O	-	-
X	X	-
X	O	-
O	-	O
X	X	X
X	O	-
O	-	O

Random agents perform badly. In the above game, although agent 2 wins finally, it is clear that a smart agent 1 would place an x at row 4 col 4 to win directly.

### 1.5.3 Train an MARL Agent

So let's start to train our Tic-Tac-Toe agent! First, import some required modules.

```
import argparse
import os
from copy import deepcopy
from typing import Optional, Tuple

import gym
import numpy as np
import torch
from pettingzoo.classic import tictactoe_v3
from torch.utils.tensorboard import SummaryWriter

from tianshou.data import Collector, VectorReplayBuffer
from tianshou.env import DummyVectorEnv
from tianshou.env.pettingzoo_env import PettingZooEnv
from tianshou.policy import (
    BasePolicy,
    DQNPolicy,
    MultiAgentPolicyManager,
    RandomPolicy,
)
from tianshou.trainer import offpolicy_trainer
from tianshou.utils import TensorboardLogger
from tianshou.utils.net.common import Net
```

The explanation of each Tianshou class/function will be deferred to their first usages. Here we define some arguments and hyperparameters of the experiment. The meaning of arguments is clear by just looking at their names.

```

def get_parser() -> argparse.ArgumentParser:
    parser = argparse.ArgumentParser()
    parser.add_argument('--seed', type=int, default=1626)
    parser.add_argument('--eps-test', type=float, default=0.05)
    parser.add_argument('--eps-train', type=float, default=0.1)
    parser.add_argument('--buffer-size', type=int, default=20000)
    parser.add_argument('--lr', type=float, default=1e-4)
    parser.add_argument(
        '--gamma', type=float, default=0.9, help='a smaller gamma favors earlier win'
    )
    parser.add_argument('--n-step', type=int, default=3)
    parser.add_argument('--target-update-freq', type=int, default=320)
    parser.add_argument('--epoch', type=int, default=50)
    parser.add_argument('--step-per-epoch', type=int, default=1000)
    parser.add_argument('--step-per-collect', type=int, default=10)
    parser.add_argument('--update-per-step', type=float, default=0.1)
    parser.add_argument('--batch-size', type=int, default=64)
    parser.add_argument(
        '--hidden-sizes', type=int, nargs='*', default=[128, 128, 128, 128]
    )
    parser.add_argument('--training-num', type=int, default=10)
    parser.add_argument('--test-num', type=int, default=10)
    parser.add_argument('--logdir', type=str, default='log')
    parser.add_argument('--render', type=float, default=0.1)
    parser.add_argument(
        '--win-rate',
        type=float,
        default=0.6,
        help='the expected winning rate: Optimal policy can get 0.7'
    )
    parser.add_argument(
        '--watch',
        default=False,
        action='store_true',
        help='no training, '
        'watch the play of pre-trained models'
    )
    parser.add_argument(
        '--agent-id',
        type=int,
        default=2,
        help='the learned agent plays as the'
        ' agent_id-th player. Choices are 1 and 2.'
    )
    parser.add_argument(
        '--resume-path',
        type=str,
        default='',
        help='the path of agent pth file '
        'for resuming from a pre-trained agent'
    )
    parser.add_argument(
        '--opponent-path',
        type=str,
        default='',
        help='the path of opponent agent pth file '
        'for resuming from a pre-trained agent'
    )

```

(continues on next page)

(continued from previous page)

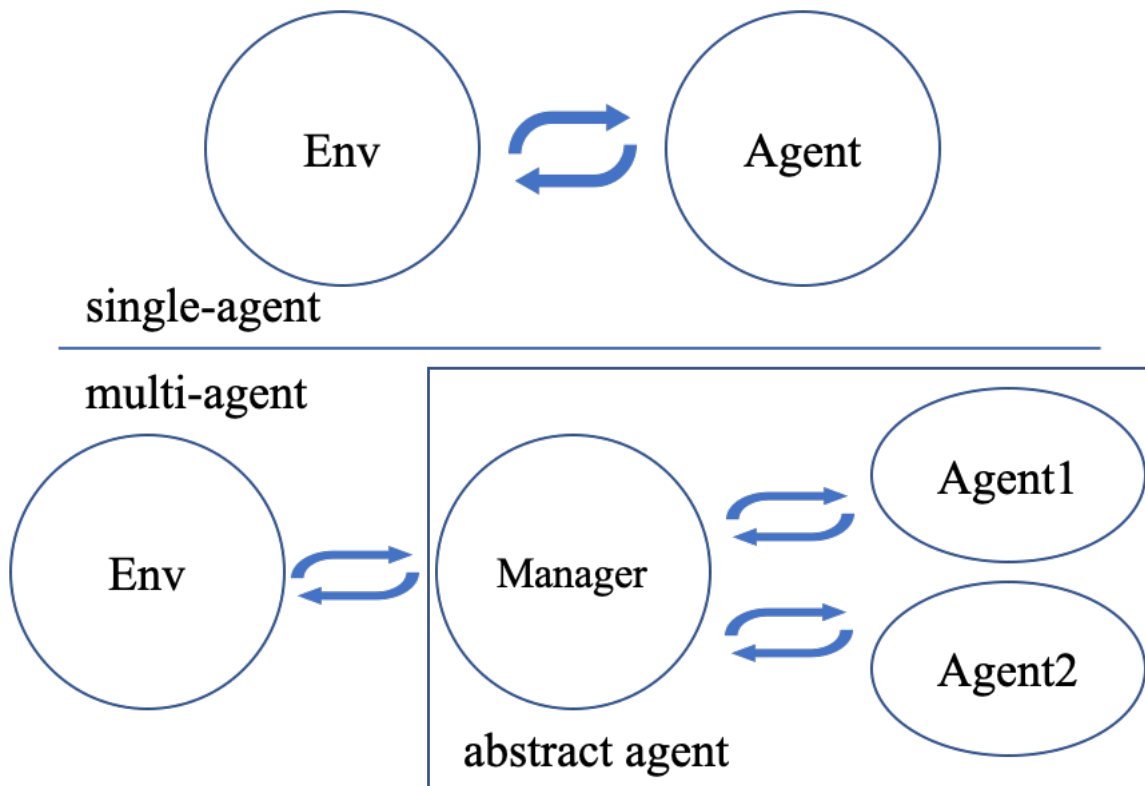
```

)
parser.add_argument(
    '--device', type=str, default='cuda' if torch.cuda.is_available() else 'cpu'
)
return parser

def get_args() -> argparse.Namespace:
    parser = get_parser()
    return parser.parse_known_args()[0]

```

### The relationship between MultiAgentPolicyManager (Manager) and BasePolicy (Agent)



The following `get_agents` function returns agents and their optimizers from either constructing a new policy, or loading from disk, or using the pass-in arguments. For the models:

- The action model we use is an instance of `Net`, essentially a multi-layer perceptron with the ReLU activation function;
- The network model is passed to a `DQNPolicy`, where actions are selected according to both the action mask and their Q-values;
- The opponent can be either a random agent `RandomPolicy` that randomly chooses an action from legal actions, or it can be a pre-trained `DQNPolicy` allowing learned agents to play with themselves.

Both agents are passed to `MultiAgentPolicyManager`, which is responsible to call the correct agent according to the `agent_id` in the observation. `MultiAgentPolicyManager` also dispatches data to each agent according to `agent_id`, so that each agent seems to play with a virtual single-agent environment.

Here it is:

```
def get_agents(
    args: argparse.Namespace = get_args(),
    agent_learn: Optional[BasePolicy] = None,
    agent_opponent: Optional[BasePolicy] = None,
    optim: Optional[torch.optim.Optimizer] = None,
) -> Tuple[BasePolicy, torch.optim.Optimizer, list]:
    env = get_env()
    observation_space = env.observation_space['observation'] if isinstance(
        env.observation_space, gym.spaces.Dict
    ) else env.observation_space
    args.state_shape = observation_space.shape or observation_space.n
    args.action_shape = env.action_space.shape or env.action_space.n
    if agent_learn is None:
        # model
        net = Net(
            args.state_shape,
            args.action_shape,
            hidden_sizes=args.hidden_sizes,
            device=args.device
        ).to(args.device)
        if optim is None:
            optim = torch.optim.Adam(net.parameters(), lr=args.lr)
        agent_learn = DQNPolicy(
            net,
            optim,
            args.gamma,
            args.n_step,
            target_update_freq=args.target_update_freq
        )
        if args.resume_path:
            agent_learn.load_state_dict(torch.load(args.resume_path))

    if agent_opponent is None:
        if args.opponent_path:
            agent_opponent = deepcopy(agent_learn)
            agent_opponent.load_state_dict(torch.load(args.opponent_path))
        else:
            agent_opponent = RandomPolicy()

    if args.agent_id == 1:
        agents = [agent_learn, agent_opponent]
    else:
        agents = [agent_opponent, agent_learn]
    policy = MultiAgentPolicyManager(agents, env)
    return policy, optim, env.agents
```

With the above preparation, we are close to the first learned agent. The following code is almost the same as the code in the DQN tutorial.

```
def get_env():
    return PettingZooEnv(tictactoe_v3.env())

def train_agent(
    args: argparse.Namespace = get_args(),
    agent_learn: Optional[BasePolicy] = None,
```

(continues on next page)



(continued from previous page)

```

agent_opponent: Optional[BasePolicy] = None,
optim: Optional[torch.optim.Optimizer] = None,
) -> Tuple[dict, BasePolicy]:

    # ===== environment setup =====
    train_envs = DummyVectorEnv([get_env for _ in range(args.training_num)])
    test_envs = DummyVectorEnv([get_env for _ in range(args.test_num)])
    # seed
    np.random.seed(args.seed)
    torch.manual_seed(args.seed)
    train_envs.seed(args.seed)
    test_envs.seed(args.seed)

    # ===== agent setup =====
    policy, optim, agents = get_agents(
        args, agent_learn=agent_learn, agent_opponent=agent_opponent, optim=optim
    )

    # ===== collector setup =====
    train_collector = Collector(
        policy,
        train_envs,
        VectorReplayBuffer(args.buffer_size, len(train_envs)),
        exploration_noise=True
    )
    test_collector = Collector(policy, test_envs, exploration_noise=True)
    # policy.set_eps(1)
    train_collector.collect(n_step=args.batch_size * args.training_num)

    # ===== tensorboard logging setup =====
    log_path = os.path.join(args.logdir, 'tic_tac_toe', 'dqn')
    writer = SummaryWriter(log_path)
    writer.add_text("args", str(args))
    logger = TensorboardLogger(writer)

    # ===== callback functions used during training =====
    def save_best_fn(policy):
        if hasattr(args, 'model_save_path'):
            model_save_path = args.model_save_path
        else:
            model_save_path = os.path.join(
                args.logdir, 'tic_tac_toe', 'dqn', 'policy.pth'
            )
        torch.save(
            policy.policies[agents[args.agent_id - 1]].state_dict(), model_save_path
        )

    def stop_fn(mean_rewards):
        return mean_rewards >= args.win_rate

    def train_fn(epoch, env_step):
        policy.policies[agents[args.agent_id - 1]].set_eps(args.eps_train)

    def test_fn(epoch, env_step):
        policy.policies[agents[args.agent_id - 1]].set_eps(args.eps_test)

    def reward_metric(rews):

```

(continues on next page)

(continued from previous page)

```

    return rews[:, args.agent_id - 1]

# trainer
result = offpolicy_trainer(
    policy,
    train_collector,
    test_collector,
    args.epoch,
    args.step_per_epoch,
    args.step_per_collect,
    args.test_num,
    args.batch_size,
    train_fn=train_fn,
    test_fn=test_fn,
    stop_fn=stop_fn,
    save_best_fn=save_best_fn,
    update_per_step=args.update_per_step,
    logger=logger,
    test_in_train=False,
    reward_metric=reward_metric
)

return result, policy.policies[agents[args.agent_id - 1]]

# ===== a test function that tests a pre-trained agent =====
def watch(
    args: argparse.Namespace = get_args(),
    agent_learn: Optional[BasePolicy] = None,
    agent_opponent: Optional[BasePolicy] = None,
) -> None:
    env = get_env()
    policy, optim, agents = get_agents(
        args, agent_learn=agent_learn, agent_opponent=agent_opponent
    )
    policy.eval()
    policy.policies[agents[args.agent_id - 1]].set_eps(args.eps_test)
    collector = Collector(policy, env, exploration_noise=True)
    result = collector.collect(n_episode=1, render=args.render)
    rews, lens = result["rews"], result["lens"]
    print(f"Final reward: {rews[:, args.agent_id - 1].mean()}, length: {lens.mean()}")

# train the agent and watch its performance in a match!
args = get_args()
result, agent = train_agent(args)
watch(args, agent)

```

That's it. By executing the code, you will see a progress bar indicating the progress of training. After about less than 1 minute, the agent has finished training, and you can see how it plays against the random agent. Here is an example:

-	-	-
-	-	X

(continues on next page)

(continued from previous page)

-	-	-
-	-	-
-	O	X
-	-	-
-	-	-
X	O	X
-	-	-
-	O	-
X	O	X
-	-	-
-	O	-
X	O	X
-	X	-
O	O	-
X	O	X
-	X	-
O	O	X
X	O	X
-	X	-

(continues on next page)

(continued from previous page)

O	O	X
X	O	X
-	X	O

Final reward: 1.0, length: 8.0

Notice that, our learned agent plays the role of agent 2, placing o on the board. The agent performs pretty well against the random opponent! It learns the rule of the game by trial and error, and learns that four consecutive o means winning, so it does!

The above code can be executed in a python shell or can be saved as a script file (we have saved it in `test/pettingzoo/test_tic_tac_toe.py`). In the latter case, you can train an agent by

```
$ python test_tic_tac_toe.py
```

By default, the trained agent is stored in `log/tic_tac_toe/dqn/policy.pth`. You can also make the trained agent play against itself, by

```
$ python test_tic_tac_toe.py --watch --resume-path log/tic_tac_toe/dqn/policy.pth --
  ↳opponent-path log/tic_tac_toe/dqn/policy.pth
```

Here is our output:

-	-	-
-	X	-
-	-	-
-	O	-
-	X	-
-	-	-
X	O	-
-	X	-
-	-	-
X	O	-

(continues on next page)

(continued from previous page)

-	X	-
-	-	O
X	O	-
-	X	-
-	X	O
X	O	O
-	X	-
-	X	O
X	O	O
-	X	-
X	X	O
X	O	O
-	X	O
X	X	O

Final reward: 1.0, length: 8.0

Well, although the learned agent plays well against the random agent, it is far away from intelligence.

Next, maybe you can try to build more intelligent agents by letting the agent learn from self-play, just like AlphaZero!

In this tutorial, we show an example of how to use Tianshou for multi-agent RL. Tianshou is a flexible and easy to use RL library. Make the best of Tianshou by yourself!

## 1.6 Logging Experiments

Tianshou comes with multiple experiment tracking and logging solutions to manage and reproduce your experiments. The dashboard loggers currently available are:

- *TensorboardLogger*
- *WandbLogger*
- *LazyLogger*

### 1.6.1 TensorboardLogger

Tensorboard tracks your experiment metrics in a local dashboard. Here is how you can use TensorboardLogger in your experiment:

```
from torch.utils.tensorboard import SummaryWriter
from tianshou.utils import TensorboardLogger

log_path = os.path.join(args.logdir, args.task, "dqn")
writer = SummaryWriter(log_path)
writer.add_text("args", str(args))
logger = TensorboardLogger(writer)
result = trainer(..., logger=logger)
```

### 1.6.2 WandbLogger

*WandbLogger* can be used to visualize your experiments in a hosted [W&B dashboard](#). It can be installed via `pip install wandb`. You can also save your checkpoints in the cloud and restore your runs from those checkpoints. Here is how you can enable WandbLogger:

```
from tianshou.utils import WandbLogger
from torch.utils.tensorboard import SummaryWriter

logger = WandbLogger(...)
writer = SummaryWriter(log_path)
writer.add_text("args", str(args))
logger.load(writer)
result = trainer(..., logger=logger)
```

Please refer to [WandbLogger](#) documentation for advanced configuration.

For logging checkpoints on any device, you need to define a `save_checkpoint_fn` which saves the experiment checkpoint and returns the path of the saved checkpoint:

```
def save_checkpoint_fn(epoch, env_step, gradient_step):
    ckpt_path = ...
    # save model
    return ckpt_path
```

Then, use this function with *WandbLogger* to automatically version your experiment checkpoints after every `save_interval` step.

For resuming runs from checkpoint artifacts on any device, pass the W&B `run_id` of the run that you want to continue in *WandbLogger*. It will then download the latest version of the checkpoint and resume your runs from the checkpoint.

The example scripts are under `test_psrl.py` and `atari_dqn.py`.

### 1.6.3 LazyLogger

This is a place-holder logger that does nothing.

## 1.7 Benchmark

### 1.7.1 Mujoco Benchmark

Tianshou's Mujoco benchmark contains state-of-the-art results.

Every experiment is conducted under 10 random seeds for 1-10M steps. Please refer to <https://github.com/thu-ml/tianshou/tree/master/examples/mujoco> for source code and detailed results.

The table below compares the performance of Tianshou against published results on OpenAI Gym MuJoCo benchmarks. We use max average return in 1M timesteps as the reward metric. ~ means the result is approximated from the plots because quantitative results are not provided. - means results are not provided. The best-performing baseline on each task is highlighted in boldface. Referenced baselines include [TD3 paper](#), [SAC paper](#), [PPO paper](#), [ACKTR paper](#), [OpenAI Baselines](#) and [Spinning Up](#).

Task		Ant	HalfCheetah	Hopper	Walker2d	Swimmer	Humanoid	Reacher	IPendulum	IDPendulum
DDPG	Tianshou	990.4	<b>11718.7</b>	<b>2197.0</b>	1400.6	<b>144.1</b>	<b>177.3</b>	<b>-3.3</b>	<b>1000.0</b>	8364.3
	TD3 Paper	<b>1005.3</b>	3305.6	<b>2020.5</b>	1843.6	/	/	-6.5	<b>1000.0</b>	<b>9355.5</b>
	TD3 Paper (Our)	888.8	8577.3	1860.0	<b>3098.1</b>	/	/	-4.0	<b>1000.0</b>	8370.0
	Spinning Up	~840	~11000	~1800	~1950	~137	/	/	/	/
TD3	Tianshou	<b>5116.4</b>	<b>10201.2</b>	3472.2	3982.4	<b>104.2</b>	<b>5189.5</b>	<b>-2.7</b>	<b>1000.0</b>	<b>9349.2</b>
	TD3 Paper	4372.4	9637.0	<b>3564.1</b>	<b>4682.8</b>	/	/	-3.6	<b>1000.0</b>	9337.5
	Spinning Up	~3800	~9750	~2860	~4000	~78	/	/	/	/
SAC	Tianshou	<b>5850.2</b>	<b>12138.8</b>	<b>3542.2</b>	<b>5007.0</b>	<b>44.4</b>	<b>5488.5</b>	<b>-2.6</b>	<b>1000.0</b>	<b>9359.5</b>
	SAC Paper	~3720	~10400	~3370	~3740	/	~5200	/	/	/
	TD3 Paper	655.4	2347.2	2996.7	1283.7	/	/	-4.4	<b>1000.0</b>	8487.2
	Spinning Up	~3980	~11520	~3150	~4250	~41.7	/	/	/	/
A2C	Tianshou	<b>3485.4</b>	<b>1829.9</b>	<b>1253.2</b>	<b>1091.6</b>	<b>36.6</b>	<b>1726.0</b>	<b>-6.7</b>	<b>1000.0</b>	<b>9257.7</b>
	PPO Paper	/	~1000	~900	~850	~31	/	~-24	~ <b>1000</b>	~7100
	PPO Paper (TR)	/	~930	~1220	~700	~ <b>36</b>	/	~-27	~ <b>1000</b>	~8100
PPO	Tianshou	<b>3258.4</b>	<b>5783.9</b>	<b>2609.3</b>	3588.5	66.7	<b>787.1</b>	<b>-4.1</b>	<b>1000.0</b>	<b>9231.3</b>
	PPO Paper	/	~1800	~2330	~3460	~108	/	~-7	~ <b>1000</b>	~8000
	TD3 Paper	1083.2	1795.4	2164.7	3317.7	/	/	-6.2	<b>1000.0</b>	8977.9
	OpenAI Baselines	/	~1700	~2400	~3510	~111	/	~-6	~940	~7350
	Spinning Up	~650	~1670	~1850	~1230	~ <b>120</b>	/	/	/	/
TRPO	Tianshou	<b>2866.7</b>	<b>4471.2</b>	2046.0	<b>3826.7</b>	40.9	<b>810.1</b>	<b>-5.1</b>	<b>1000.0</b>	<b>8435.2</b>
	ACKTR paper	~0	~400	~1400	~550	~40	/	-8	~ <b>1000</b>	~800
	PPO Paper	/	~0	~2100	~1100	~ <b>121</b>	/	~-115	~ <b>1000</b>	~200
	TD3 paper	-75.9	-15.6	<b>2471.3</b>	2321.5	/	/	-111.4	985.4	205.9
	OpenAI Baselines	/	~1350	~ <b>2200</b>	~2350	~95	/	~-5	~910	~7000
	Spinning Up (TF)	~150	~850	~1200	~600	~85	/	/	/	/

Runtime averaged on 8 MuJoCo benchmark tasks is listed below. All results are obtained using a single Nvidia TITAN X GPU and up to 48 CPU cores (at most one CPU core for each thread).



Algorithm	# of Envs	1M timesteps	Collecting (%)	Updating (%)	Evaluating (%)	Others (%)
DDPG	1	2.9h	12.0	80.2	2.4	5.4
TD3	1	3.3h	11.4	81.7	1.7	5.2
SAC	1	5.2h	10.9	83.8	1.8	3.5
REIN-FORCE	64	4min	84.9	1.8	12.5	0.8
A2C	16	7min	62.5	28.0	6.6	2.9
PPO	64	24min	11.4	85.3	3.2	0.2
NPG	16	7min	65.1	24.9	9.5	0.6
TRPO	16	7min	62.9	26.5	10.1	0.6

### 1.7.2 Atari Benchmark

Tianshou also provides reliable and reproducible Atari 10M benchmark.

Every experiment is conducted under 10 random seeds for 10M steps. Please refer to <https://github.com/thu-ml/tianshou/tree/master/examples/atari> for source code and refer to <https://wandb.ai/tianshou/atari.benchmark/reports/Atari-Benchmark-VmldzoxOTA1NzA5> for detailed results hosted on wandb.

The table below compares the performance of Tianshou against published results on Atari games. We use max average return in 10M timesteps as the reward metric (**to be consistent with Mujoco**). / means results are not provided. The best-performing baseline on each task is highlighted in boldface. Referenced baselines include [Google Dopamine](#) and [OpenAI Baselines](#).

Task		Pong	Break-out	Enduro	Qbert	MsPac-man	Seaquest	SpaceInvaders
DQN	Tianshou	<b>20.2</b> $\pm$ <b>2.3</b>	<b>133.5</b> $\pm$ <b>44.6</b>	997.9 $\pm$ 180.6	<b>11620.2</b> $\pm$ <b>786.1</b>	2324.8 $\pm$ 359.8	<b>3213.9</b> $\pm$ <b>381.6</b>	947.9 $\pm$ 155.3
	Dopamine	9.8	92.2	<b>2126.9</b>	6836.7	<b>2451.3</b>	1406.6	<b>1559.1</b>
	OpenAI Baselines	16.5	131.5	479.8	3254.8	/	1164.1	1129.5 $\pm$ 145.3
C51	Tianshou	<b>20.6</b> $\pm$ <b>2.4</b>	<b>412.9</b> $\pm$ <b>35.8</b>	<b>940.8</b> $\pm$ <b>133.9</b>	<b>12513.2</b> $\pm$ <b>1274.6</b>	2254.9 $\pm$ 201.2	<b>3305.4</b> $\pm$ <b>1524.3</b>	557.3
	Dopamine	17.4	222.4	665.3	9924.5	<b>2860.4</b>	1706.6	<b>604.6</b> $\pm$ <b>157.5</b>
Rainbow	Tianshou	<b>20.2</b> $\pm$ <b>3.0</b>	<b>440.4</b> $\pm$ <b>50.1</b>	1496.1 $\pm$ 112.3	14224.8 $\pm$ 1230.1	2524.2 $\pm$ 338.8	1934.6 $\pm$ 376.4	<b>1178.4</b>
	Dopamine	19.1	47.9	<b>2185.1</b>	<b>15682.2</b>	<b>3161.7</b>	<b>3328.9</b>	459.9
IQN	Tianshou	<b>20.7</b> $\pm$ <b>2.9</b>	<b>355.9</b> $\pm$ <b>22.7</b>	<b>1252.7</b> $\pm$ <b>118.1</b>	<b>14409.2</b> $\pm$ <b>808.6</b>	2228.6 $\pm$ 253.1	5341.2 $\pm$ 670.2	667.8 $\pm$ 81.5
	Dopamine	19.6	96.3	1227.6	12496.7	<b>4422.7</b>	<b>16418</b>	<b>1358.2</b> $\pm$ <b>267.6</b>
PPO	Tianshou	<b>20.3</b> $\pm$ <b>1.2</b>	<b>283.0</b> $\pm$ <b>74.3</b>	<b>1098.9</b> $\pm$ <b>110.5</b>	<b>12341.8</b> $\pm$ <b>1760.7</b>	1699.4 $\pm$ 248.0	1035.2 $\pm$ 353.6	1641.3
	OpenAI Baselines	13.7	114.3	350.2	7012.1	/	<b>1218.9</b>	<b>1787.5</b> $\pm$ <b>340.8</b>
QR-DQN	Tianshou	20.7 $\pm$ 2.0	228.3 $\pm$ 27.3	951.7 $\pm$ 333.5	14761.5 $\pm$ 862.9	2259.3 $\pm$ 269.2	4187.6 $\pm$ 725.7	1114.7 $\pm$ 116.9
FQF	Tianshou	20.4 $\pm$ 2.5	382.6 $\pm$ 29.5	1816.8 $\pm$ 314.3	15301.2 $\pm$ 684.1	2506.6 $\pm$ 402.5	8051.5 $\pm$ 3155.6	2558.3

Please note that the comparison table for both two benchmarks could NOT be used to prove which implementation is “better”. The hyperparameters of the algorithms vary across different implementations. Also, the reward metric is not strictly the same (e.g. Tianshou uses max average return in 10M steps but OpenAI Baselines only report average return at 10M steps, which is unfair). Lastly, Tianshou always uses 10 random seeds while others might use fewer. The comparison is here only to show Tianshou’s reliability.

## 1.8 Cheat Sheet

This page shows some code snippets of how to use Tianshou to develop new algorithms / apply algorithms to new scenarios.

By the way, some of these issues can be resolved by using a `gym.Wrapper`. It could be a universal solution in the policy-environment interaction. But you can also use the batch processor *Handle Batched Data Stream in Collector*.

### 1.8.1 Build Policy Network

See *Build the Network*.

### 1.8.2 Build New Policy

See *BasePolicy*.

### 1.8.3 Customize Training Process

See *Train a Policy with Customized Codes*.

### 1.8.4 Resume Training Process

This is related to [Issue 349](#).

To resume training process from an existing checkpoint, you need to do the following things in the training process:

1. Make sure you write `save_checkpoint_fn` which saves everything needed in the training process, i.e., policy, optim, buffer; pass it to trainer;
2. Use `TensorboardLogger`;
3. To adjust the save frequency, specify `save_interval` when initializing `TensorboardLogger`.

And to successfully resume from a checkpoint:

1. Load everything needed in the training process **before trainer initialization**, i.e., policy, optim, buffer;
2. Set `resume_from_log=True` with trainer;

We provide an example to show how these steps work: checkout `test_c51.py`, `test_ppo.py` or `test_il_bcq.py` by running

```
$ python3 test/discrete/test_c51.py # train some epoch
$ python3 test/discrete/test_c51.py --resume # restore from existing log and_
↪ continuing training
```

To correctly render the data (including several tfevent files), we highly recommend using `tensorboard >= 2.5.0` (see [here](#) for the reason). Otherwise, it may cause overlapping issue that you need to manually handle with.

### 1.8.5 Parallel Sampling

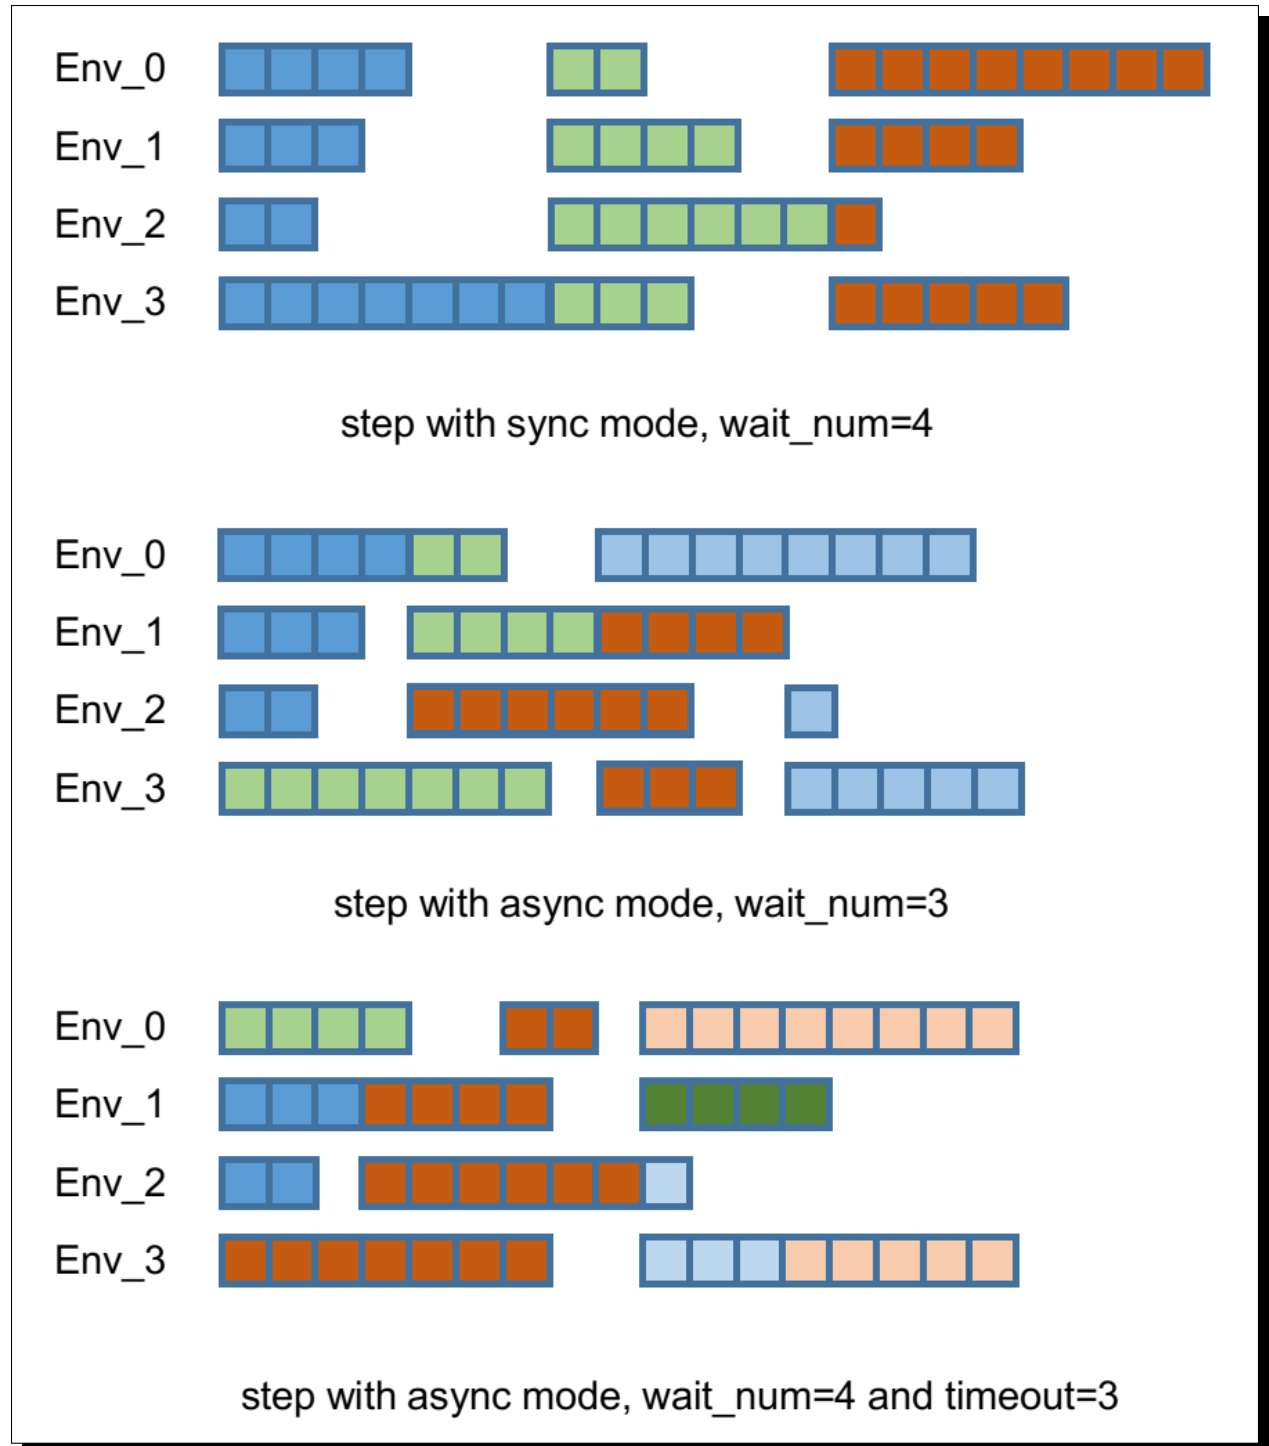
Tianshou provides the following classes for vectorized environment:

- `DummyVectorEnv` is for pseudo-parallel simulation (implemented with a for-loop, useful for debugging).
- `SubprocVectorEnv` uses multiple processes for parallel simulation. This is the most often choice for parallel simulation.
- `ShmemVectorEnv` has a similar implementation to `SubprocVectorEnv`, but is optimized (in terms of both memory footprint and simulation speed) for environments with large observations such as images.
- `RayVectorEnv` is currently the only choice for parallel simulation in a cluster with multiple machines.

Although these classes are optimized for different scenarios, they have exactly the same APIs because they are subclasses of `BaseVectorEnv`. Just provide a list of functions who return environments upon called, and it is all set.

```
env_fns = [lambda x=i: MyTestEnv(size=x) for i in [2, 3, 4, 5]]
venv = SubprocVectorEnv(env_fns)  # DummyVectorEnv, ShmemVectorEnv, or RayVectorEnv, ↵
↵whichever you like.
venv.reset()  # returns the initial observations of each environment
venv.step(actions)  # provide actions for each environment and get their results
```

**An example of sync/async VectorEnv (steps with the same color end up in one batch that is disposed by the policy at the same time).**



By default, parallel environment simulation is synchronous: a step is done after all environments have finished a step. Synchronous simulation works well if each step of environments costs roughly the same time.

In case the time cost of environments varies a lot (e.g. 90% step cost 1s, but 10% cost 10s) where slow environments lag fast environments behind, async simulation can be used (related to [Issue 103](#)). The idea is to start those finished environments without waiting for slow environments.

Asynchronous simulation is a built-in functionality of *BaseVectorEnv*. Just provide `wait_num` or `timeout` (or both) and async simulation works.

```
env_fns = [lambda x=i: MyTestEnv(size=x, sleep=x) for i in [2, 3, 4, 5]]
# DummyVectorEnv, ShmemVectorEnv, or RayVectorEnv, whichever you like.
venv = SubprocVectorEnv(env_fns, wait_num=3, timeout=0.2)
venv.reset() # returns the initial observations of each environment
# returns "wait_num" steps or finished steps after "timeout" seconds,
# whichever occurs first.
venv.step(actions, ready_id)
```

If we have 4 envs and set `wait_num = 3`, each of the step only returns 3 results of these 4 envs.

You can treat the `timeout` parameter as a dynamic `wait_num`. In each vectorized step it only returns the environments finished within the given time. If there is no such environment, it will wait until any of them finished.

The figure in the right gives an intuitive comparison among synchronous/asynchronous simulation.

**Warning:** If you use your own environment, please make sure the `seed` method is set up properly, e.g.,

```
def seed(self, seed):
    np.random.seed(seed)
```

Otherwise, the outputs of these envs may be the same with each other.

## 1.8.6 EnvPool Integration

`EnvPool` is a C++-based vectorized environment implementation and is way faster than the above solutions. The APIs are almost the same as above four classes, so that means you can directly switch the vectorized environment to `envpool` and get immediate speed-up.

Currently it supports Atari, VizDoom, `toy_text` and `classic_control` environments. For more information, please refer to `EnvPool`'s documentation.

```
# install envpool: pip3 install envpool

import envpool
envs = envpool.make_gym("CartPole-v0", num_envs=10)
collector = Collector(policy, envs, buffer)
```

Here are some examples: [https://github.com/sail-sg/envpool/tree/master/examples/tianshou\\_examples](https://github.com/sail-sg/envpool/tree/master/examples/tianshou_examples)

## 1.8.7 Handle Batched Data Stream in Collector

This is related to [Issue 42](#).

If you want to get log stat from data stream / pre-process batch-image / modify the reward with given env info, use `preproces_fn` in `Collector`. This is a hook which will be called before the data adding into the buffer.

It will receive with "obs" and "env\_id" when the collector resets the environment, and will receive six keys "obs\_next", "rew", "done", "info", "policy", "env\_id" in a normal env step. It returns either a dict or a `Batch` with the modified keys and values.

These variables are intended to gather all the information requires to keep track of a simulation step, namely the (observation, action, reward, done flag, next observation, info, intermediate result of the policy) at time  $t$ , for the whole duration of the simulation.

For example, you can write your hook as:

```

import numpy as np
from collections import deque

class MyProcessor:
    def __init__(self, size=100):
        self.episode_log = None
        self.main_log = deque(maxlen=size)
        self.main_log.append(0)
        self.baseline = 0

    def preprocess_fn(**kwargs):
        """change reward to zero mean"""
        # if obs && env_id exist -> reset
        # if obs_next/act/rew/done/policy/env_id exist -> normal step
        if 'rew' not in kwargs:
            # means that it is called after env.reset(), it can only process the obs
            return Batch() # none of the variables are needed to be updated
        else:
            n = len(kwargs['rew']) # the number of envs in collector
            if self.episode_log is None:
                self.episode_log = [[] for i in range(n)]
            for i in range(n):
                self.episode_log[i].append(kwargs['rew'][i])
                kwargs['rew'][i] -= self.baseline
            for i in range(n):
                if kwargs['done']:
                    self.main_log.append(np.mean(self.episode_log[i]))
                    self.episode_log[i] = []
                    self.baseline = np.mean(self.main_log)
            return Batch(rew=kwargs['rew'])

```

And finally,

```

test_processor = MyProcessor(size=100)
collector = Collector(policy, env, buffer, preprocess_fn=test_processor.preprocess_fn)

```

Some examples are in `test/base/test_collector.py`.

## 1.8.8 RNN-style Training

This is related to [Issue 19](#).

First, add an argument “stack\_num” to `ReplayBuffer`, `VectorReplayBuffer`, or other types of buffer you are using, like:

```

buf = ReplayBuffer(size=size, stack_num=stack_num)

```

Then, change the network to recurrent-style, for example, `Recurrent`, `RecurrentActorProb` and `RecurrentCritic`.

The above code supports only stacked-observation. If you want to use stacked-action (for Q(stacked-s, stacked-a)), stacked-reward, or other stacked variables, you can add a `gym.Wrapper` to modify the state representation. For example, if we add a wrapper that map `[s, a]` pair to a new state:

- Before: `(s, a, s', r, d)` stored in replay buffer, and get stacked `s`;

- After applying wrapper:  $([s, a], a, [s', a'], r, d)$  stored in replay buffer, and get both stacked  $s$  and  $a$ .

## 1.8.9 Multi-GPU Training

To enable training an RL agent with multiple GPUs for a standard environment (i.e., without nested observation) with default networks provided by Tianshou:

1. Import `DataParallelNet` from `tianshou.utils.net.common`;
2. Change the device argument to `None` in the existing networks such as `Net`, `Actor`, `Critic`, `ActorProb`
3. Apply `DataParallelNet` wrapper to these networks.

```
from tianshou.utils.net.common import Net, DataParallelNet
from tianshou.utils.net.discrete import Actor, Critic

actor = DataParallelNet(Actor(net, args.action_shape, device=None).to(args.device))
critic = DataParallelNet(Critic(net, device=None).to(args.device))
```

Yes, that's all! This general approach can be applied to almost all kinds of algorithms implemented in Tianshou. We provide a complete script to show how to run multi-GPU: [test/discrete/test\\_ppo.py](#)

As for other cases such as customized network or environments that have a nested observation, here are the rules:

1. The data format transformation (numpy -> cuda) is done in the `DataParallelNet` wrapper; your customized network should not apply any kinds of data format transformation;
2. Create a similar class that inherit `DataParallelNet`, which is only in charge of data format transformation (numpy -> cuda);
3. Do the same things above.

## 1.8.10 User-defined Environment and Different State Representation

This is related to [Issue 38](#) and [Issue 69](#).

First of all, your self-defined environment must follow the Gym's API, some of them are listed below:

- `reset()` -> state
- `step(action)` -> state, reward, done, info
- `seed(s)` -> List[int]
- `render(mode)` -> Any
- `close()` -> None
- `observation_space`: `gym.Space`
- `action_space`: `gym.Space`

The state can be a `numpy.ndarray` or a Python dictionary. Take "FetchReach-v1" as an example:

```
>>> e = gym.make('FetchReach-v1')
>>> e.reset()
{'observation': array([ 1.34183265e+00,  7.49100387e-01,  5.34722720e-01,  1.
↪ 97805133e-04,
        7.15193042e-05,  7.73933014e-06,  5.51992816e-08, -2.42927453e-06,
        4.73325650e-06, -2.28455228e-06]),
```

(continues on next page)

(continued from previous page)

```
'achieved_goal': array([1.34183265, 0.74910039, 0.53472272]),
'desired_goal': array([1.24073906, 0.77753463, 0.63457791])}
```

It shows that the state is a dictionary which has 3 keys. It will stored in *ReplayBuffer* as:

```
>>> from tianshou.data import Batch, ReplayBuffer
>>> b = ReplayBuffer(size=3)
>>> b.add(Batch(obs=e.reset(), act=0, rew=0, done=0))
>>> print(b)
ReplayBuffer(
  act: array([0, 0, 0]),
  done: array([False, False, False]),
  obs: Batch(
    achieved_goal: array([[1.34183265, 0.74910039, 0.53472272],
                          [0.          , 0.          , 0.          ],
                          [0.          , 0.          , 0.          ]]),
    desired_goal: array([[1.42154265, 0.62505137, 0.62929863],
                        [0.          , 0.          , 0.          ],
                        [0.          , 0.          , 0.          ]]),
    observation: array([[ 1.34183265e+00,  7.49100387e-01,  5.34722720e-01,
                        1.97805133e-04,  7.15193042e-05,  7.73933014e-06,
                        5.51992816e-08, -2.42927453e-06,  4.73325650e-06,
                        -2.28455228e-06],
                        [ 0.00000000e+00,  0.00000000e+00,  0.00000000e+00,
                        0.00000000e+00,  0.00000000e+00,  0.00000000e+00,
                        0.00000000e+00,  0.00000000e+00,  0.00000000e+00,
                        0.00000000e+00],
                        [ 0.00000000e+00,  0.00000000e+00,  0.00000000e+00,
                        0.00000000e+00,  0.00000000e+00,  0.00000000e+00,
                        0.00000000e+00,  0.00000000e+00,  0.00000000e+00,
                        0.00000000e+00]]),
    ),
  rew: array([0, 0, 0]),
)
>>> print(b.obs.achieved_goal)
[[1.34183265 0.74910039 0.53472272]
 [0.          0.          0.          ]
 [0.          0.          0.          ]]
```

And the data batch sampled from this replay buffer:

```
>>> batch, indices = b.sample(2)
>>> batch.keys()
['act', 'done', 'info', 'obs', 'obs_next', 'policy', 'rew']
>>> batch.obs[-1]
Batch(
  achieved_goal: array([1.34183265, 0.74910039, 0.53472272]),
  desired_goal: array([1.42154265, 0.62505137, 0.62929863]),
  observation: array([ 1.34183265e+00,  7.49100387e-01,  5.34722720e-01,  1.
↪ 97805133e-04,
                        7.15193042e-05,  7.73933014e-06,  5.51992816e-08, -2.
↪ 42927453e-06,
                        4.73325650e-06, -2.28455228e-06]),
)
>>> batch.obs.desired_goal[-1] # recommended
array([1.42154265, 0.62505137, 0.62929863])
>>> batch.obs[-1].desired_goal # not recommended
```

(continues on next page)



(continued from previous page)

```
array([1.42154265, 0.62505137, 0.62929863])
>>> batch[-1].obs.desired_goal # not recommended
array([1.42154265, 0.62505137, 0.62929863])
```

Thus, in your self-defined network, just change the forward function as:

```
def forward(self, s, ...):
    # s is a batch
    observation = s.observation
    achieved_goal = s.achieved_goal
    desired_goal = s.desired_goal
    ...
```

For self-defined class, the replay buffer will store the reference into a `numpy.ndarray`, e.g.:

```
>>> import networkx as nx
>>> b = ReplayBuffer(size=3)
>>> b.add(Batch(obs=nx.Graph(), act=0, rew=0, done=0))
>>> print(b)
ReplayBuffer(
  act: array([0, 0, 0]),
  done: array([0, 0, 0]),
  info: Batch(),
  obs: array([<networkx.classes.graph.Graph object at 0x7f5c607826a0>, None,
             None], dtype=object),
  policy: Batch(),
  rew: array([0, 0, 0]),
)
```

But the state stored in the buffer may be a shallow-copy. To make sure each of your state stored in the buffer is distinct, please return the deep-copy version of your state in your env:

```
def reset():
    return copy.deepcopy(self.graph)
def step(action):
    ...
    return copy.deepcopy(self.graph), reward, done, {}
```

**Note:** Please make sure this variable is numpy-compatible, e.g., `np.array([variable])` will not result in an empty array. Otherwise, `ReplayBuffer` cannot create a numpy array to store it.

## 1.8.11 Multi-Agent Reinforcement Learning

This is related to [Issue 121](#). The discussion is still goes on.

With the flexible core APIs, Tianshou can support multi-agent reinforcement learning with minimal efforts.

Currently, we support three types of multi-agent reinforcement learning paradigms:

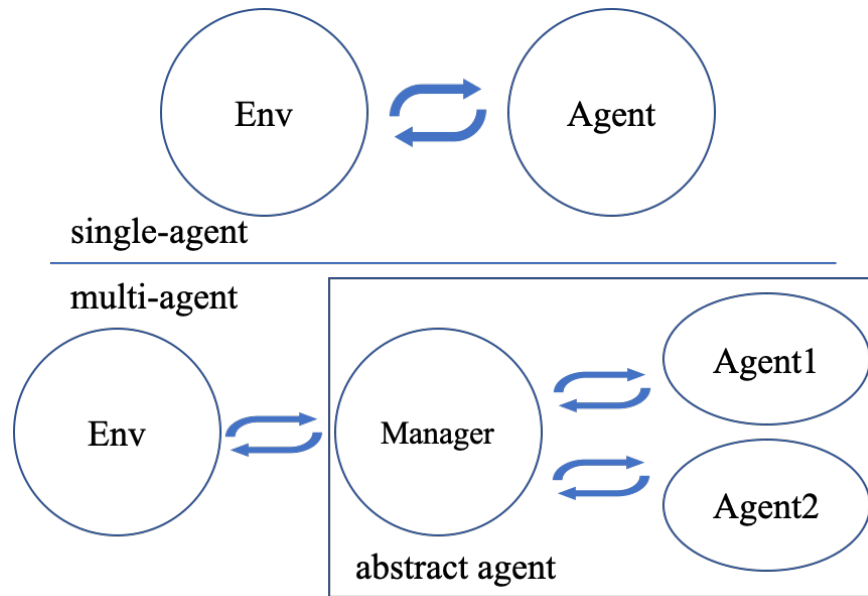
1. Simultaneous move: at each timestep, all the agents take their actions (example: MOBA games)
2. Cyclic move: players take action in turn (example: Go game)
3. Conditional move, at each timestep, the environment conditionally selects an agent to take action. (example: [Pig Game](#))

We mainly address these multi-agent RL problems by converting them into traditional RL formulations.

For simultaneous move, the solution is simple: we can just add a `num_agent` dimension to state, action, and reward. Nothing else is going to change.

For 2 & 3 (cyclic move and conditional move), they can be unified into a single framework: at each timestep, the environment selects an agent with id `agent_id` to play. Since multi-agents are usually wrapped into one object (which we call “abstract agent”), we can pass the `agent_id` to the “abstract agent”, leaving it to further call the specific agent.

In addition, legal actions in multi-agent RL often vary with timestep (just like Go games), so the environment should also pass the legal action mask to the “abstract agent”, where the mask is a boolean array that “True” for available actions and “False” for illegal actions at the current step. Below is a figure that explains the abstract agent.



The above description gives rise to the following formulation of multi-agent RL:

```
act = policy(state, agent_id, mask)
(next_state, next_agent_id, next_mask), reward = env.step(act)
```

By constructing a new state `state_ = (state, agent_id, mask)`, essentially we can return to the typical formulation of RL:

```
act = policy(state_)
next_state_, reward = env.step(act)
```

Following this idea, we write a tiny example of playing **Tic Tac Toe** against a random player by using a Q-learning algorithm. The tutorial is at [Multi-Agent RL](#).

## 1.9 tianshou.data

### 1.9.1 Batch

**class** tianshou.data.Batch (batch\_dict: Optional[Union[dict, tianshou.data.batch.Batch, Sequence[Union[dict, tianshou.data.batch.Batch]], numpy.ndarray]] = None, copy: bool = False, \*\*kwargs: Any)

Bases: object

The internal data structure in Tianshou.

Batch is a kind of supercharged array (of temporal data) stored individually in a (recursive) dictionary of object that can be either numpy array, torch tensor, or batch themselves. It is designed to make it extremely easily to access, manipulate and set partial view of the heterogeneous data conveniently.

For a detailed description, please refer to [Understand Batch](#).

**\_\_getitem\_\_** (index: Union[str, slice, int, numpy.ndarray, List[int]]) → Any  
Return self[index].

**\_\_setitem\_\_** (index: Union[str, slice, int, numpy.ndarray, List[int]], value: Any) → None  
Assign value to self[index].

**to\_numpy** () → None  
Change all torch.Tensor to numpy.ndarray in-place.

**to\_torch** (dtype: Optional[torch.dtype] = None, device: Union[str, int, torch.device] = 'cpu') → None  
Change all numpy.ndarray to torch.Tensor in-place.

**cat** (batches: Union[tianshou.data.batch.Batch, Sequence[Union[dict, tianshou.data.batch.Batch]]]) → None  
Concatenate a list of (or one) Batch objects into current batch.

**static cat** (batches: Sequence[Union[dict, tianshou.data.batch.Batch]]) → tianshou.data.batch.Batch  
Concatenate a list of Batch object into a single new batch.

For keys that are not shared across all batches, batches that do not have these keys will be padded by zeros with appropriate shapes. E.g.

```
>>> a = Batch(a=np.zeros([3, 4]), common=Batch(c=np.zeros([3, 5])))
>>> b = Batch(b=np.zeros([4, 3]), common=Batch(c=np.zeros([4, 5])))
>>> c = Batch.cat([a, b])
>>> c.a.shape
(7, 4)
>>> c.b.shape
(7, 3)
>>> c.common.c.shape
(7, 5)
```

**stack** (batches: Sequence[Union[dict, tianshou.data.batch.Batch]], axis: int = 0) → None  
Stack a list of Batch object into current batch.

**static stack** (batches: Sequence[Union[dict, tianshou.data.batch.Batch]], axis: int = 0) → tianshou.data.batch.Batch  
Stack a list of Batch object into a single new batch.

For keys that are not shared across all batches, batches that do not have these keys will be padded by zeros. E.g.

```

>>> a = Batch(a=np.zeros([4, 4]), common=Batch(c=np.zeros([4, 5])))
>>> b = Batch(b=np.zeros([4, 6]), common=Batch(c=np.zeros([4, 5])))
>>> c = Batch.stack([a, b])
>>> c.a.shape
(2, 4, 4)
>>> c.b.shape
(2, 4, 6)
>>> c.common.c.shape
(2, 4, 5)

```

**Note:** If there are keys that are not shared across all batches, `stack` with `axis != 0` is undefined, and will cause an exception.

**empty\_**(*index: Optional[Union[slice, int, numpy.ndarray, List[int]]] = None*) → `tianshou.data.batch.Batch`  
 Return an empty Batch object with 0 or None filled.

If “index” is specified, it will only reset the specific indexed-data.

```

>>> data.empty_()
>>> print(data)
Batch(
  a: array([[0., 0.],
            [0., 0.]])
  b: array([None, None], dtype=object),
)
>>> b={'c': [2., 'st'], 'd': [1., 0.]}
>>> data = Batch(a=[False, True], b=b)
>>> data[0] = Batch.empty_(data[1])
>>> data
Batch(
  a: array([False, True]),
  b: Batch(
    c: array([None, 'st']),
    d: array([0., 0.]),
  ),
)

```

**static empty**(*batch: tianshou.data.batch.Batch, index: Optional[Union[slice, int, numpy.ndarray, List[int]]] = None*) → `tianshou.data.batch.Batch`  
 Return an empty Batch object with 0 or None filled.

The shape is the same as the given Batch.

**update**(*batch: Optional[Union[dict, tianshou.data.batch.Batch]] = None, \*\*kwargs: Any*) → `None`  
 Update this batch from another dict/Batch.

**\_\_len\_\_**() → `int`  
 Return `len(self)`.

**is\_empty**(*recurse: bool = False*) → `bool`  
 Test if a Batch is empty.

If `recurse=True`, it further tests the values of the object; else it only tests the existence of any key.

`b.is_empty(recurse=True)` is mainly used to distinguish `Batch(a=Batch(a=Batch()))` and `Batch(a=1)`. They both raise exceptions when applied to `len()`, but the former can be used in `cat`, while the latter is a scalar and cannot be used in `cat`.

Another usage is in `__len__`, where we have to skip checking the length of recursively empty Batch.

```
>>> Batch().is_empty()
True
>>> Batch(a=Batch(), b=Batch(c=Batch())).is_empty()
False
>>> Batch(a=Batch(), b=Batch(c=Batch())).is_empty(recurse=True)
True
>>> Batch(d=1).is_empty()
False
>>> Batch(a=np.float64(1.0)).is_empty()
False
```

### property shape

Return self.shape.

**split** (*size: int, shuffle: bool = True, merge\_last: bool = False*) → `Iterator[tianshou.data.batch.Batch]`  
Split whole data into multiple small batches.

#### Parameters

- **size** (*int*) – divide the data batch with the given size, but one batch if the length of the batch is smaller than “size”.
- **shuffle** (*bool*) – randomly shuffle the entire data batch if it is True, otherwise remain in the same. Default to True.
- **merge\_last** (*bool*) – merge the last batch into the previous one. Default to False.

## 1.9.2 Buffer

### ReplayBuffer

**class** `tianshou.data.ReplayBuffer` (*size: int, stack\_num: int = 1, ignore\_obs\_next: bool = False, save\_only\_last\_obs: bool = False, sample\_avail: bool = False, \*\*kwargs: Any*)

Bases: `object`

*ReplayBuffer* stores data generated from interaction between the policy and environment.

ReplayBuffer can be considered as a specialized form (or management) of Batch. It stores all the data in a batch with circular-queue style.

For the example usage of ReplayBuffer, please check out Section Buffer in *Basic concepts in Tianshou*.

#### Parameters

- **size** (*int*) – the maximum size of replay buffer.
- **stack\_num** (*int*) – the frame-stack sampling argument, should be greater than or equal to 1. Default to 1 (no stacking).
- **ignore\_obs\_next** (*bool*) – whether to store obs\_next. Default to False.
- **save\_only\_last\_obs** (*bool*) – only save the last obs/obs\_next when it has a shape of (timestamp, ...) because of temporal stacking. Default to False.
- **sample\_avail** (*bool*) – the parameter indicating sampling only available index when using frame-stack sampling method. Default to False.

`__len__()` → `int`  
Return len(self).

**save\_hdf5** (*path*: str, *compression*: Optional[str] = None) → None  
Save replay buffer to HDF5 file.

**classmethod load\_hdf5** (*path*: str, *device*: Optional[str] = None) → tianshou.data.buffer.base.ReplayBuffer  
Load replay buffer from HDF5 file.

**classmethod from\_data** (*obs*: h5py.\_hl.dataset.Dataset, *act*: h5py.\_hl.dataset.Dataset, *rew*: h5py.\_hl.dataset.Dataset, *done*: h5py.\_hl.dataset.Dataset, *obs\_next*: h5py.\_hl.dataset.Dataset) → tianshou.data.buffer.base.ReplayBuffer

**reset** (*keep\_statistics*: bool = False) → None  
Clear all the data in replay buffer and episode statistics.

**set\_batch** (*batch*: tianshou.data.batch.Batch) → None  
Manually choose the batch you want the ReplayBuffer to manage.

**unfinished\_index** () → numpy.ndarray  
Return the index of unfinished episode.

**prev** (*index*: Union[int, numpy.ndarray]) → numpy.ndarray  
Return the index of previous transition.  
  
The index won't be modified if it is the beginning of an episode.

**next** (*index*: Union[int, numpy.ndarray]) → numpy.ndarray  
Return the index of next transition.  
  
The index won't be modified if it is the end of an episode.

**update** (*buffer*: tianshou.data.buffer.base.ReplayBuffer) → numpy.ndarray  
Move the data from the given buffer to current buffer.  
  
Return the updated indices. If update fails, return an empty array.

**add** (*batch*: tianshou.data.batch.Batch, *buffer\_ids*: Optional[Union[numpy.ndarray, List[int]]] = None) → Tuple[numpy.ndarray, numpy.ndarray, numpy.ndarray, numpy.ndarray]  
Add a batch of data into replay buffer.

#### Parameters

- **batch** (`Batch`) – the input data batch. Its keys must belong to the 7 reserved keys, and “obs”, “act”, “rew”, “done” is required.
- **buffer\_ids** – to make consistent with other buffer's add function; if it is not None, we assume the input batch's first dimension is always 1.

Return (current\_index, episode\_reward, episode\_length, episode\_start\_index). If the episode is not finished, the return value of episode\_length and episode\_reward is 0.

**sample\_indices** (*batch\_size*: int) → numpy.ndarray  
Get a random sample of index with size = batch\_size.  
  
Return all available indices in the buffer if batch\_size is 0; return an empty numpy array if batch\_size < 0 or no available index can be sampled.

**sample** (*batch\_size*: int) → Tuple[tianshou.data.batch.Batch, numpy.ndarray]  
Get a random sample from buffer with size = batch\_size.  
  
Return all the data in the buffer if batch\_size is 0.

**Returns** Sample data and its corresponding index inside the buffer.

**get** (*index*: Union[int, List[int], numpy.ndarray], *key*: str, *default\_value*: Optional[Any] = None, *stack\_num*: Optional[int] = None) → Union[tianshou.data.batch.Batch, numpy.ndarray]  
Return the stacked result.

E.g., if you set `key = "obs"`, `stack_num = 4`, `index = t`, it returns the stacked result as `[obs[t-3], obs[t-2], obs[t-1], obs[t]]`.

#### Parameters

- **index** – the index for getting stacked data.
- **key** (*str*) – the key to get, should be one of the reserved\_keys.
- **default\_value** – if the given key’s data is not found and default\_value is set, return this default\_value.
- **stack\_num** (*int*) – Default to self.stack\_num.

**\_\_getitem\_\_** (*index: Union[slice, int, List[int], numpy.ndarray]*) → `tianshou.data.batch.Batch`  
Return a data batch: `self[index]`.

If `stack_num` is larger than 1, return the stacked obs and obs\_next with shape (batch, len, ...).

### PrioritizedReplayBuffer

**class** `tianshou.data.PrioritizedReplayBuffer` (*size: int, alpha: float, beta: float, weight\_norm: bool = True, \*\*kwargs: Any*)  
Bases: `tianshou.data.buffer.base.ReplayBuffer`

Implementation of Prioritized Experience Replay. arXiv:1511.05952.

#### Parameters

- **alpha** (*float*) – the prioritization exponent.
- **beta** (*float*) – the importance sample soft coefficient.
- **weight\_norm** (*bool*) – whether to normalize returned weights with the maximum weight value within the batch. Default to True.

#### See also:

Please refer to [ReplayBuffer](#) for other APIs’ usage.

**init\_weight** (*index: Union[int, numpy.ndarray]*) → None

**update** (*buffer: tianshou.data.buffer.base.ReplayBuffer*) → `numpy.ndarray`  
Move the data from the given buffer to current buffer.

Return the updated indices. If update fails, return an empty array.

**add** (*batch: tianshou.data.batch.Batch, buffer\_ids: Optional[Union[numpy.ndarray, List[int]]] = None*)  
→ `Tuple[numpy.ndarray, numpy.ndarray, numpy.ndarray, numpy.ndarray]`  
Add a batch of data into replay buffer.

#### Parameters

- **batch** (*Batch*) – the input data batch. Its keys must belong to the 7 reserved keys, and “obs”, “act”, “rew”, “done” is required.
- **buffer\_ids** – to make consistent with other buffer’s add function; if it is not None, we assume the input batch’s first dimension is always 1.

Return (current\_index, episode\_reward, episode\_length, episode\_start\_index). If the episode is not finished, the return value of episode\_length and episode\_reward is 0.

**sample\_indices** (*batch\_size: int*) → `numpy.ndarray`  
Get a random sample of index with size = batch\_size.

Return all available indices in the buffer if `batch_size` is 0; return an empty numpy array if `batch_size < 0` or no available index can be sampled.

**get\_weight** (*index*: *Union[int, numpy.ndarray]*) → *Union[float, numpy.ndarray]*

Get the importance sampling weight.

The “weight” in the returned Batch is the weight on loss function to debias the sampling process (some transition tuples are sampled more often so their losses are weighted less).

**update\_weight** (*index*: *numpy.ndarray*, *new\_weight*: *Union[numpy.ndarray, torch.Tensor]*) → *None*

Update priority weight by index in this buffer.

#### Parameters

- **index** (*np.ndarray*) – index you want to update weight.
- **new\_weight** (*np.ndarray*) – new priority weight you want to update.

**\_\_getitem\_\_** (*index*: *Union[slice, int, List[int], numpy.ndarray]*) → *tianshou.data.batch.Batch*

Return a data batch: `self[index]`.

If `stack_num` is larger than 1, return the stacked obs and obs\_next with shape (batch, len, ...).

**set\_beta** (*beta*: *float*) → *None*

## ReplayBufferManager

**class** `tianshou.data.ReplayBufferManager` (*buffer\_list*: *List[tianshou.data.buffer.base.ReplayBuffer]*)

Bases: `tianshou.data.buffer.base.ReplayBuffer`

ReplayBufferManager contains a list of ReplayBuffer with exactly the same configuration.

These replay buffers have contiguous memory layout, and the storage space each buffer has is a shallow copy of the topmost memory.

**Parameters** **buffer\_list** – a list of ReplayBuffer needed to be handled.

See also:

Please refer to [ReplayBuffer](#) for other APIs’ usage.

**\_\_len\_\_** () → *int*

Return `len(self)`.

**reset** (*keep\_statistics*: *bool = False*) → *None*

Clear all the data in replay buffer and episode statistics.

**set\_batch** (*batch*: *tianshou.data.batch.Batch*) → *None*

Manually choose the batch you want the ReplayBuffer to manage.

**unfinished\_index** () → *numpy.ndarray*

Return the index of unfinished episode.

**prev** (*index*: *Union[int, numpy.ndarray]*) → *numpy.ndarray*

Return the index of previous transition.

The index won’t be modified if it is the beginning of an episode.

**next** (*index*: *Union[int, numpy.ndarray]*) → *numpy.ndarray*

Return the index of next transition.

The index won’t be modified if it is the end of an episode.

**update** (*buffer*: *tianshou.data.buffer.base.ReplayBuffer*) → *numpy.ndarray*

The ReplayBufferManager cannot be updated by any buffer.



**add** (*batch*: *tianshou.data.batch.Batch*, *buffer\_ids*: *Optional[Union[numpy.ndarray, List[int]]]* = *None*)  
 → *Tuple[numpy.ndarray, numpy.ndarray, numpy.ndarray, numpy.ndarray]*  
 Add a batch of data into *ReplayBufferManager*.

Each of the data's length (first dimension) must equal to the length of *buffer\_ids*. By default *buffer\_ids* is *[0, 1, ..., buffer\_num - 1]*.

Return (*current\_index*, *episode\_reward*, *episode\_length*, *episode\_start\_index*). If the episode is not finished, the return value of *episode\_length* and *episode\_reward* is 0.

**sample\_indices** (*batch\_size*: *int*) → *numpy.ndarray*  
 Get a random sample of index with size = *batch\_size*.

Return all available indices in the buffer if *batch\_size* is 0; return an empty numpy array if *batch\_size* < 0 or no available index can be sampled.

### PrioritizedReplayBufferManager

**class** *tianshou.data.PrioritizedReplayBufferManager* (*buffer\_list*: *Sequence[tianshou.data.buffer.prio.PrioritizedReplayBuffer]*, *Se-*  
*quence[tianshou.data.buffer.prio.PrioritizedReplayBuffer]*)  
 Bases: *tianshou.data.buffer.prio.PrioritizedReplayBuffer*, *tianshou.data.buffer.manager.ReplayBufferManager*

*PrioritizedReplayBufferManager* contains a list of *PrioritizedReplayBuffer* with exactly the same configuration.

These replay buffers have contiguous memory layout, and the storage space each buffer has is a shallow copy of the topmost memory.

**Parameters** *buffer\_list* – a list of *PrioritizedReplayBuffer* needed to be handled.

**See also:**

Please refer to [ReplayBuffer](#) for other APIs' usage.

### VectorReplayBuffer

**class** *tianshou.data.VectorReplayBuffer* (*total\_size*: *int*, *buffer\_num*: *int*, *\*\*kwargs*: *Any*)  
 Bases: *tianshou.data.buffer.manager.ReplayBufferManager*

*VectorReplayBuffer* contains *n* *ReplayBuffer* with the same size.

It is used for storing transition from different environments yet keeping the order of time.

**Parameters**

- **total\_size** (*int*) – the total size of *VectorReplayBuffer*.
- **buffer\_num** (*int*) – the number of *ReplayBuffer* it uses, which are under the same configuration.

Other input arguments (*stack\_num*/*ignore\_obs\_next*/*save\_only\_last\_obs*/*sample\_avail*) are the same as [ReplayBuffer](#).

**See also:**

Please refer to [ReplayBuffer](#) for other APIs' usage.

## PrioritizedVectorReplayBuffer

```
class tianshou.data.PrioritizedVectorReplayBuffer (total_size: int, buffer_num: int,  
                                                    **kwargs: Any)
```

Bases: `tianshou.data.buffer.manager.PrioritizedReplayBufferManager`

PrioritizedVectorReplayBuffer contains `n` PrioritizedReplayBuffer with same size.

It is used for storing transition from different environments yet keeping the order of time.

### Parameters

- **total\_size** (*int*) – the total size of PrioritizedVectorReplayBuffer.
- **buffer\_num** (*int*) – the number of PrioritizedReplayBuffer it uses, which are under the same configuration.

Other input arguments (`alpha/beta/stack_num/ignore_obs_next/save_only_last_obs/ sample_avail`) are the same as *PrioritizedReplayBuffer*.

### See also:

Please refer to *ReplayBuffer* for other APIs' usage.

**set\_beta** (*beta: float*) → None

## CachedReplayBuffer

```
class tianshou.data.CachedReplayBuffer (main_buffer: tianshou.data.buffer.base.ReplayBuffer,  
                                          cached_buffer_num: int, max_episode_length: int)
```

Bases: `tianshou.data.buffer.manager.ReplayBufferManager`

CachedReplayBuffer contains a given main buffer and `n` cached buffers, `cached_buffer_num * ReplayBuffer (size=max_episode_length)`.

The memory layout is: | `main_buffer` | `cached_buffers[0]` | `cached_buffers[1]` | ... | `cached_buffers[cached_buffer_num - 1]` |.

The data is first stored in cached buffers. When an episode is terminated, the data will move to the main buffer and the corresponding cached buffer will be reset.

### Parameters

- **main\_buffer** (*ReplayBuffer*) – the main buffer whose `.update()` function behaves normally.
- **cached\_buffer\_num** (*int*) – number of ReplayBuffer needs to be created for cached buffer.
- **max\_episode\_length** (*int*) – the maximum length of one episode, used in each cached buffer's maxsize.

### See also:

Please refer to *ReplayBuffer* for other APIs' usage.

**add** (*batch: tianshou.data.batch.Batch, buffer\_ids: Optional[Union[numpy.ndarray, List[int]]] = None*)  
→ `Tuple[numpy.ndarray, numpy.ndarray, numpy.ndarray, numpy.ndarray]`  
Add a batch of data into CachedReplayBuffer.

Each of the data's length (first dimension) must equal to the length of `buffer_ids`. By default the `buffer_ids` is `[0, 1, ..., cached_buffer_num - 1]`.

Return (current\_index, episode\_reward, episode\_length, episode\_start\_index) with each of the shape (len(buffer\_ids), ...), where (current\_index[i], episode\_reward[i], episode\_length[i], episode\_start\_index[i]) refers to the cached\_buffer\_ids[i]th cached buffer’s corresponding episode result.

### 1.9.3 Collector

#### Collector

```
class tianshou.data.Collector (policy:          tianshou.policy.base.BasePolicy,          env:
                                Union[gym.core.Env, tianshou.env.venvs.BaseVectorEnv], buffer:
                                Optional[tianshou.data.buffer.base.ReplayBuffer] = None, pre-
                                process_fn: Optional[Callable[[...], tianshou.data.batch.Batch]]
                                = None, exploration_noise: bool = False)
```

Bases: object

Collector enables the policy to interact with different types of envs with exact number of steps or episodes.

#### Parameters

- **policy** – an instance of the [BasePolicy](#) class.
- **env** – a gym.Env environment or an instance of the [BaseVectorEnv](#) class.
- **buffer** – an instance of the [ReplayBuffer](#) class. If set to None, it will not store the data. Default to None.
- **preprocess\_fn** (*function*) – a function called before the data has been added to the buffer, see issue #42 and [Handle Batched Data Stream in Collector](#). Default to None.
- **exploration\_noise** (*bool*) – determine whether the action needs to be modified with corresponding policy’s exploration noise. If so, “policy. exploration\_noise(act, batch)” will be called automatically to add the exploration noise into action. Default to False.

The “preprocess\_fn” is a function called before the data has been added to the buffer with batch format. It will receive only “obs” and “env\_id” when the collector resets the environment, and will receive six keys “obs\_next”, “rew”, “done”, “info”, “policy” and “env\_id” in a normal env step. It returns either a dict or a [Batch](#) with the modified keys and values. Examples are in “test/base/test\_collector.py”.

---

**Note:** Please make sure the given environment has a time limitation if using n\_episode collect option.

---



---

**Note:** In past versions of Tianshou, the replay buffer that was passed to `__init__` was automatically reset. This is not done in the current implementation.

---

**reset** (*reset\_buffer: bool = True*) → None

Reset the environment, statistics, current data and possibly replay memory.

**Parameters** **reset\_buffer** (*bool*) – if true, reset the replay buffer that is attached to the collector.

**reset\_stat** () → None

Reset the statistic variables.

**reset\_buffer** (*keep\_statistics: bool = False*) → None

Reset the data buffer.

**reset\_env** () → None

Reset all of the environments.

**collect** (*n\_step*: Optional[int] = None, *n\_episode*: Optional[int] = None, *random*: bool = False, *render*: Optional[float] = None, *no\_grad*: bool = True) → Dict[str, Any]

Collect a specified number of step or episode.

To ensure unbiased sampling result with *n\_episode* option, this function will first collect *n\_episode* - *env\_num* episodes, then for the last *env\_num* episodes, they will be collected evenly from each env.

#### Parameters

- **n\_step** (*int*) – how many steps you want to collect.
- **n\_episode** (*int*) – how many episodes you want to collect.
- **random** (*bool*) – whether to use random policy for collecting data. Default to False.
- **render** (*float*) – the sleep time between rendering consecutive frames. Default to None (no rendering).
- **no\_grad** (*bool*) – whether to retain gradient in `policy.forward()`. Default to True (no gradient retaining).

---

**Note:** One and only one collection number specification is permitted, either *n\_step* or *n\_episode*.

---

#### Returns

A dict including the following keys

- *n/ep* collected number of episodes.
- *n/st* collected number of steps.
- *rews* array of episode reward over collected episodes.
- *lens* array of episode length over collected episodes.
- *idxs* array of episode start index in buffer over collected episodes.
- *rew* mean of episodic rewards.
- *len* mean of episodic lengths.
- *rew\_std* standard error of episodic rewards.
- *len\_std* standard error of episodic lengths.

### AsyncCollector

```
class tianshou.data.AsyncCollector (policy:          tianshou.policy.base.BasePolicy,      env:
                                     tianshou.env.venvs.BaseVectorEnv,      buffer:      Op-
                                     tional[tianshou.data.buffer.base.ReplayBuffer]      =
                                     None, preprocess_fn: Optional[Callable[[...], tian-
                                     shou.data.batch.Batch]] = None, exploration_noise: bool
                                     = False)
```

Bases: `tianshou.data.collector.Collector`

Async Collector handles async vector environment.

The arguments are exactly the same as `Collector`, please refer to `Collector` for more detailed explanation.

**reset\_env** () → None

Reset all of the environments.

**collect** (*n\_step*: Optional[int] = None, *n\_episode*: Optional[int] = None, *random*: bool = False, *render*: Optional[float] = None, *no\_grad*: bool = True) → Dict[str, Any]

Collect a specified number of step or episode with async env setting.

This function doesn't collect exactly *n\_step* or *n\_episode* number of transitions. Instead, in order to support async setting, it may collect more than given *n\_step* or *n\_episode* transitions and save into buffer.

#### Parameters

- **n\_step** (*int*) – how many steps you want to collect.
- **n\_episode** (*int*) – how many episodes you want to collect.
- **random** (*bool*) – whether to use random policy for collecting data. Default to False.
- **render** (*float*) – the sleep time between rendering consecutive frames. Default to None (no rendering).
- **no\_grad** (*bool*) – whether to retain gradient in policy.forward(). Default to True (no gradient retaining).

---

**Note:** One and only one collection number specification is permitted, either *n\_step* or *n\_episode*.

---

#### Returns

A dict including the following keys

- *n/ep* collected number of episodes.
- *n/st* collected number of steps.
- *rews* array of episode reward over collected episodes.
- *lens* array of episode length over collected episodes.
- *idxs* array of episode start index in buffer over collected episodes.
- *rew* mean of episodic rewards.
- *len* mean of episodic lengths.
- *rew\_std* standard error of episodic rewards.
- *len\_std* standard error of episodic lengths.

## 1.10 tianshou.env

### 1.10.1 VectorEnv

#### BaseVectorEnv

```
class tianshou.env.BaseVectorEnv (env_fns: List[Callable[], gym.core.Env], worker_fn:
    Callable[[Callable[], gym.core.Env], tian-
    shou.env.worker.base.EnvWorker], wait_num: Optional[int]
    = None, timeout: Optional[float] = None)
```

Bases: object

Base class for vectorized environments.

Usage:

```
env_num = 8
envs = DummyVectorEnv([lambda: gym.make(task) for _ in range(env_num)])
assert len(envs) == env_num
```

It accepts a list of environment generators. In other words, an environment generator `efn` of a specific task means that `efn()` returns the environment of the given task, for example, `gym.make(task)`.

All of the `VectorEnv` must inherit `BaseVectorEnv`. Here are some other usages:

```
envs.seed(2) # which is equal to the next line
envs.seed([2, 3, 4, 5, 6, 7, 8, 9]) # set specific seed for each env
obs = envs.reset() # reset all environments
obs = envs.reset([0, 5, 7]) # reset 3 specific environments
obs, rew, done, info = envs.step([1] * 8) # step synchronously
envs.render() # render all environments
envs.close() # close all environments
```

**Warning:** If you use your own environment, please make sure the `seed` method is set up properly, e.g.,

```
def seed(self, seed):
    np.random.seed(seed)
```

Otherwise, the outputs of these envs may be the same with each other.

### Parameters

- **env\_fns** – a list of callable envs, `env_fns[i]()` generates the *i*-th env.
- **worker\_fn** – a callable worker, `worker_fn(env_fns[i])` generates a worker which contains the *i*-th env.
- **wait\_num** (*int*) – use in asynchronous simulation if the time cost of `env.step` varies with time and synchronously waiting for all environments to finish a step is time-wasting. In that case, we can return when `wait_num` environments finish a step and keep on simulation in these environments. If `None`, asynchronous simulation is disabled; else,  $1 \leq \text{wait\_num} \leq \text{env\_num}$ .
- **timeout** (*float*) – use in asynchronous simulation same as above, in each vectorized step it only deal with those environments spending time within `timeout` seconds.

`__len__()` → *int*

Return `len(self)`, which is the number of environments.

**get\_env\_attr** (*key: str, id: Optional[Union[int, List[int], numpy.ndarray]] = None*) → *List[Any]*

Get an attribute from the underlying environments.

If `id` is an *int*, retrieve the attribute denoted by `key` from the environment underlying the worker at index `id`. The result is returned as a list with one element. Otherwise, retrieve the attribute for all workers at indices `id` and return a list that is ordered correspondingly to `id`.

### Parameters

- **key** (*str*) – The key of the desired attribute.
- **id** – Indices of the desired worker(s). Default to `None` for all `env_id`.

**Return list** The list of environment attributes.

**set\_env\_attr** (*key: str, value: Any, id: Optional[Union[int, List[int], numpy.ndarray]] = None*) → *None*  
 Set an attribute in the underlying environments.

If *id* is an int, set the attribute denoted by *key* from the environment underlying the worker at index *id* to *value*. Otherwise, set the attribute for all workers at indices *id*.

#### Parameters

- **key** (*str*) – The key of the desired attribute.
- **value** (*Any*) – The new value of the attribute.
- **id** – Indices of the desired worker(s). Default to *None* for all *env\_id*.

**reset** (*id: Optional[Union[int, List[int], numpy.ndarray]] = None*) → *numpy.ndarray*  
 Reset the state of some envs and return initial observations.

If *id* is *None*, reset the state of all the environments and return initial observations, otherwise reset the specific environments with the given *id*, either an int or a list.

**step** (*action: numpy.ndarray, id: Optional[Union[int, List[int], numpy.ndarray]] = None*) → *Tuple[numpy.ndarray, numpy.ndarray, numpy.ndarray, numpy.ndarray]*  
 Run one timestep of some environments' dynamics.

If *id* is *None*, run one timestep of all the environments' dynamics; otherwise run one timestep for some environments with given *id*, either an int or a list. When the end of episode is reached, you are responsible for calling `reset(id)` to reset this environment's state.

Accept a batch of action and return a tuple (*batch\_obs*, *batch\_rew*, *batch\_done*, *batch\_info*) in *numpy* format.

**Parameters** **action** (*numpy.ndarray*) – a batch of action provided by the agent.

#### Returns

A tuple including four items:

- *obs* a *numpy.ndarray*, the agent's observation of current environments
- *rew* a *numpy.ndarray*, the amount of rewards returned after previous actions
- *done* a *numpy.ndarray*, whether these episodes have ended, in which case further `step()` calls will return undefined results
- *info* a *numpy.ndarray*, contains auxiliary diagnostic information (helpful for debugging, and sometimes learning)

For the async simulation:

Provide the given action to the environments. The action sequence should correspond to the *id* argument, and the *id* argument should be a subset of the *env\_id* in the last returned *info* (initially they are *env\_ids* of all the environments). If action is *None*, fetch unfinished `step()` calls instead.

**seed** (*seed: Optional[Union[int, List[int]]] = None*) → *List[Optional[List[int]]]*  
 Set the seed for all environments.

Accept *None*, an int (which will extend *i* to [*i*, *i* + 1, *i* + 2, ...]) or a list.

**Returns** The list of seeds used in this env's random number generators. The first value in the list should be the "main" seed, or the value which a reproducer pass to "seed".

**render** (*\*\*kwargs: Any*) → *List[Any]*  
 Render all of the environments.

**close()** → None

Close all of the environments.

This function will be called only once (if not, it will be called during garbage collected). This way, `close` of all workers can be assured.

### DummyVectorEnv

**class** `tianshou.env.DummyVectorEnv` (*env\_fns: List[Callable[], gym.core.Env]*, *\*\*kwargs: Any*)  
Bases: `tianshou.env.venvs.BaseVectorEnv`

Dummy vectorized environment wrapper, implemented in for-loop.

**See also:**

Please refer to [BaseVectorEnv](#) for other APIs' usage.

### SubprocVectorEnv

**class** `tianshou.env.SubprocVectorEnv` (*env\_fns: List[Callable[], gym.core.Env]*, *\*\*kwargs: Any*)  
Bases: `tianshou.env.venvs.BaseVectorEnv`

Vectorized environment wrapper based on subprocess.

**See also:**

Please refer to [BaseVectorEnv](#) for other APIs' usage.

### ShmemVectorEnv

**class** `tianshou.env.ShmemVectorEnv` (*env\_fns: List[Callable[], gym.core.Env]*, *\*\*kwargs: Any*)  
Bases: `tianshou.env.venvs.BaseVectorEnv`

Optimized SubprocVectorEnv with shared buffers to exchange observations.

ShmemVectorEnv has exactly the same API as SubprocVectorEnv.

**See also:**

Please refer to [BaseVectorEnv](#) for other APIs' usage.

### RayVectorEnv

**class** `tianshou.env.RayVectorEnv` (*env\_fns: List[Callable[], gym.core.Env]*, *\*\*kwargs: Any*)  
Bases: `tianshou.env.venvs.BaseVectorEnv`

Vectorized environment wrapper based on ray.

This is a choice to run distributed environments in a cluster.

**See also:**

Please refer to [BaseVectorEnv](#) for other APIs' usage.



## 1.10.2 Worker

### EnvWorker

```
class tianshou.env.worker.EnvWorker (env_fn: Callable[], gym.core.Env)
    Bases: abc.ABC

    An abstract worker for an environment.

    abstract get_env_attr (key: str) → Any

    abstract set_env_attr (key: str, value: Any) → None

    send (action: Optional[numpy.ndarray]) → None
        Send action signal to low-level worker.

        When action is None, it indicates sending “reset” signal; otherwise it indicates “step” signal. The paired
        return value from “recv” function is determined by such kind of different signal.

    recv () → Union[Tuple[numpy.ndarray, numpy.ndarray, numpy.ndarray, numpy.ndarray],
        numpy.ndarray]
        Receive result from low-level worker.

        If the last “send” function sends a NULL action, it only returns a single observation; otherwise it returns a
        tuple of (obs, rew, done, info).

    reset () → numpy.ndarray

    step (action: numpy.ndarray) → Tuple[numpy.ndarray, numpy.ndarray, numpy.ndarray,
        numpy.ndarray]
        Perform one timestep of the environment’s dynamic.

        “send” and “recv” are coupled in sync simulation, so users only call “step” function. But they can be called
        separately in async simulation, i.e. someone calls “send” first, and calls “recv” later.

    static wait (workers: List[tianshou.env.worker.base.EnvWorker], wait_num: int, timeout: Op-
        tional[float] = None) → List[tianshou.env.worker.base.EnvWorker]
        Given a list of workers, return those ready ones.

    seed (seed: Optional[int] = None) → Optional[List[int]]

    abstract render (**kwargs: Any) → Any
        Render the environment.

    abstract close_env () → None

    close () → None
```

### DummyEnvWorker

```
class tianshou.env.worker.DummyEnvWorker (env_fn: Callable[], gym.core.Env)
    Bases: tianshou.env.worker.base.EnvWorker

    Dummy worker used in sequential vector environments.

    get_env_attr (key: str) → Any

    set_env_attr (key: str, value: Any) → None

    reset () → Any

    static wait (workers: List[tianshou.env.worker.dummy.DummyEnvWorker], wait_num: int, timeout:
        Optional[float] = None) → List[tianshou.env.worker.dummy.DummyEnvWorker]
        Given a list of workers, return those ready ones.
```

**send** (*action: Optional[numpy.ndarray]*) → None

Send action signal to low-level worker.

When action is None, it indicates sending “reset” signal; otherwise it indicates “step” signal. The paired return value from “recv” function is determined by such kind of different signal.

**seed** (*seed: Optional[int] = None*) → List[int]

**render** (*\*\*kwargs: Any*) → Any

Render the environment.

**close\_env** () → None

## SubprocEnvWorker

**class** tianshou.env.worker.**SubprocEnvWorker** (*env\_fn: Callable[], gym.core.Env,*  
*share\_memory: bool = False*)

Bases: tianshou.env.worker.base.EnvWorker

Subprocess worker used in SubprocVectorEnv and ShmemVectorEnv.

**get\_env\_attr** (*key: str*) → Any

**set\_env\_attr** (*key: str, value: Any*) → None

**static wait** (*workers: List[tianshou.env.worker.subproc.SubprocEnvWorker], wait\_num: int, time-*  
*out: Optional[float] = None*) → List[tianshou.env.worker.subproc.SubprocEnvWorker]

Given a list of workers, return those ready ones.

**send** (*action: Optional[numpy.ndarray]*) → None

Send action signal to low-level worker.

When action is None, it indicates sending “reset” signal; otherwise it indicates “step” signal. The paired return value from “recv” function is determined by such kind of different signal.

**recv** () → Union[Tuple[numpy.ndarray, numpy.ndarray, numpy.ndarray, numpy.ndarray],  
numpy.ndarray]

Receive result from low-level worker.

If the last “send” function sends a NULL action, it only returns a single observation; otherwise it returns a tuple of (obs, rew, done, info).

**seed** (*seed: Optional[int] = None*) → Optional[List[int]]

**render** (*\*\*kwargs: Any*) → Any

Render the environment.

**close\_env** () → None

## RayEnvWorker

**class** tianshou.env.worker.**RayEnvWorker** (*env\_fn: Callable[], gym.core.Env]*)

Bases: tianshou.env.worker.base.EnvWorker

Ray worker used in RayVectorEnv.

**get\_env\_attr** (*key: str*) → Any

**set\_env\_attr** (*key: str, value: Any*) → None

**reset** () → Any

**static wait** (*workers*: *List*[*tianshou.env.worker.ray.RayEnvWorker*], *wait\_num*: *int*, *timeout*: *Optional*[*float*] = *None*) → *List*[*tianshou.env.worker.ray.RayEnvWorker*]  
 Given a list of workers, return those ready ones.

**send** (*action*: *Optional*[*numpy.ndarray*]) → *None*  
 Send action signal to low-level worker.

When action is *None*, it indicates sending “reset” signal; otherwise it indicates “step” signal. The paired return value from “recv” function is determined by such kind of different signal.

**recv** () → *Union*[*Tuple*[*numpy.ndarray*, *numpy.ndarray*, *numpy.ndarray*, *numpy.ndarray*], *numpy.ndarray*]  
 Receive result from low-level worker.

If the last “send” function sends a *NULL* action, it only returns a single observation; otherwise it returns a tuple of (obs, rew, done, info).

**seed** (*seed*: *Optional*[*int*] = *None*) → *List*[*int*]

**render** (*\*\*kwargs*: *Any*) → *Any*  
 Render the environment.

**close\_env** () → *None*

## 1.11 tianshou.policy

### 1.11.1 Base

**class** *tianshou.policy.BasePolicy* (*observation\_space*: *Optional*[*gym.spaces.space.Space*] = *None*, *action\_space*: *Optional*[*gym.spaces.space.Space*] = *None*, *action\_scaling*: *bool* = *False*, *action\_bound\_method*: *str* = *"*, *lr\_scheduler*: *Optional*[*Union*[*torch.optim.lr\_scheduler.LambdaLR*, *tianshou.utils.lr\_scheduler.MultipleLRSchedulers*]] = *None*)

Bases: *abc.ABC*, *torch.nn.modules.module.Module*

The base class for any RL policy.

Tianshou aims to modularize RL algorithms. It comes into several classes of policies in Tianshou. All of the policy classes must inherit *BasePolicy*.

A policy class typically has the following parts:

- *\_\_init\_\_*(): initialize the policy, including coping the target network and so on;
- *forward*(): compute action with given observation;
- *process\_fn*(): pre-process data from the replay buffer (this function can interact with replay buffer);
- *learn*(): update policy with a given batch of data.
- *post\_process\_fn*(): update the replay buffer from the learning process (e.g., prioritized replay buffer needs to update the weight);
- *update*(): the main interface for training, i.e., *process\_fn* -> *learn* -> *post\_process\_fn*.

Most of the policy needs a neural network to predict the action and an optimizer to optimize the policy. The rules of self-defined networks are:

1. Input: observation “obs” (may be a *numpy.ndarray*, a *torch.Tensor*, a dict or any others), hidden state “state” (for RNN usage), and other information “info” provided by the environment.

2. Output: some “logits”, the next hidden state “state”, and the intermediate result during policy forwarding procedure “policy”. The “logits” could be a tuple instead of a `torch.Tensor`. It depends on how the policy process the network output. For example, in PPO, the return of the network might be `(mu, sigma)`, `state` for Gaussian policy. The “policy” can be a Batch of `torch.Tensor` or other things, which will be stored in the replay buffer, and can be accessed in the policy update process (e.g. in “`policy.learn()`”, the “`batch.policy`” is what you need).

Since `BasePolicy` inherits `torch.nn.Module`, you can use `BasePolicy` almost the same as `torch.nn.Module`, for instance, loading and saving the model:

```
torch.save(policy.state_dict(), "policy.pth")
policy.load_state_dict(torch.load("policy.pth"))
```

**set\_agent\_id** (*agent\_id*: int) → None  
Set self.agent\_id = agent\_id, for MARL.

**exploration\_noise** (*act*: Union[numpy.ndarray, tianshou.data.batch.Batch], *batch*: tianshou.data.batch.Batch) → Union[numpy.ndarray, tianshou.data.batch.Batch]  
Modify the action from policy.forward with exploration noise.

#### Parameters

- **act** – a data batch or `numpy.ndarray` which is the action taken by `policy.forward`.
- **batch** – the input batch for `policy.forward`, kept for advanced usage.

**Returns** action in the same form of input “act” but with added exploration noise.

**soft\_update** (*tgt*: torch.nn.modules.module.Module, *src*: torch.nn.modules.module.Module, *tau*: float) → None  
Softly update the parameters of target module towards the parameters of source module.

**abstract forward** (*batch*: tianshou.data.batch.Batch, *state*: Optional[Union[dict, tianshou.data.batch.Batch, numpy.ndarray]] = None, *\*\*kwargs*: Any) → tianshou.data.batch.Batch  
Compute action over the given batch data.

#### Returns

A `Batch` which MUST have the following keys:

- **act** an `numpy.ndarray` or a `torch.Tensor`, the action over given batch data.
- **state** a dict, an `numpy.ndarray` or a `torch.Tensor`, the internal state of the policy, None as default.

Other keys are user-defined. It depends on the algorithm. For example,

```
# some code
return Batch(logits=..., act=..., state=None, dist=...)
```

The keyword `policy` is reserved and the corresponding data will be stored into the replay buffer. For instance,

```
# some code
return Batch(..., policy=Batch(log_prob=dist.log_prob(act)))
# and in the sampled data batch, you can directly use
# batch.policy.log_prob to get your data.
```

**Note:** In continuous action space, you should do another step “`map_action`” to get the real action:

```
act = policy(batch).act # doesn't map to the target action range
act = policy.map_action(act, batch)
```

**map\_action** (*act*: Union[tianshou.data.batch.Batch, numpy.ndarray]) → Union[tianshou.data.batch.Batch, numpy.ndarray]  
 Map raw network output to action range in gym's env.action\_space.

This function is called in `collect()` and only affects action sending to env. Remapped action will not be stored in buffer and thus can be viewed as a part of env (a black box action transformation).

Action mapping includes 2 standard procedures: bounding and scaling. Bounding procedure expects original action range is  $(-\infty, \infty)$  and maps it to  $[-1, 1]$ , while scaling procedure expects original action range is  $(-1, 1)$  and maps it to  $[action\_space.low, action\_space.high]$ . Bounding procedure is applied first.

**Parameters** *act* – a data batch or numpy.ndarray which is the action taken by policy.forward.

**Returns** action in the same form of input “act” but remap to the target action space.

**map\_action\_inverse** (*act*: Union[tianshou.data.batch.Batch, List, numpy.ndarray]) → Union[tianshou.data.batch.Batch, List, numpy.ndarray]  
 Inverse operation to `map_action()`.

This function is called in `collect()` for random initial steps. It scales  $[action\_space.low, action\_space.high]$  to the value ranges of policy.forward.

**Parameters** *act* – a data batch, list or numpy.ndarray which is the action taken by `gym.spaces.Box.sample()`.

**Returns** action remapped.

**process\_fn** (*batch*: tianshou.data.batch.Batch, *buffer*: tianshou.data.buffer.base.ReplayBuffer, *indices*: numpy.ndarray) → tianshou.data.batch.Batch  
 Pre-process the data from the provided replay buffer.

Used in `update()`. Check out `policy.process_fn` for more information.

**abstract learn** (*batch*: tianshou.data.batch.Batch, *\*\*kwargs*: Any) → Dict[str, Any]  
 Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**post\_process\_fn** (*batch*: tianshou.data.batch.Batch, *buffer*: tianshou.data.buffer.base.ReplayBuffer, *indices*: numpy.ndarray) → None  
 Post-process the data from the provided replay buffer.

Typical usage is to update the sampling weight in prioritized experience replay. Used in `update()`.

**update** (*sample\_size*: int, *buffer*: Optional[tianshou.data.buffer.base.ReplayBuffer], *\*\*kwargs*: Any) → Dict[str, Any]

Update the policy network and replay buffer.

It includes 3 function steps: `process_fn`, `learn`, and `post_process_fn`. In addition, this function will change the value of `self.updating`: it will be False before this function and will be True when executing `update()`. Please refer to [States for policy](#) for more detailed explanation.

#### Parameters

- **sample\_size** (int) – 0 means it will extract all the data from the buffer, otherwise it will sample a batch with given sample\_size.
- **buffer** (ReplayBuffer) – the corresponding replay buffer.

**Returns** A dict, including the data needed to be logged (e.g., loss) from `policy.learn()`.

**static value\_mask** (*buffer*: tianshou.data.buffer.base.ReplayBuffer, *indices*: numpy.ndarray) → numpy.ndarray

Value mask determines whether the `obs_next` of `buffer[indices]` is valid.

For instance, usually “obs\_next” after “done” flag is considered to be invalid, and its q/advantage value can provide meaningless (even misleading) information, and should be set to 0 by hand. But if “done” flag is generated because timelimit of game length (info[“TimeLimit.truncated”] is set to True in gym’s settings), “obs\_next” will instead be valid. Value mask is typically used for assisting in calculating the correct q/advantage value.

#### Parameters

- **buffer** (ReplayBuffer) – the corresponding replay buffer.
- **indices** (numpy.ndarray) – indices of replay buffer whose “obs\_next” will be judged.

**Returns** A bool type numpy.ndarray in the same shape with indices. “True” means “obs\_next” of that `buffer[indices]` is valid.

**static compute\_episodic\_return** (*batch*: tianshou.data.batch.Batch, *buffer*: tianshou.data.buffer.base.ReplayBuffer, *indices*: numpy.ndarray, *v\_s*: Optional[Union[numpy.ndarray, torch.Tensor]] = None, *v\_s*: Optional[Union[numpy.ndarray, torch.Tensor]] = None, *gamma*: float = 0.99, *gae\_lambda*: float = 0.95) → Tuple[numpy.ndarray, numpy.ndarray]

Compute returns over given batch.

Use Implementation of Generalized Advantage Estimator (arXiv:1506.02438) to calculate q/advantage value of given batch.

#### Parameters

- **batch** (Batch) – a data batch which contains several episodes of data in sequential order. Mind that the end of each finished episode of batch should be marked by done flag, unfinished (or collecting) episodes will be recognized by `buffer.unfinished_index()`.
- **indices** (numpy.ndarray) – tell batch’s location in buffer, batch is equal to `buffer[indices]`.
- **v\_s** (np.ndarray) – the value function of all next states  $V(s')$ .
- **gamma** (float) – the discount factor, should be in [0, 1]. Default to 0.99.
- **gae\_lambda** (float) – the parameter for Generalized Advantage Estimation, should be in [0, 1]. Default to 0.95.

**Returns** two numpy arrays (returns, advantage) with each shape (bsz, ).

**static compute\_nstep\_return** (*batch*: *tianshou.data.batch.Batch*, *buffer*: *tianshou.data.buffer.base.ReplayBuffer*, *indice*: *numpy.ndarray*, *target\_q\_fn*: *Callable[[tianshou.data.buffer.base.ReplayBuffer, numpy.ndarray], torch.Tensor]*, *gamma*: *float* = 0.99, *n\_step*: *int* = 1, *rew\_norm*: *bool* = False) → *tianshou.data.batch.Batch*

Compute n-step return for Q-learning targets.

$$G_t = \sum_{i=t}^{t+n-1} \gamma^{i-t} (1 - d_i) r_i + \gamma^n (1 - d_{t+n}) Q_{\text{target}}(s_{t+n})$$

where  $\gamma$  is the discount factor,  $\gamma \in [0, 1]$ ,  $d_t$  is the done flag of step  $t$ .

#### Parameters

- **batch** (*Batch*) – a data batch, which is equal to `buffer[indice]`.
- **buffer** (*ReplayBuffer*) – the data buffer.
- **target\_q\_fn** (*function*) – a function which compute target Q value of “obs\_next” given data buffer and wanted indices.
- **gamma** (*float*) – the discount factor, should be in  $[0, 1]$ . Default to 0.99.
- **n\_step** (*int*) – the number of estimation step, should be an int greater than 0. Default to 1.
- **rew\_norm** (*bool*) – normalize the reward to Normal(0, 1), Default to False.

**Returns** a *Batch*. The result will be stored in `batch.returns` as a *torch.Tensor* with the same shape as `target_q_fn`’s return tensor.

**training:** `bool`

**class** `tianshou.policy.RandomPolicy` (*observation\_space*: *Optional[gym.spaces.space.Space]* = None, *action\_space*: *Optional[gym.spaces.space.Space]* = None, *action\_scaling*: *bool* = False, *action\_bound\_method*: *str* = "", *lr\_scheduler*: *Optional[Union[torch.optim.lr\_scheduler.LambdaLR, tianshou.utils.lr\_scheduler.MultipleLRSchedulers]]* = None)

Bases: `tianshou.policy.base.BasePolicy`

A random agent used in multi-agent learning.

It randomly chooses an action from the legal action.

**forward** (*batch*: *tianshou.data.batch.Batch*, *state*: *Optional[Union[dict, tianshou.data.batch.Batch, numpy.ndarray]]* = None, *\*\*kwargs*: Any) → *tianshou.data.batch.Batch*

Compute the random action over the given batch data.

The input should contain a mask in `batch.obs`, with “True” to be available and “False” to be unavailable. For example, `batch.obs.mask == np.array([[False, True, False]])` means with batch size 1, action “1” is available but action “0” and “2” are unavailable.

**Returns** A *Batch* with “act” key, containing the random action.

**See also:**

Please refer to `forward()` for more detailed explanation.

**learn** (*batch*: *tianshou.data.batch.Batch*, *\*\*kwargs*: Any) → Dict[str, float]

Since a random agent learns nothing, it returns an empty dict.

`training: bool`

### 1.11.2 Model-free

#### DQN Family

```
class tianshou.policy.DQNPoly (model: torch.nn.modules.module.Module, optim:
    torch.optim.optimizer.Optimizer, discount_factor: float =
    0.99, estimation_step: int = 1, target_update_freq: int = 0,
    reward_normalization: bool = False, is_double: bool = True,
    **kwargs: Any)
```

Bases: `tianshou.policy.base.BasePolicy`

Implementation of Deep Q Network. arXiv:1312.5602.

Implementation of Double Q-Learning. arXiv:1509.06461.

Implementation of Dueling DQN. arXiv:1511.06581 (the dueling DQN is implemented in the network side, not here).

##### Parameters

- **model** (`torch.nn.Module`) – a model following the rules in `BasePolicy`. (s -> logits)
- **optim** (`torch.optim.Optimizer`) – a torch.optim for optimizing the model.
- **discount\_factor** (`float`) – in [0, 1].
- **estimation\_step** (`int`) – the number of steps to look ahead. Default to 1.
- **target\_update\_freq** (`int`) – the target network update frequency (0 if you do not use the target network). Default to 0.
- **reward\_normalization** (`bool`) – normalize the reward to Normal(0, 1). Default to False.
- **is\_double** (`bool`) – use double dqn. Default to True.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).

See also:

Please refer to `BasePolicy` for more detailed explanation.

**set\_eps** (`eps: float`) → None

Set the eps for epsilon-greedy exploration.

**train** (`mode: bool = True`) → `tianshou.policy.modelfree.dqn.DQNPoly`

Set the module in training mode, except for the target network.

**sync\_weight** () → None

Synchronize the weight for the target network.

**process\_fn** (`batch: tianshou.data.batch.Batch`, `buffer: tianshou.data.buffer.base.ReplayBuffer`, `indices: numpy.ndarray`) → `tianshou.data.batch.Batch`

Compute the n-step return for Q-learning targets.

More details can be found at `compute_nstep_return()`.

**compute\_q\_value** (`logits: torch.Tensor`, `mask: Optional[numpy.ndarray]`) → `torch.Tensor`

Compute the q value based on the network's raw output and action mask.



**forward** (*batch*: *tianshou.data.batch.Batch*, *state*: *Optional[Union[dict, tianshou.data.batch.Batch, numpy.ndarray]] = None*, *model*: *str = 'model'*, *input*: *str = 'obs'*, *\*\*kwargs*: *Any*) → *tianshou.data.batch.Batch*  
 Compute action over the given batch data.

If you need to mask the action, please add a “mask” into *batch.obs*, for example, if we have an environment that has “0/1/2” three actions:

```
batch == Batch(
    obs=Batch(
        obs="original obs, with batch_size=1 for demonstration",
        mask=np.array([[False, True, False]]),
        # action 1 is available
        # action 0 and 2 are unavailable
    ),
    ...
)
```

**Parameters** *eps* (*float*) – in [0, 1], for epsilon-greedy exploration method.

**Returns**

A *Batch* which has 3 keys:

- *act* the action.
- *logits* the network’s raw output.
- *state* the hidden state.

**See also:**

Please refer to *forward()* for more detailed explanation.

**learn** (*batch*: *tianshou.data.batch.Batch*, *\*\*kwargs*: *Any*) → *Dict[str, float]*

Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by *self.training* and *self.updating*. Please refer to *States for policy* for more detailed explanation.

**Warning:** If you use *torch.distributions.Normal* and *torch.distributions.Categorical* to calculate the *log\_prob*, please be careful about the shape: *Categorical* distribution gives “[batch\_size]” shape while *Normal* distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**exploration\_noise** (*act*: *Union[numpy.ndarray, tianshou.data.batch.Batch]*, *batch*: *tianshou.data.batch.Batch*) → *Union[numpy.ndarray, tianshou.data.batch.Batch]*

Modify the action from *policy.forward* with exploration noise.

**Parameters**

- *act* – a data batch or *numpy.ndarray* which is the action taken by *policy.forward*.
- *batch* – the input batch for *policy.forward*, kept for advanced usage.

**Returns** action in the same form of input “act” but with added exploration noise.

**training:** bool

```
class tianshou.policy.C51Policy(model: torch.nn.modules.module.Module, optim:
    torch.optim.optimizer.Optimizer, discount_factor: float =
    0.99, num_atoms: int = 51, v_min: float = -10.0, v_max: float
    = 10.0, estimation_step: int = 1, target_update_freq: int = 0,
    reward_normalization: bool = False, **kwargs: Any)
```

Bases: `tianshou.policy.modelfree.dqn.DQNPoly`

Implementation of Categorical Deep Q-Network. arXiv:1707.06887.

### Parameters

- **model** (`torch.nn.Module`) – a model following the rules in `BasePolicy`. (s -> logits)
- **optim** (`torch.optim.Optimizer`) – a torch.optim for optimizing the model.
- **discount\_factor** (`float`) – in [0, 1].
- **num\_atoms** (`int`) – the number of atoms in the support set of the value distribution. Default to 51.
- **v\_min** (`float`) – the value of the smallest atom in the support set. Default to -10.0.
- **v\_max** (`float`) – the value of the largest atom in the support set. Default to 10.0.
- **estimation\_step** (`int`) – the number of steps to look ahead. Default to 1.
- **target\_update\_freq** (`int`) – the target network update frequency (0 if you do not use the target network). Default to 0.
- **reward\_normalization** (`bool`) – normalize the reward to Normal(0, 1). Default to False.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).

See also:

Please refer to `DQNPoly` for more detailed explanation.

**compute\_q\_value** (logits: `torch.Tensor`, mask: `Optional[numpy.ndarray]`) → `torch.Tensor`  
Compute the q value based on the network’s raw output and action mask.

**learn** (batch: `tianshou.data.batch.Batch`, `**kwargs: Any`) → `Dict[str, float]`  
Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

---

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the log\_prob, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** bool

```
class tianshou.policy.RainbowPolicy(model: torch.nn.modules.module.Module, optim:
    torch.optim.optimizer.Optimizer, discount_factor: float
    = 0.99, num_atoms: int = 51, v_min: float = - 10.0,
    v_max: float = 10.0, estimation_step: int = 1, tar-
    get_update_freq: int = 0, reward_normalization: bool =
    False, **kwargs: Any)
```

Bases: `tianshou.policy.modelfree.c51.C51Policy`

Implementation of Rainbow DQN. arXiv:1710.02298.

#### Parameters

- **model** (`torch.nn.Module`) – a model following the rules in `BasePolicy`. (s -> logits)
- **optim** (`torch.optim.Optimizer`) – a torch.optim for optimizing the model.
- **discount\_factor** (`float`) – in [0, 1].
- **num\_atoms** (`int`) – the number of atoms in the support set of the value distribution. Default to 51.
- **v\_min** (`float`) – the value of the smallest atom in the support set. Default to -10.0.
- **v\_max** (`float`) – the value of the largest atom in the support set. Default to 10.0.
- **estimation\_step** (`int`) – the number of steps to look ahead. Default to 1.
- **target\_update\_freq** (`int`) – the target network update frequency (0 if you do not use the target network). Default to 0.
- **reward\_normalization** (`bool`) – normalize the reward to Normal(0, 1). Default to False.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).

See also:

Please refer to `C51Policy` for more detailed explanation.

**learn** (`batch: tianshou.data.batch.Batch`, `**kwargs: Any`) → Dict[str, float]

Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

---

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the log\_prob, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** bool

```
class tianshou.policy.QRDQNPoly (model: torch.nn.modules.module.Module, optim:
    torch.optim.optimizer.Optimizer, discount_factor: float
    = 0.99, num_quantiles: int = 200, estimation_step: int = 1,
    target_update_freq: int = 0, reward_normalization: bool =
    False, **kwargs: Any)
```

Bases: `tianshou.policy.modelfree.dqn.DQNPoly`

Implementation of Quantile Regression Deep Q-Network. arXiv:1710.10044.

#### Parameters

- **model** (*torch.nn.Module*) – a model following the rules in *BasePolicy*. (s -> logits)
- **optim** (*torch.optim.Optimizer*) – a torch.optim for optimizing the model.
- **discount\_factor** (*float*) – in [0, 1].
- **num\_quantiles** (*int*) – the number of quantile midpoints in the inverse cumulative distribution function of the value. Default to 200.
- **estimation\_step** (*int*) – the number of steps to look ahead. Default to 1.
- **target\_update\_freq** (*int*) – the target network update frequency (0 if you do not use the target network).
- **reward\_normalization** (*bool*) – normalize the reward to Normal(0, 1). Default to False.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).

See also:

Please refer to *DQNPoly* for more detailed explanation.

**compute\_q\_value** (*logits: torch.Tensor, mask: Optional[numpy.ndarray]*) → *torch.Tensor*  
Compute the q value based on the network’s raw output and action mask.

**learn** (*batch: tianshou.data.batch.Batch, \*\*kwargs: Any*) → *Dict[str, float]*  
Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to *States for policy* for more detailed explanation.

---

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** `bool`

```
class tianshou.policy.IQNPolicy(model: torch.nn.modules.module.Module, optim:
    torch.optim.optimizer.Optimizer, discount_factor: float =
    0.99, sample_size: int = 32, online_sample_size: int = 8,
    target_sample_size: int = 8, estimation_step: int = 1, tar-
    get_update_freq: int = 0, reward_normalization: bool = False,
    **kwargs: Any)
```

Bases: `tianshou.policy.modelfree.qrdqn.QRDQNPolicy`

Implementation of Implicit Quantile Network. arXiv:1806.06923.

### Parameters

- **model** (`torch.nn.Module`) – a model following the rules in `BasePolicy`. (s -> logits)
- **optim** (`torch.optim.Optimizer`) – a torch.optim for optimizing the model.
- **discount\_factor** (`float`) – in [0, 1].
- **sample\_size** (`int`) – the number of samples for policy evaluation. Default to 32.
- **online\_sample\_size** (`int`) – the number of samples for online model in training. Default to 8.
- **target\_sample\_size** (`int`) – the number of samples for target model in training. Default to 8.
- **estimation\_step** (`int`) – the number of steps to look ahead. Default to 1.
- **target\_update\_freq** (`int`) – the target network update frequency (0 if you do not use the target network).
- **reward\_normalization** (`bool`) – normalize the reward to Normal(0, 1). Default to False.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).

See also:

Please refer to `QRDQNPolicy` for more detailed explanation.

**forward** (`batch: tianshou.data.batch.Batch`, `state: Optional[Union[dict, tianshou.data.batch.Batch, numpy.ndarray]] = None`, `model: str = 'model'`, `input: str = 'obs'`, `**kwargs: Any`) → `tianshou.data.batch.Batch`  
 Compute action over the given batch data.

If you need to mask the action, please add a “mask” into batch.obs, for example, if we have an environment that has “0/1/2” three actions:

```
batch == Batch(
  obs=Batch(
    obs="original obs, with batch_size=1 for demonstration",
    mask=np.array([[False, True, False]]),
    # action 1 is available
    # action 0 and 2 are unavailable
  ),
  ...
)
```

**Parameters** **eps** (`float`) – in [0, 1], for epsilon-greedy exploration method.

**Returns**

A *Batch* which has 3 keys:

- `act` the action.
- `logits` the network's raw output.
- `state` the hidden state.

**See also:**

Please refer to `forward()` for more detailed explanation.

**learn** (*batch*: *tianshou.data.batch.Batch*, *\*\*kwargs*: *Any*) → Dict[str, float]

Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to *States for policy* for more detailed explanation.

---

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** `bool`

```
class tianshou.policy.FQFPolicy(model: tianshou.utils.net.discrete.FullQuantileFunction,
                                optim: torch.optim.optimizer.Optimizer, fraction_model:
                                tianshou.utils.net.discrete.FractionProposalNetwork, frac-
                                tion_optim: torch.optim.optimizer.Optimizer, discount_factor:
                                float = 0.99, num_fractions: int = 32, ent_coef: float =
                                0.0, estimation_step: int = 1, target_update_freq: int = 0,
                                reward_normalization: bool = False, **kwargs: Any)
```

Bases: `tianshou.policy.modelfree.qrdqn.QRDQNPoly`

Implementation of Fully-parameterized Quantile Function. arXiv:1911.02140.

**Parameters**

- **model** (*torch.nn.Module*) – a model following the rules in *BasePolicy*. (*s* → logits)
- **optim** (*torch.optim.Optimizer*) – a torch.optim for optimizing the model.
- **fraction\_model** (*FractionProposalNetwork*) – a FractionProposalNetwork for proposing fractions/quantiles given state.
- **fraction\_optim** (*torch.optim.Optimizer*) – a torch.optim for optimizing the fraction model above.
- **discount\_factor** (*float*) – in [0, 1].
- **num\_fractions** (*int*) – the number of fractions to use. Default to 32.
- **ent\_coef** (*float*) – the coefficient for entropy loss. Default to 0.

- **estimation\_step** (*int*) – the number of steps to look ahead. Default to 1.
- **target\_update\_freq** (*int*) – the target network update frequency (0 if you do not use the target network).
- **reward\_normalization** (*bool*) – normalize the reward to Normal(0, 1). Default to False.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).

See also:

Please refer to [QRDQNPolicy](#) for more detailed explanation.

**forward** (*batch*: *tianshou.data.batch.Batch*, *state*: *Optional[Union[dict, tianshou.data.batch.Batch, numpy.ndarray]] = None*, *model*: *str = 'model'*, *input*: *str = 'obs'*, *fractions*: *Optional[tianshou.data.batch.Batch] = None*, *\*\*kwargs*: *Any*) → *tianshou.data.batch.Batch*  
 Compute action over the given batch data.

If you need to mask the action, please add a “mask” into batch.obs, for example, if we have an environment that has “0/1/2” three actions:

```
batch == Batch(
    obs=Batch(
        obs="original obs, with batch_size=1 for demonstration",
        mask=np.array([[False, True, False]]),
        # action 1 is available
        # action 0 and 2 are unavailable
    ),
    ...
)
```

**Parameters** **eps** (*float*) – in [0, 1], for epsilon-greedy exploration method.

**Returns**

A *Batch* which has 3 keys:

- **act** the action.
- **logits** the network’s raw output.
- **state** the hidden state.

See also:

Please refer to [forward\(\)](#) for more detailed explanation.

**learn** (*batch*: *tianshou.data.batch.Batch*, *\*\*kwargs*: *Any*) → *Dict[str, float]*  
 Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

---

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** `bool`

## On-policy

```
class tianshou.policy.PGPolicy(model: torch.nn.modules.module.Module, op-  
                                tim: torch.optim.optimizer.Optimizer, dist_fn:  
                                Type[torch.distributions.distribution.Distribution], dis-  
                                count_factor: float = 0.99, reward_normalization: bool =  
                                False, action_scaling: bool = True, action_bound_method: str  
                                = 'clip', deterministic_eval: bool = False, **kwargs: Any)
```

Bases: `tianshou.policy.base.BasePolicy`

Implementation of REINFORCE algorithm.

### Parameters

- **model** (`torch.nn.Module`) – a model following the rules in `BasePolicy`. (s -> logits)
- **optim** (`torch.optim.Optimizer`) – a torch.optim for optimizing the model.
- **dist\_fn** (`Type[torch.distributions.Distribution]`) – distribution class for computing the action.
- **discount\_factor** (`float`) – in [0, 1]. Default to 0.99.
- **action\_scaling** (`bool`) – whether to map actions from range [-1, 1] to range [action\_spaces.low, action\_spaces.high]. Default to True.
- **action\_bound\_method** (`str`) – method to bound action to range [-1, 1], can be either “clip” (for simply clipping the action), “tanh” (for applying tanh squashing) for now, or empty string for no bounding. Default to “clip”.
- **action\_space** (`Optional[gym.Space]`) – env’s action space, mandatory if you want to use option “action\_scaling” or “action\_bound\_method”. Default to None.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).
- **deterministic\_eval** (`bool`) – whether to use deterministic action instead of stochastic action sampled by the policy. Default to False.

See also:

Please refer to `BasePolicy` for more detailed explanation.

**process\_fn** (*batch*: `tianshou.data.batch.Batch`, *buffer*: `tianshou.data.buffer.base.ReplayBuffer`, *indices*: `numpy.ndarray`) → `tianshou.data.batch.Batch`  
Compute the discounted returns for each transition.

$$G_t = \sum_{i=t}^T \gamma^{i-t} r_i$$

where  $T$  is the terminal time step,  $\gamma$  is the discount factor,  $\gamma \in [0, 1]$ .



**forward** (*batch*: *tianshou.data.batch.Batch*, *state*: *Optional[Union[dict, tianshou.data.batch.Batch, numpy.ndarray]] = None*, *\*\*kwargs*: *Any*) → *tianshou.data.batch.Batch*  
 Compute action over the given batch data.

#### Returns

A *Batch* which has 4 keys:

- *act* the action.
- *logits* the network's raw output.
- *dist* the action distribution.
- *state* the hidden state.

#### See also:

Please refer to *forward()* for more detailed explanation.

**learn** (*batch*: *tianshou.data.batch.Batch*, *batch\_size*: *int*, *repeat*: *int*, *\*\*kwargs*: *Any*) → *Dict[str, List[float]]*  
 Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by *self.training* and *self.updating*. Please refer to *States for policy* for more detailed explanation.

---

**Warning:** If you use *torch.distributions.Normal* and *torch.distributions.Categorical* to calculate the *log\_prob*, please be careful about the shape: *Categorical* distribution gives "[batch\_size]" shape while *Normal* distribution gives "[batch\_size, 1]" shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** *bool*

```
class tianshou.policy.NPGPolicy (actor: torch.nn.modules.module.Module,
                                critic: torch.nn.modules.module.Module,      op-
                                tim: torch.optim.optimizer.Optimizer,         tim
                                dist_fn: Type[torch.distributions.distribution.Distribution], dist_fn
                                advantage_normalization: bool = True, optim_critic_iters: int = 5,
                                actor_step_size: float = 0.5, **kwargs: Any)
```

Bases: *tianshou.policy.modelfree.a2c.A2CPolicy*

Implementation of Natural Policy Gradient.

<https://proceedings.neurips.cc/paper/2001/file/4b86abe48d358ecf194c56c69108433e-Paper.pdf>

#### Parameters

- **actor** (*torch.nn.Module*) – the actor network following the rules in *BasePolicy*. (s -> logits)
- **critic** (*torch.nn.Module*) – the critic network. (s -> V(s))
- **optim** (*torch.optim.Optimizer*) – the optimizer for actor and critic network.
- **dist\_fn** (*Type[torch.distributions.Distribution]*) – distribution class for computing the action.

- **advantage\_normalization** (*bool*) – whether to do per mini-batch advantage normalization. Default to True.
- **optim\_critic\_iters** (*int*) – Number of times to optimize critic network per update. Default to 5.
- **gae\_lambda** (*float*) – in  $[0, 1]$ , param for Generalized Advantage Estimation. Default to 0.95.
- **reward\_normalization** (*bool*) – normalize estimated values to have std close to 1. Default to False.
- **max\_batchsize** (*int*) – the maximum size of the batch when computing GAE, depends on the size of available memory and the memory cost of the model; should be as large as possible within the memory constraint. Default to 256.
- **action\_scaling** (*bool*) – whether to map actions from range  $[-1, 1]$  to range  $[\text{action\_spaces.low}, \text{action\_spaces.high}]$ . Default to True.
- **action\_bound\_method** (*str*) – method to bound action to range  $[-1, 1]$ , can be either “clip” (for simply clipping the action), “tanh” (for applying tanh squashing) for now, or empty string for no bounding. Default to “clip”.
- **action\_space** (*Optional[gym.Space]*) – env’s action space, mandatory if you want to use option “action\_scaling” or “action\_bound\_method”. Default to None.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).
- **deterministic\_eval** (*bool*) – whether to use deterministic action instead of stochastic action sampled by the policy. Default to False.

**process\_fn** (*batch: tianshou.data.batch.Batch, buffer: tianshou.data.buffer.base.ReplayBuffer, indices: numpy.ndarray*)  $\rightarrow$  *tianshou.data.batch.Batch*  
 Compute the discounted returns for each transition.

$$G_t = \sum_{i=t}^T \gamma^{i-t} r_i$$

where  $T$  is the terminal time step,  $\gamma$  is the discount factor,  $\gamma \in [0, 1]$ .

**learn** (*batch: tianshou.data.batch.Batch, batch\_size: int, repeat: int, \*\*kwargs: Any*)  $\rightarrow$  *Dict[str, List[float]]*  
 Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

---

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** *bool*

```
class tianshou.policy.A2CPolicy(actor: torch.nn.modules.module.Module,
                                critic: torch.nn.modules.module.Module, op-
                                tim: torch.optim.optimizer.Optimizer, dist_fn:
                                Type[torch.distributions.distribution.Distribution], vf_coef:
                                float = 0.5, ent_coef: float = 0.01, max_grad_norm:
                                Optional[float] = None, gae_lambda: float = 0.95,
                                max_batchsize: int = 256, **kwargs: Any)
```

Bases: `tianshou.policy.modelfree.pg.PGPolicy`

Implementation of Synchronous Advantage Actor-Critic. arXiv:1602.01783.

### Parameters

- **actor** (`torch.nn.Module`) – the actor network following the rules in `BasePolicy`. (s -> logits)
- **critic** (`torch.nn.Module`) – the critic network. (s -> V(s))
- **optim** (`torch.optim.Optimizer`) – the optimizer for actor and critic network.
- **dist\_fn** (`Type[torch.distributions.Distribution]`) – distribution class for computing the action.
- **discount\_factor** (`float`) – in [0, 1]. Default to 0.99.
- **vf\_coef** (`float`) – weight for value loss. Default to 0.5.
- **ent\_coef** (`float`) – weight for entropy loss. Default to 0.01.
- **max\_grad\_norm** (`float`) – clipping gradients in back propagation. Default to None.
- **gae\_lambda** (`float`) – in [0, 1], param for Generalized Advantage Estimation. Default to 0.95.
- **reward\_normalization** (`bool`) – normalize estimated values to have std close to 1. Default to False.
- **max\_batchsize** (`int`) – the maximum size of the batch when computing GAE, depends on the size of available memory and the memory cost of the model; should be as large as possible within the memory constraint. Default to 256.
- **action\_scaling** (`bool`) – whether to map actions from range [-1, 1] to range [action\_spaces.low, action\_spaces.high]. Default to True.
- **action\_bound\_method** (`str`) – method to bound action to range [-1, 1], can be either “clip” (for simply clipping the action), “tanh” (for applying tanh squashing) for now, or empty string for no bounding. Default to “clip”.
- **action\_space** (`Optional[gym.Space]`) – env’s action space, mandatory if you want to use option “action\_scaling” or “action\_bound\_method”. Default to None.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).
- **deterministic\_eval** (`bool`) – whether to use deterministic action instead of stochastic action sampled by the policy. Default to False.

See also:

Please refer to `BasePolicy` for more detailed explanation.

**process\_fn** (*batch*: *tianshou.data.batch.Batch*, *buffer*: *tianshou.data.buffer.base.ReplayBuffer*, *indices*: *numpy.ndarray*) → *tianshou.data.batch.Batch*  
 Compute the discounted returns for each transition.

$$G_t = \sum_{i=t}^T \gamma^{i-t} r_i$$

where  $T$  is the terminal time step,  $\gamma$  is the discount factor,  $\gamma \in [0, 1]$ .

**learn** (*batch*: *tianshou.data.batch.Batch*, *batch\_size*: *int*, *repeat*: *int*, *\*\*kwargs*: *Any*) → Dict[str, List[float]]  
 Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** `bool`

**class** `tianshou.policy.TRPOPoly` (*actor*: *torch.nn.modules.module.Module*,  
*critic*: *torch.nn.modules.module.Module*, *optim*: *torch.optim.optimizer.Optimizer*, *dist\_fn*:  
*Type[torch.distributions.distribution.Distribution]*, *max\_kl*:  
*float = 0.01*, *backtrack\_coeff*: *float = 0.8*, *max\_backtracks*:  
*int = 10*, *\*\*kwargs*: *Any*)

Bases: `tianshou.policy.modelfree.npg.NPGPolicy`

Implementation of Trust Region Policy Optimization. arXiv:1502.05477.

#### Parameters

- **actor** (*torch.nn.Module*) – the actor network following the rules in [BasePolicy](#). (s → logits)
- **critic** (*torch.nn.Module*) – the critic network. (s → V(s))
- **optim** (*torch.optim.Optimizer*) – the optimizer for actor and critic network.
- **dist\_fn** (*Type[torch.distributions.Distribution]*) – distribution class for computing the action.
- **advantage\_normalization** (*bool*) – whether to do per mini-batch advantage normalization. Default to True.
- **optim\_critic\_iters** (*int*) – Number of times to optimize critic network per update. Default to 5.
- **max\_kl** (*int*) – max kl-divergence used to constrain each actor network update. Default to 0.01.

- **backtrack\_coeff** (*float*) – Coefficient to be multiplied by step size when constraints are not met. Default to 0.8.
- **max\_backtracks** (*int*) – Max number of backtracking times in linesearch. Default to 10.
- **gae\_lambda** (*float*) – in [0, 1], param for Generalized Advantage Estimation. Default to 0.95.
- **reward\_normalization** (*bool*) – normalize estimated values to have std close to 1. Default to False.
- **max\_batchsize** (*int*) – the maximum size of the batch when computing GAE, depends on the size of available memory and the memory cost of the model; should be as large as possible within the memory constraint. Default to 256.
- **action\_scaling** (*bool*) – whether to map actions from range [-1, 1] to range [action\_spaces.low, action\_spaces.high]. Default to True.
- **action\_bound\_method** (*str*) – method to bound action to range [-1, 1], can be either “clip” (for simply clipping the action), “tanh” (for applying tanh squashing) for now, or empty string for no bounding. Default to “clip”.
- **action\_space** (*Optional[gym.Space]*) – env’s action space, mandatory if you want to use option “action\_scaling” or “action\_bound\_method”. Default to None.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).
- **deterministic\_eval** (*bool*) – whether to use deterministic action instead of stochastic action sampled by the policy. Default to False.

**learn** (*batch: tianshou.data.batch.Batch, batch\_size: int, repeat: int, \*\*kwargs: Any*) → Dict[str, List[float]]  
Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the log\_prob, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** bool

```
class tianshou.policy.PPOPoly (actor: torch.nn.modules.module.Module,
                                critic: torch.nn.modules.module.Module, op-
                                tim: torch.optim.optimizer.Optimizer, dist_fn:
                                Type[torch.distributions.distribution.Distribution], eps_clip:
                                float = 0.2, dual_clip: Optional[float] = None, value_clip:
                                bool = False, advantage_normalization: bool = True,
                                recompute_advantage: bool = False, **kwargs: Any)
Bases: tianshou.policy.modelfree.a2c.A2CPolicy
```

Implementation of Proximal Policy Optimization. arXiv:1707.06347.

### Parameters

- **actor** (*torch.nn.Module*) – the actor network following the rules in *BasePolicy*. (s -> logits)
- **critic** (*torch.nn.Module*) – the critic network. (s -> V(s))
- **optim** (*torch.optim.Optimizer*) – the optimizer for actor and critic network.
- **dist\_fn** (*Type[torch.distributions.Distribution]*) – distribution class for computing the action.
- **discount\_factor** (*float*) – in  $[0, 1]$ . Default to 0.99.
- **eps\_clip** (*float*) –  $\epsilon$  in  $L_{CLIP}$  in the original paper. Default to 0.2.
- **dual\_clip** (*float*) – a parameter  $c$  mentioned in arXiv:1912.09729 Equ. 5, where  $c > 1$  is a constant indicating the lower bound. Default to 5.0 (set None if you do not want to use it).
- **value\_clip** (*bool*) – a parameter mentioned in arXiv:1811.02553v3 Sec. 4.1. Default to True.
- **advantage\_normalization** (*bool*) – whether to do per mini-batch advantage normalization. Default to True.
- **recompute\_advantage** (*bool*) – whether to recompute advantage every update repeat according to <https://arxiv.org/pdf/2006.05990.pdf> Sec. 3.5. Default to False.
- **vf\_coef** (*float*) – weight for value loss. Default to 0.5.
- **ent\_coef** (*float*) – weight for entropy loss. Default to 0.01.
- **max\_grad\_norm** (*float*) – clipping gradients in back propagation. Default to None.
- **gae\_lambda** (*float*) – in  $[0, 1]$ , param for Generalized Advantage Estimation. Default to 0.95.
- **reward\_normalization** (*bool*) – normalize estimated values to have std close to 1, also normalize the advantage to Normal(0, 1). Default to False.
- **max\_batchsize** (*int*) – the maximum size of the batch when computing GAE, depends on the size of available memory and the memory cost of the model; should be as large as possible within the memory constraint. Default to 256.
- **action\_scaling** (*bool*) – whether to map actions from range  $[-1, 1]$  to range  $[\text{action\_spaces.low}, \text{action\_spaces.high}]$ . Default to True.
- **action\_bound\_method** (*str*) – method to bound action to range  $[-1, 1]$ , can be either “clip” (for simply clipping the action), “tanh” (for applying tanh squashing) for now, or empty string for no bounding. Default to “clip”.
- **action\_space** (*Optional[gym.Space]*) – env’s action space, mandatory if you want to use option “action\_scaling” or “action\_bound\_method”. Default to None.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).
- **deterministic\_eval** (*bool*) – whether to use deterministic action instead of stochastic action sampled by the policy. Default to False.

See also:

Please refer to *BasePolicy* for more detailed explanation.

**process\_fn** (*batch*: *tianshou.data.batch.Batch*, *buffer*: *tianshou.data.buffer.base.ReplayBuffer*, *indices*: *numpy.ndarray*) → *tianshou.data.batch.Batch*  
 Compute the discounted returns for each transition.

$$G_t = \sum_{i=t}^T \gamma^{i-t} r_i$$

where  $T$  is the terminal time step,  $\gamma$  is the discount factor,  $\gamma \in [0, 1]$ .

**learn** (*batch*: *tianshou.data.batch.Batch*, *batch\_size*: *int*, *repeat*: *int*, *\*\*kwargs*: *Any*) → *Dict[str, List[float]]*  
 Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: `Categorical` distribution gives “[batch\_size]” shape while `Normal` distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** `bool`

## Off-policy

**class** `tianshou.policy.DDPGPolicy` (*actor*: *Optional[torch.nn.modules.module.Module]*, *actor\_optim*: *Optional[torch.optim.optimizer.Optimizer]*, *critic*: *Optional[torch.nn.modules.module.Module]*, *critic\_optim*: *Optional[torch.optim.optimizer.Optimizer]*, *tau*: *float* = 0.005, *gamma*: *float* = 0.99, *exploration\_noise*: *Optional[tianshou.exploration.random.BaseNoise]* = `<tianshou.exploration.random.GaussianNoise object>`, *reward\_normalization*: *bool* = `False`, *estimation\_step*: *int* = 1, *action\_scaling*: *bool* = `True`, *action\_bound\_method*: *str* = ‘clip’, *\*\*kwargs*: *Any*)

Bases: `tianshou.policy.base.BasePolicy`

Implementation of Deep Deterministic Policy Gradient. arXiv:1509.02971.

### Parameters

- **actor** (*torch.nn.Module*) – the actor network following the rules in `BasePolicy`. (s → logits)
- **actor\_optim** (*torch.optim.Optimizer*) – the optimizer for actor network.
- **critic** (*torch.nn.Module*) – the critic network. (s, a → Q(s, a))
- **critic\_optim** (*torch.optim.Optimizer*) – the optimizer for critic network.
- **tau** (*float*) – param for soft update of the target network. Default to 0.005.
- **gamma** (*float*) – discount factor, in [0, 1]. Default to 0.99.

- **exploration\_noise** (*BaseNoise*) – the exploration noise, add to the action. Default to `GaussianNoise(sigma=0.1)`.
- **reward\_normalization** (*bool*) – normalize the reward to `Normal(0, 1)`, Default to `False`.
- **estimation\_step** (*int*) – the number of steps to look ahead. Default to 1.
- **action\_scaling** (*bool*) – whether to map actions from range `[-1, 1]` to range `[action_spaces.low, action_spaces.high]`. Default to `True`.
- **action\_bound\_method** (*str*) – method to bound action to range `[-1, 1]`, can be either “clip” (for simply clipping the action) or empty string for no bounding. Default to “clip”.
- **action\_space** (*Optional[gym.Space]*) – env’s action space, mandatory if you want to use option “action\_scaling” or “action\_bound\_method”. Default to `None`.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each `policy.update()`. Default to `None` (no `lr_scheduler`).

See also:

Please refer to [BasePolicy](#) for more detailed explanation.

**set\_exp\_noise** (*noise: Optional[tianshou.exploration.random.BaseNoise]*) → `None`  
Set the exploration noise.

**train** (*mode: bool = True*) → `tianshou.policy.modelfree.ddpg.DDPGPolicy`  
Set the module in training mode, except for the target network.

**sync\_weight** () → `None`  
Soft-update the weight for the target network.

**process\_fn** (*batch: tianshou.data.batch.Batch, buffer: tianshou.data.buffer.base.ReplayBuffer, indices: numpy.ndarray*) → `tianshou.data.batch.Batch`  
Pre-process the data from the provided replay buffer.

Used in `update()`. Check out [policy.process\\_fn](#) for more information.

**forward** (*batch: tianshou.data.batch.Batch, state: Optional[Union[dict, tianshou.data.batch.Batch, numpy.ndarray]] = None, model: str = 'actor', input: str = 'obs', \*\*kwargs: Any*) → `tianshou.data.batch.Batch`  
Compute action over the given batch data.

#### Returns

A `Batch` which has 2 keys:

- `act` the action.
- `state` the hidden state.

See also:

Please refer to [forward\(\)](#) for more detailed explanation.

**learn** (*batch: tianshou.data.batch.Batch, \*\*kwargs: Any*) → `Dict[str, float]`  
Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

---



**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**exploration\_noise** (*act*: Union[numpy.ndarray, tianshou.data.batch.Batch], *batch*: tianshou.data.batch.Batch) → Union[numpy.ndarray, tianshou.data.batch.Batch]  
 Modify the action from policy.forward with exploration noise.

#### Parameters

- **act** – a data batch or numpy.ndarray which is the action taken by policy.forward.
- **batch** – the input batch for policy.forward, kept for advanced usage.

**Returns** action in the same form of input “act” but with added exploration noise.

**training:** bool

```
class tianshou.policy.TD3Policy (actor: torch.nn.modules.module.Module, actor_optim: torch.optim.optimizer.Optimizer, critic1: torch.nn.modules.module.Module, critic1_optim: torch.optim.optimizer.Optimizer, critic2: torch.nn.modules.module.Module, critic2_optim: torch.optim.optimizer.Optimizer, tau: float = 0.005, gamma: float = 0.99, exploration_noise: Optional[tianshou.exploration.random.BaseNoise] = <tianshou.exploration.random.GaussianNoise object>, policy_noise: float = 0.2, update_actor_freq: int = 2, noise_clip: float = 0.5, reward_normalization: bool = False, estimation_step: int = 1, **kwargs: Any)
```

Bases: tianshou.policy.modelfree.ddpg.DDPGPolicy

Implementation of TD3, arXiv:1802.09477.

#### Parameters

- **actor** (*torch.nn.Module*) – the actor network following the rules in *BasePolicy*. (s -> logits)
- **actor\_optim** (*torch.optim.Optimizer*) – the optimizer for actor network.
- **critic1** (*torch.nn.Module*) – the first critic network. (s, a -> Q(s, a))
- **critic1\_optim** (*torch.optim.Optimizer*) – the optimizer for the first critic network.
- **critic2** (*torch.nn.Module*) – the second critic network. (s, a -> Q(s, a))
- **critic2\_optim** (*torch.optim.Optimizer*) – the optimizer for the second critic network.
- **tau** (*float*) – param for soft update of the target network. Default to 0.005.
- **gamma** (*float*) – discount factor, in [0, 1]. Default to 0.99.
- **exploration\_noise** (*float*) – the exploration noise, add to the action. Default to GaussianNoise(sigma=0.1)
- **policy\_noise** (*float*) – the noise used in updating policy network. Default to 0.2.
- **update\_actor\_freq** (*int*) – the update frequency of actor network. Default to 2.

- **noise\_clip** (*float*) – the clipping range used in updating policy network. Default to 0.5.
- **reward\_normalization** (*bool*) – normalize the reward to Normal(0, 1). Default to False.
- **action\_scaling** (*bool*) – whether to map actions from range [-1, 1] to range [action\_spaces.low, action\_spaces.high]. Default to True.
- **action\_bound\_method** (*str*) – method to bound action to range [-1, 1], can be either “clip” (for simply clipping the action) or empty string for no bounding. Default to “clip”.
- **action\_space** (*Optional[gym.Space]*) – env’s action space, mandatory if you want to use option “action\_scaling” or “action\_bound\_method”. Default to None.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).

See also:

Please refer to [BasePolicy](#) for more detailed explanation.

**train** (*mode: bool = True*) → `tianshou.policy.modelfree.td3.TD3Policy`  
Set the module in training mode, except for the target network.

**sync\_weight** () → None  
Soft-update the weight for the target network.

**learn** (*batch: tianshou.data.batch.Batch, \*\*kwargs: Any*) → `Dict[str, float]`  
Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

---

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

```
training: bool
actor: torch.nn.Module
actor_optim: torch.optim.Optimizer
critic: torch.nn.Module
critic_optim: torch.optim.Optimizer
```

```

class tianshou.policy.SACPolicy(actor: torch.nn.modules.module.Module, ac-
    tor_optim: torch.optim.optimizer.Optimizer,
    critic1: torch.nn.modules.module.Module,
    critic1_optim: torch.optim.optimizer.Optimizer, critic2:
    torch.nn.modules.module.Module, critic2_optim:
    torch.optim.optimizer.Optimizer, tau: float = 0.005,
    gamma: float = 0.99, alpha: Union[float, Tu-
    ple[float, torch.Tensor, torch.optim.optimizer.Optimizer]]
    = 0.2, reward_normalization: bool = False, es-
    timation_step: int = 1, exploration_noise: Op-
    tional[tianshou.exploration.random.BaseNoise] = None,
    deterministic_eval: bool = True, **kwargs: Any)

```

Bases: `tianshou.policy.modelfree.ddpg.DDPGPolicy`

Implementation of Soft Actor-Critic. arXiv:1812.05905.

### Parameters

- **actor** (`torch.nn.Module`) – the actor network following the rules in *BasePolicy*. (s -> logits)
- **actor\_optim** (`torch.optim.Optimizer`) – the optimizer for actor network.
- **critic1** (`torch.nn.Module`) – the first critic network. (s, a -> Q(s, a))
- **critic1\_optim** (`torch.optim.Optimizer`) – the optimizer for the first critic network.
- **critic2** (`torch.nn.Module`) – the second critic network. (s, a -> Q(s, a))
- **critic2\_optim** (`torch.optim.Optimizer`) – the optimizer for the second critic network.
- **tau** (`float`) – param for soft update of the target network. Default to 0.005.
- **gamma** (`float`) – discount factor, in [0, 1]. Default to 0.99.
- **torch.Tensor, torch.optim.Optimizer** or **float alpha** (`(float,)`) – entropy regularization coefficient. Default to 0.2. If a tuple (target\_entropy, log\_alpha, alpha\_optim) is provided, then alpha is automatically tuned.
- **reward\_normalization** (`bool`) – normalize the reward to Normal(0, 1). Default to False.
- **exploration\_noise** (`BaseNoise`) – add a noise to action for exploration. Default to None. This is useful when solving hard-exploration problem.
- **deterministic\_eval** (`bool`) – whether to use deterministic action (mean of Gaussian policy) instead of stochastic action sampled by the policy. Default to True.
- **action\_scaling** (`bool`) – whether to map actions from range [-1, 1] to range [action\_spaces.low, action\_spaces.high]. Default to True.
- **action\_bound\_method** (`str`) – method to bound action to range [-1, 1], can be either “clip” (for simply clipping the action) or empty string for no bounding. Default to “clip”.
- **action\_space** (`Optional[gym.Space]`) – env’s action space, mandatory if you want to use option “action\_scaling” or “action\_bound\_method”. Default to None.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).

**See also:**

Please refer to [BasePolicy](#) for more detailed explanation.

**actor:** `torch.nn.Module`

**actor\_optim:** `torch.optim.Optimizer`

**train** (*mode: bool = True*) → `tianshou.policy.modelfree.sac.SACPolicy`  
Set the module in training mode, except for the target network.

**sync\_weight** () → `None`  
Soft-update the weight for the target network.

**forward** (*batch: tianshou.data.batch.Batch, state: Optional[Union[dict, tianshou.data.batch.Batch, numpy.ndarray]] = None, input: str = 'obs', \*\*kwargs: Any*) → `tianshou.data.batch.Batch`  
Compute action over the given batch data.

**Returns**

A [Batch](#) which has 2 keys:

- `act` the action.
- `state` the hidden state.

**See also:**

Please refer to [forward\(\)](#) for more detailed explanation.

**training:** `bool`

**critic:** `torch.nn.Module`

**critic\_optim:** `torch.optim.Optimizer`

**learn** (*batch: tianshou.data.batch.Batch, \*\*kwargs: Any*) → `Dict[str, float]`  
Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

---

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

```

class tianshou.policy.REDQPolicy(
    actor: torch.nn.modules.module.Module,
    actor_optim: torch.optim.optimizer.Optimizer,
    critics: torch.nn.modules.module.Module,
    critics_optim: torch.optim.optimizer.Optimizer,
    ensemble_size: int = 10,
    subset_size: int = 2,
    tau: float = 0.005,
    gamma: float = 0.99,
    alpha: Union[float, Tuple[float, torch.Tensor, torch.optim.optimizer.Optimizer]] = 0.2,
    reward_normalization: bool = False,
    estimation_step: int = 1,
    actor_delay: int = 20,
    exploration_noise: Optional[tianshou.exploration.random.BaseNoise] = None,
    deterministic_eval: bool = True,
    target_mode: str = 'min',
    **kwargs: Any
)

```

Bases: `tianshou.policy.modelfree.ddpg.DDPGPolicy`

Implementation of REDQ. arXiv:2101.05982.

### Parameters

- **actor** (`torch.nn.Module`) – the actor network following the rules in *BasePolicy*. (s -> logits)
- **actor\_optim** (`torch.optim.Optimizer`) – the optimizer for actor network.
- **critics** (`torch.nn.Module`) – critic ensemble networks.
- **critics\_optim** (`torch.optim.Optimizer`) – the optimizer for the critic networks.
- **ensemble\_size** (`int`) – Number of sub-networks in the critic ensemble. Default to 10.
- **subset\_size** (`int`) – Number of networks in the subset. Default to 2.
- **tau** (`float`) – param for soft update of the target network. Default to 0.005.
- **gamma** (`float`) – discount factor, in [0, 1]. Default to 0.99.
- **torch.Tensor, torch.optim.Optimizer** or **float alpha** (`(float,)`) – entropy regularization coefficient. Default to 0.2. If a tuple (target\_entropy, log\_alpha, alpha\_optim) is provided, then alpha is automatically tuned.
- **reward\_normalization** (`bool`) – normalize the reward to Normal(0, 1). Default to False.
- **actor\_delay** (`int`) – Number of critic updates before an actor update. Default to 20.
- **exploration\_noise** (`BaseNoise`) – add a noise to action for exploration. Default to None. This is useful when solving hard-exploration problem.
- **deterministic\_eval** (`bool`) – whether to use deterministic action (mean of Gaussian policy) instead of stochastic action sampled by the policy. Default to True.
- **target\_mode** (`str`) – methods to integrate critic values in the subset, currently support minimum and average. Default to min.
- **action\_scaling** (`bool`) – whether to map actions from range [-1, 1] to range [action\_spaces.low, action\_spaces.high]. Default to True.
- **action\_bound\_method** (`str`) – method to bound action to range [-1, 1], can be either “clip” (for simply clipping the action) or empty string for no bounding. Default to “clip”.
- **action\_space** (`Optional[gym.Space]`) – env’s action space, mandatory if you want to use option “action\_scaling” or “action\_bound\_method”. Default to None.

See also:

Please refer to [BasePolicy](#) for more detailed explanation.

**actor:** `torch.nn.Module`

**actor\_optim:** `torch.optim.Optimizer`

**train** (*mode: bool = True*) → `tianshou.policy.modelfree.redq.REDQPolicy`  
Set the module in training mode, except for the target network.

**sync\_weight** () → `None`  
Soft-update the weight for the target network.

**forward** (*batch: tianshou.data.batch.Batch, state: Optional[Union[dict, tianshou.data.batch.Batch, numpy.ndarray]] = None, input: str = 'obs', \*\*kwargs: Any*) → `tianshou.data.batch.Batch`  
Compute action over the given batch data.

**Returns**

A [Batch](#) which has 2 keys:

- `act` the action.
- `state` the hidden state.

See also:

Please refer to [forward\(\)](#) for more detailed explanation.

**training:** `bool`

**critic:** `torch.nn.Module`

**critic\_optim:** `torch.optim.Optimizer`

**learn** (*batch: tianshou.data.batch.Batch, \*\*kwargs: Any*) → `Dict[str, float]`  
Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

---

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

```

class tianshou.policy.DiscreteSACPolicy(actor: torch.nn.modules.module.Module, ac-
tor_optim: torch.optim.optimizer.Optimizer,
critic1: torch.nn.modules.module.Module,
critic1_optim: torch.optim.optimizer.Optimizer,
critic2: torch.nn.modules.module.Module,
critic2_optim: torch.optim.optimizer.Optimizer,
tau: float = 0.005, gamma: float = 0.99, al-
pha: Union[float, Tuple[float, torch.Tensor,
torch.optim.optimizer.Optimizer]] = 0.2, re-
ward_normalization: bool = False, estima-
tion_step: int = 1, **kwargs: Any)

```

Bases: `tianshou.policy.modelfree.sac.SACPolicy`

Implementation of SAC for Discrete Action Settings. arXiv:1910.07207.

### Parameters

- **actor** (`torch.nn.Module`) – the actor network following the rules in *BasePolicy*. (s -> logits)
- **actor\_optim** (`torch.optim.Optimizer`) – the optimizer for actor network.
- **critic1** (`torch.nn.Module`) – the first critic network. (s -> Q(s))
- **critic1\_optim** (`torch.optim.Optimizer`) – the optimizer for the first critic network.
- **critic2** (`torch.nn.Module`) – the second critic network. (s -> Q(s))
- **critic2\_optim** (`torch.optim.Optimizer`) – the optimizer for the second critic network.
- **tau** (`float`) – param for soft update of the target network. Default to 0.005.
- **gamma** (`float`) – discount factor, in [0, 1]. Default to 0.99.
- **torch.Tensor, torch.optim.Optimizer) or float alpha** (`(float,)`) – entropy regularization coefficient. Default to 0.2. If a tuple (target\_entropy, log\_alpha, alpha\_optim) is provided, the alpha is automatically tuned.
- **reward\_normalization** (`bool`) – normalize the reward to Normal(0, 1). Default to False.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).

### See also:

Please refer to *BasePolicy* for more detailed explanation.

**forward** (`batch: tianshou.data.batch.Batch, state: Optional[Union[dict, tianshou.data.batch.Batch, numpy.ndarray]] = None, input: str = 'obs', **kwargs: Any`) → `tianshou.data.batch.Batch`  
 Compute action over the given batch data.

### Returns

A *Batch* which has 2 keys:

- `act` the action.
- `state` the hidden state.

### See also:

Please refer to *forward()* for more detailed explanation.

**learn** (*batch*: *tianshou.data.batch.Batch*, *\*\*kwargs*: *Any*) → Dict[str, float]

Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

---

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**exploration\_noise** (*act*: *Union[numpy.ndarray, tianshou.data.batch.Batch]*, *batch*: *tianshou.data.batch.Batch*) → *Union[numpy.ndarray, tianshou.data.batch.Batch]*

Modify the action from `policy.forward` with exploration noise.

**Parameters**

- **act** – a data batch or `numpy.ndarray` which is the action taken by `policy.forward`.
- **batch** – the input batch for `policy.forward`, kept for advanced usage.

**Returns** action in the same form of input “act” but with added exploration noise.

**training:** bool

**actor:** torch.nn.Module

**actor\_optim:** torch.optim.Optimizer

**critic:** torch.nn.Module

**critic\_optim:** torch.optim.Optimizer

### 1.11.3 Imitation

**class** `tianshou.policy.ImitationPolicy` (*model*: *torch.nn.modules.module.Module*, *optim*: *torch.optim.optimizer.Optimizer*, *\*\*kwargs*: *Any*)

Bases: `tianshou.policy.base.BasePolicy`

Implementation of vanilla imitation learning.

**Parameters**

- **model** (*torch.nn.Module*) – a model following the rules in [BasePolicy](#). (s → a)
- **optim** (*torch.optim.Optimizer*) – for optimizing the model.
- **action\_space** (*gym.Space*) – env’s action space.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each `policy.update()`. Default to None (no `lr_scheduler`).

**See also:**

Please refer to [BasePolicy](#) for more detailed explanation.



**forward** (*batch*: *tianshou.data.batch.Batch*, *state*: *Optional[Union[dict, tianshou.data.batch.Batch, numpy.ndarray]] = None*, *\*\*kwargs*: *Any*) → *tianshou.data.batch.Batch*  
 Compute action over the given batch data.

#### Returns

A *Batch* which MUST have the following keys:

- *act* an *numpy.ndarray* or a *torch.Tensor*, the action over given batch data.
- *state* a dict, an *numpy.ndarray* or a *torch.Tensor*, the internal state of the policy, *None* as default.

Other keys are user-defined. It depends on the algorithm. For example,

```
# some code
return Batch(logits=..., act=..., state=None, dist=...)
```

The keyword *policy* is reserved and the corresponding data will be stored into the replay buffer. For instance,

```
# some code
return Batch(..., policy=Batch(log_prob=dist.log_prob(act)))
# and in the sampled data batch, you can directly use
# batch.policy.log_prob to get your data.
```

**Note:** In continuous action space, you should do another step “map\_action” to get the real action:

```
act = policy(batch).act # doesn't map to the target action range
act = policy.map_action(act, batch)
```

**learn** (*batch*: *tianshou.data.batch.Batch*, *\*\*kwargs*: *Any*) → *Dict[str, float]*

Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by *self.training* and *self.updating*. Please refer to [States for policy](#) for more detailed explanation.

**Warning:** If you use *torch.distributions.Normal* and *torch.distributions.Categorical* to calculate the *log\_prob*, please be careful about the shape: *Categorical* distribution gives “[batch\_size]” shape while *Normal* distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** *bool*

```

class tianshou.policy.BCQPolicy(actor: torch.nn.modules.module.Module, ac-
    tor_optim: torch.optim.optimizer.Optimizer,
    critic1: torch.nn.modules.module.Module,
    critic1_optim: torch.optim.optimizer.Optimizer,
    critic2: torch.nn.modules.module.Module,
    critic2_optim: torch.optim.optimizer.Optimizer,
    vae: tianshou.utils.net.continuous.VAE, vae_optim:
    torch.optim.optimizer.Optimizer, device: Union[str,
    torch.device] = 'cpu', gamma: float = 0.99, tau: float =
    0.005, lmbda: float = 0.75, forward_sampled_times: int = 100,
    num_sampled_action: int = 10, **kwargs: Any)

```

Bases: `tianshou.policy.base.BasePolicy`

Implementation of BCQ algorithm. arXiv:1812.02900.

### Parameters

- **actor** (`Perturbation`) – the actor perturbation. (s, a -> perturbed a)
- **actor\_optim** (`torch.optim.Optimizer`) – the optimizer for actor network.
- **critic1** (`torch.nn.Module`) – the first critic network. (s, a -> Q(s, a))
- **critic1\_optim** (`torch.optim.Optimizer`) – the optimizer for the first critic network.
- **critic2** (`torch.nn.Module`) – the second critic network. (s, a -> Q(s, a))
- **critic2\_optim** (`torch.optim.Optimizer`) – the optimizer for the second critic network.
- **vae** (`VAE`) – the VAE network, generating actions similar to those in batch. (s, a -> generated a)
- **vae\_optim** (`torch.optim.Optimizer`) – the optimizer for the VAE network.
- **torch.device** **device** (`Union[str,)` – which device to create this model on. Default to “cpu”.
- **gamma** (`float`) – discount factor, in [0, 1]. Default to 0.99.
- **tau** (`float`) – param for soft update of the target network. Default to 0.005.
- **lmbda** (`float`) – param for Clipped Double Q-learning. Default to 0.75.
- **forward\_sampled\_times** (`int`) – the number of sampled actions in forward function. The policy samples many actions and takes the action with the max value. Default to 100.
- **num\_sampled\_action** (`int`) – the number of sampled actions in calculating target Q. The algorithm samples several actions using VAE, and perturbs each action to get the target Q. Default to 10.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each `policy.update()`. Default to None (no `lr_scheduler`).

### See also:

Please refer to `BasePolicy` for more detailed explanation.

**train** (`mode: bool = True`) → `tianshou.policy.imitation.bcq.BCQPolicy`

Set the module in training mode, except for the target network.

**forward** (`batch: tianshou.data.batch.Batch, state: Optional[Union[dict, tianshou.data.batch.Batch, numpy.ndarray]] = None, **kwargs: Any`) → `tianshou.data.batch.Batch`

Compute action over the given batch data.

**sync\_weight** () → None

Soft-update the weight for the target network.

**learn** (batch: *tianshou.data.batch.Batch*, \*\*kwargs: Any) → Dict[str, float]

Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

---

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: `Categorical` distribution gives “[batch\_size]” shape while `Normal` distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** bool

```
class tianshou.policy.CQLPolicy (actor: tianshou.utils.net.continuous.ActorProb,
                                  actor_optim: torch.optim.optimizer.Optimizer,
                                  critic1: torch.nn.modules.module.Module,
                                  critic1_optim: torch.optim.optimizer.Optimizer, critic2:
                                  torch.nn.modules.module.Module, critic2_optim:
                                  torch.optim.optimizer.Optimizer, cql_alpha_lr: float =
                                  0.0001, cql_weight: float = 1.0, tau: float = 0.005, gamma:
                                  float = 0.99, alpha: Union[float, Tuple[float, torch.Tensor,
                                  torch.optim.optimizer.Optimizer]] = 0.2, temperature: float =
                                  1.0, with_lagrange: bool = True, lagrange_threshold:
                                  float = 10.0, min_action: float = - 1.0, max_action: float =
                                  1.0, num_repeat_actions: int = 10, alpha_min: float = 0.0,
                                  alpha_max: float = 1000000.0, clip_grad: float = 1.0, device:
                                  Union[str, torch.device] = 'cpu', **kwargs: Any)
```

Bases: `tianshou.policy.modelfree.sac.SACPolicy`

Implementation of CQL algorithm. arXiv:2006.04779.

#### Parameters

- **actor** (`ActorProb`) – the actor network following the rules in `BasePolicy`. (s → a)
- **actor\_optim** (`torch.optim.Optimizer`) – the optimizer for actor network.
- **critic1** (`torch.nn.Module`) – the first critic network. (s, a → Q(s, a))
- **critic1\_optim** (`torch.optim.Optimizer`) – the optimizer for the first critic network.
- **critic2** (`torch.nn.Module`) – the second critic network. (s, a → Q(s, a))
- **critic2\_optim** (`torch.optim.Optimizer`) – the optimizer for the second critic network.
- **cql\_alpha\_lr** (`float`) – the learning rate of `cql_log_alpha`. Default to 1e-4.
- **cql\_weight** (`float`) – the value of `alpha`. Default to 1.0.
- **tau** (`float`) – param for soft update of the target network. Default to 0.005.

- **gamma** (*float*) – discount factor, in  $[0, 1]$ . Default to 0.99.
- **torch.Tensor, torch.optim.Optimizer** or **float alpha** (*float*,) – entropy regularization coefficient. Default to 0.2. If a tuple (target\_entropy, log\_alpha, alpha\_optim) is provided, then alpha is automatically tuned.
- **temperature** (*float*) – the value of temperature. Default to 1.0.
- **with\_lagrange** (*bool*) – whether to use Lagrange. Default to True.
- **lagrange\_threshold** (*float*) – the value of tau in CQL(Lagrange). Default to 10.0.
- **min\_action** (*float*) – The minimum value of each dimension of action. Default to -1.0.
- **max\_action** (*float*) – The maximum value of each dimension of action. Default to 1.0.
- **num\_repeat\_actions** (*int*) – The number of times the action is repeated when calculating log-sum-exp. Default to 10.
- **alpha\_min** (*float*) – lower bound for clipping cql\_alpha. Default to 0.0.
- **alpha\_max** (*float*) – upper bound for clipping cql\_alpha. Default to 1e6.
- **clip\_grad** (*float*) – clip\_grad for updating critic network. Default to 1.0.
- **torch.device]** **device** (*Union[str,]*) – which device to create this model on. Default to “cpu”.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).

See also:

Please refer to [BasePolicy](#) for more detailed explanation.

**train** (*mode: bool = True*) → `tianshou.policy.imitation.cql.CQLPolicy`  
Set the module in training mode, except for the target network.

**sync\_weight** () → None  
Soft-update the weight for the target network.

**actor\_pred** (*obs: torch.Tensor*) → `Tuple[torch.Tensor, torch.Tensor]`

**calc\_actor\_loss** (*obs: torch.Tensor*) → `Tuple[torch.Tensor, torch.Tensor]`

**calc\_pi\_values** (*obs\_pi: torch.Tensor, obs\_to\_pred: torch.Tensor*) → `Tuple[torch.Tensor, torch.Tensor]`

**calc\_random\_values** (*obs: torch.Tensor, act: torch.Tensor*) → `Tuple[torch.Tensor, torch.Tensor]`

**process\_fn** (*batch: tianshou.data.batch.Batch, buffer: tianshou.data.buffer.base.ReplayBuffer, indices: numpy.ndarray*) → `tianshou.data.batch.Batch`  
Pre-process the data from the provided replay buffer.

Used in `update()`. Check out [policy.process\\_fn](#) for more information.

**learn** (*batch: tianshou.data.batch.Batch, \*\*kwargs: Any*) → `Dict[str, float]`  
Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

---

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

```

training: bool

actor: torch.nn.Module

actor_optim: torch.optim.Optimizer

critic: torch.nn.Module

critic_optim: torch.optim.Optimizer

class tianshou.policy.DiscreteBCQPolicy(model: torch.nn.modules.module.Module, imi-
                                     tator: torch.nn.modules.module.Module, optim:
                                     torch.optim.optimizer.Optimizer, discount_factor:
                                     float = 0.99, estimation_step: int = 1, tar-
                                     get_update_freq: int = 8000, eval_eps: float
                                     = 0.001, unlikely_action_threshold: float =
                                     0.3, imitation_logits_penalty: float = 0.01, re-
                                     ward_normalization: bool = False, **kwargs:
                                     Any)
Bases: tianshou.policy.modelfree.dqn.DQNPolicy
Implementation of discrete BCQ algorithm. arXiv:1910.01708.

```

#### Parameters

- **model** (*torch.nn.Module*) – a model following the rules in *BasePolicy*. (s -> q\_value)
- **imitator** (*torch.nn.Module*) – a model following the rules in *BasePolicy*. (s -> imitation\_logits)
- **optim** (*torch.optim.Optimizer*) – a torch.optim for optimizing the model.
- **discount\_factor** (*float*) – in [0, 1].
- **estimation\_step** (*int*) – the number of steps to look ahead. Default to 1.
- **target\_update\_freq** (*int*) – the target network update frequency.
- **eval\_eps** (*float*) – the epsilon-greedy noise added in evaluation.
- **unlikely\_action\_threshold** (*float*) – the threshold (tau) for unlikely actions, as shown in Equ. (17) in the paper. Default to 0.3.
- **imitation\_logits\_penalty** (*float*) – regularization weight for imitation logits. Default to 1e-2.
- **reward\_normalization** (*bool*) – normalize the reward to Normal(0, 1). Default to False.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).

See also:

Please refer to *BasePolicy* for more detailed explanation.

**train** (*mode: bool = True*) → `tianshou.policy.imitation.discrete_bcq.DiscreteBCQPolicy`

Set the module in training mode, except for the target network.

**forward** (*batch: tianshou.data.batch.Batch, state: Optional[Union[dict, tianshou.data.batch.Batch, numpy.ndarray]] = None, input: str = 'obs', \*\*kwargs: Any*) → `tianshou.data.batch.Batch`

Compute action over the given batch data.

If you need to mask the action, please add a “mask” into `batch.obs`, for example, if we have an environment that has “0/1/2” three actions:

```
batch == Batch(
  obs=Batch(
    obs="original obs, with batch_size=1 for demonstration",
    mask=np.array([[False, True, False]]),
    # action 1 is available
    # action 0 and 2 are unavailable
  ),
  ...
)
```

**Parameters** `eps` (*float*) – in [0, 1], for epsilon-greedy exploration method.

#### Returns

A `Batch` which has 3 keys:

- `act` the action.
- `logits` the network’s raw output.
- `state` the hidden state.

#### See also:

Please refer to `forward()` for more detailed explanation.

**learn** (*batch: tianshou.data.batch.Batch, \*\*kwargs: Any*) → `Dict[str, float]`

Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

---

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** `bool`

```
class tianshou.policy.DiscreteCQLPolicy (model: torch.nn.modules.module.Module, optim:
    torch.optim.optimizer.Optimizer, discount_factor:
    float = 0.99, num_quantiles: int = 200, estima-
    tion_step: int = 1, target_update_freq: int = 0, re-
    ward_normalization: bool = False, min_q_weight:
    float = 10.0, **kwargs: Any)
```

Bases: `tianshou.policy.modelfree.qrdqn.QRDQNPolicy`

Implementation of discrete Conservative Q-Learning algorithm. arXiv:2006.04779.

#### Parameters

- **model** (`torch.nn.Module`) – a model following the rules in `BasePolicy`. (s -> logits)
- **optim** (`torch.optim.Optimizer`) – a torch.optim for optimizing the model.
- **discount\_factor** (`float`) – in [0, 1].
- **num\_quantiles** (`int`) – the number of quantile midpoints in the inverse cumulative distribution function of the value. Default to 200.
- **estimation\_step** (`int`) – the number of steps to look ahead. Default to 1.
- **target\_update\_freq** (`int`) – the target network update frequency (0 if you do not use the target network).
- **reward\_normalization** (`bool`) – normalize the reward to Normal(0, 1). Default to False.
- **min\_q\_weight** (`float`) – the weight for the cql loss.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).

#### See also:

Please refer to `QRDQNPolicy` for more detailed explanation.

**learn** (`batch: tianshou.data.batch.Batch`, `**kwargs: Any`) → Dict[str, float]

Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

---

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** bool

```
class tianshou.policy.DiscreteCRRPolicy (actor: torch.nn.modules.module.Module,  
                                         critic: torch.nn.modules.module.Module, optim:  
                                         torch.optim.optimizer.Optimizer, discount_factor:  
                                         float = 0.99, policy_improvement_mode: str =  
                                         'exp', ratio_upper_bound: float = 20.0, beta:  
                                         float = 1.0, min_q_weight: float = 10.0, tar-  
                                         get_update_freq: int = 0, reward_normalization:  
                                         bool = False, **kwargs: Any)
```

Bases: `tianshou.policy.modelfree.pg.PGPolicy`

Implementation of discrete Critic Regularized Regression. arXiv:2006.15134.

### Parameters

- **actor** (`torch.nn.Module`) – the actor network following the rules in *BasePolicy*. (s -> logits)
- **critic** (`torch.nn.Module`) – the action-value critic (i.e., Q function) network. (s -> Q(s, \*))
- **optim** (`torch.optim.Optimizer`) – a torch.optim for optimizing the model.
- **discount\_factor** (`float`) – in [0, 1]. Default to 0.99.
- **policy\_improvement\_mode** (`str`) – type of the weight function f. Possible values: “binary”/“exp”/“all”. Default to “exp”.
- **ratio\_upper\_bound** (`float`) – when policy\_improvement\_mode is “exp”, the value of the exp function is upper-bounded by this parameter. Default to 20.
- **beta** (`float`) – when policy\_improvement\_mode is “exp”, this is the denominator of the exp function. Default to 1.
- **min\_q\_weight** (`float`) – weight for CQL loss/regularizer. Default to 10.
- **target\_update\_freq** (`int`) – the target network update frequency (0 if you do not use the target network). Default to 0.
- **reward\_normalization** (`bool`) – normalize the reward to Normal(0, 1). Default to False.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).

### See also:

Please refer to *PGPolicy* for more detailed explanation.

**sync\_weight** () → None

**training**: bool

**learn** (*batch*: `tianshou.data.batch.Batch`, **\*\*kwargs**: Any) → Dict[str, float]

Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to *States for policy* for more detailed explanation.

---



**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

```
class tianshou.policy.GAILPolicy(actor: torch.nn.modules.module.Module,
                                critic: torch.nn.modules.module.Module, op-
                                tim: torch.optim.optimizer.Optimizer, dist_fn:
                                Type[torch.distributions.distribution.Distribution], ex-
                                pert_buffer: tianshou.data.buffer.base.ReplayBuffer,
                                disc_net: torch.nn.modules.module.Module, disc_optim:
                                torch.optim.optimizer.Optimizer, disc_update_num: int = 4,
                                eps_clip: float = 0.2, dual_clip: Optional[float] = None,
                                value_clip: bool = False, advantage_normalization: bool =
                                True, recompute_advantage: bool = False, **kwargs: Any)
Bases: tianshou.policy.modelfree.ppo.PPOPolicy
```

Implementation of Generative Adversarial Imitation Learning. arXiv:1606.03476.

#### Parameters

- **actor** (`torch.nn.Module`) – the actor network following the rules in `BasePolicy`. (s -> logits)
- **critic** (`torch.nn.Module`) – the critic network. (s -> V(s))
- **optim** (`torch.optim.Optimizer`) – the optimizer for actor and critic network.
- **dist\_fn** (`Type[torch.distributions.Distribution]`) – distribution class for computing the action.
- **expert\_buffer** (`ReplayBuffer`) – the replay buffer contains expert experience.
- **disc\_net** (`torch.nn.Module`) – the discriminator network with input dim equals state dim plus action dim and output dim equals 1.
- **disc\_optim** (`torch.optim.Optimizer`) – the optimizer for the discriminator network.
- **disc\_update\_num** (`int`) – the number of discriminator grad steps per model grad step. Default to 4.
- **discount\_factor** (`float`) – in [0, 1]. Default to 0.99.
- **eps\_clip** (`float`) –  $\epsilon$  in  $L_{CLIP}$  in the original paper. Default to 0.2.
- **dual\_clip** (`float`) – a parameter  $c$  mentioned in arXiv:1912.09729 Equ. 5, where  $c > 1$  is a constant indicating the lower bound. Default to 5.0 (set None if you do not want to use it).
- **value\_clip** (`bool`) – a parameter mentioned in arXiv:1811.02553 Sec. 4.1. Default to True.
- **advantage\_normalization** (`bool`) – whether to do per mini-batch advantage normalization. Default to True.
- **recompute\_advantage** (`bool`) – whether to recompute advantage every update repeat according to <https://arxiv.org/pdf/2006.05990.pdf> Sec. 3.5. Default to False.
- **vf\_coef** (`float`) – weight for value loss. Default to 0.5.
- **ent\_coef** (`float`) – weight for entropy loss. Default to 0.01.

- **max\_grad\_norm** (*float*) – clipping gradients in back propagation. Default to None.
- **gae\_lambda** (*float*) – in  $[0, 1]$ , param for Generalized Advantage Estimation. Default to 0.95.
- **reward\_normalization** (*bool*) – normalize estimated values to have std close to 1, also normalize the advantage to  $\text{Normal}(0, 1)$ . Default to False.
- **max\_batchsize** (*int*) – the maximum size of the batch when computing GAE, depends on the size of available memory and the memory cost of the model; should be as large as possible within the memory constraint. Default to 256.
- **action\_scaling** (*bool*) – whether to map actions from range  $[-1, 1]$  to range  $[\text{action\_spaces.low}, \text{action\_spaces.high}]$ . Default to True.
- **action\_bound\_method** (*str*) – method to bound action to range  $[-1, 1]$ , can be either “clip” (for simply clipping the action), “tanh” (for applying tanh squashing) for now, or empty string for no bounding. Default to “clip”.
- **action\_space** (*Optional[gym.Space]*) – env’s action space, mandatory if you want to use option “action\_scaling” or “action\_bound\_method”. Default to None.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).
- **deterministic\_eval** (*bool*) – whether to use deterministic action instead of stochastic action sampled by the policy. Default to False.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).

See also:

Please refer to [PPOPolicy](#) for more detailed explanation.

**process\_fn** (*batch: tianshou.data.batch.Batch, buffer: tianshou.data.buffer.base.ReplayBuffer, indices: numpy.ndarray*)  $\rightarrow$  *tianshou.data.batch.Batch*  
Pre-process the data from the provided replay buffer.

Used in `update()`. Check out [policy.process\\_fn](#) for more information.

**disc** (*batch: tianshou.data.batch.Batch*)  $\rightarrow$  *torch.Tensor*

**learn** (*batch: tianshou.data.batch.Batch, batch\_size: int, repeat: int, \*\*kwargs: Any*)  $\rightarrow$  *Dict[str, List[float]]*  
Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

---

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: `Categorical` distribution gives “[batch\_size]” shape while `Normal` distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** *bool*

### 1.11.4 Model-based

```
class tianshou.policy.PSRLPolicy(trans_count_prior: numpy.ndarray, rew_mean_prior:
                                numpy.ndarray, rew_std_prior: numpy.ndarray, dis-
                                count_factor: float = 0.99, epsilon: float = 0.01,
                                add_done_loop: bool = False, **kwargs: Any)
```

Bases: `tianshou.policy.base.BasePolicy`

Implementation of Posterior Sampling Reinforcement Learning.

Reference: Strens M. A Bayesian framework for reinforcement learning [C] //ICML. 2000, 2000: 943-950.

#### Parameters

- **trans\_count\_prior** (`np.ndarray`) – dirichlet prior (alphas), with shape (n\_state, n\_action, n\_state).
- **rew\_mean\_prior** (`np.ndarray`) – means of the normal priors of rewards, with shape (n\_state, n\_action).
- **rew\_std\_prior** (`np.ndarray`) – standard deviations of the normal priors of rewards, with shape (n\_state, n\_action).
- **discount\_factor** (`float`) – in [0, 1].
- **epsilon** (`float`) – for precision control in value iteration.
- **add\_done\_loop** (`bool`) – whether to add an extra self-loop for the terminal state in MDP. Default to False.

See also:

Please refer to [BasePolicy](#) for more detailed explanation.

**forward** (*batch*: `tianshou.data.batch.Batch`, *state*: `Optional[Union[dict, tianshou.data.batch.Batch, numpy.ndarray]] = None`, *\*\*kwargs*: `Any`) → `tianshou.data.batch.Batch`  
 Compute action over the given batch data with PSRL model.

**Returns** A `Batch` with “act” key containing the action.

See also:

Please refer to `forward()` for more detailed explanation.

**learn** (*batch*: `tianshou.data.batch.Batch`, *\*args*: `Any`, *\*\*kwargs*: `Any`) → `Dict[str, float]`  
 Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

---

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

---

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: `Categorical` distribution gives “[batch\_size]” shape while `Normal` distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** `bool`

```
class tianshou.policy.ICMPolicy (policy: tianshou.policy.base.BasePolicy, model: tianshou.
utils.net.discrete.IntrinsicCuriosityModule, optim: torch.optim.optimizer.Optimizer, lr_scale: float, reward_scale:
float, forward_loss_weight: float, **kwargs: Any)
```

Bases: `tianshou.policy.base.BasePolicy`

Implementation of Intrinsic Curiosity Module. arXiv:1705.05363.

#### Parameters

- **policy** (`BasePolicy`) – a base policy to add ICM to.
- **model** (`IntrinsicCuriosityModule`) – the ICM model.
- **optim** (`torch.optim.Optimizer`) – a torch.optim for optimizing the model.
- **lr\_scale** (`float`) – the scaling factor for ICM learning.
- **forward\_loss\_weight** (`float`) – the weight for forward model loss.
- **lr\_scheduler** – a learning rate scheduler that adjusts the learning rate in optimizer in each policy.update(). Default to None (no lr\_scheduler).

#### See also:

Please refer to `BasePolicy` for more detailed explanation.

**train** (*mode: bool = True*) → `tianshou.policy.modelbased.icm.ICMPolicy`  
Set the module in training mode.

**forward** (*batch: tianshou.data.batch.Batch, state: Optional[Union[dict, tianshou.data.batch.Batch, numpy.ndarray]] = None, \*\*kwargs: Any*) → `tianshou.data.batch.Batch`  
Compute action over the given batch data by inner policy.

#### See also:

Please refer to `forward()` for more detailed explanation.

**exploration\_noise** (*act: Union[numpy.ndarray, tianshou.data.batch.Batch], batch: tianshou.data.batch.Batch*) → `Union[numpy.ndarray, tianshou.data.batch.Batch]`  
Modify the action from policy.forward with exploration noise.

#### Parameters

- **act** – a data batch or numpy.ndarray which is the action taken by policy.forward.
- **batch** – the input batch for policy.forward, kept for advanced usage.

**Returns** action in the same form of input “act” but with added exploration noise.

**set\_eps** (*eps: float*) → None  
Set the eps for epsilon-greedy exploration.

**process\_fn** (*batch: tianshou.data.batch.Batch, buffer: tianshou.data.buffer.base.ReplayBuffer, indices: numpy.ndarray*) → `tianshou.data.batch.Batch`  
Pre-process the data from the provided replay buffer.

Used in `update()`. Check out `policy.process_fn` for more information.

**post\_process\_fn** (*batch: tianshou.data.batch.Batch, buffer: tianshou.data.buffer.base.ReplayBuffer, indices: numpy.ndarray*) → None  
Post-process the data from the provided replay buffer.

Typical usage is to update the sampling weight in prioritized experience replay. Used in `update()`.

**learn** (*batch: tianshou.data.batch.Batch, \*\*kwargs: Any*) → `Dict[str, float]`  
Update policy with a given batch of data.

**Returns** A dict, including the data needed to be logged (e.g., loss).

**Note:** In order to distinguish the collecting state, updating state and testing state, you can check the policy state by `self.training` and `self.updating`. Please refer to [States for policy](#) for more detailed explanation.

**Warning:** If you use `torch.distributions.Normal` and `torch.distributions.Categorical` to calculate the `log_prob`, please be careful about the shape: Categorical distribution gives “[batch\_size]” shape while Normal distribution gives “[batch\_size, 1]” shape. The auto-broadcasting of numerical operation with torch tensors will amplify this error.

**training:** bool

### 1.11.5 Multi-agent

**class** `tianshou.policy.MultiAgentPolicyManager` (*policies: List[tianshou.policy.base.BasePolicy], env: None, \*\*kwargs: Any*)

Bases: `tianshou.policy.base.BasePolicy`

Multi-agent policy manager for MARL.

This multi-agent policy manager accepts a list of [BasePolicy](#). It dispatches the batch data to each of these policies when the “forward” is called. The same as “process\_fn” and “learn”: it splits the data and feeds them to each policy. A figure in [Multi-Agent Reinforcement Learning](#) can help you better understand this procedure.

**replace\_policy** (*policy: tianshou.policy.base.BasePolicy, agent\_id: int*) → None

Replace the “agent\_id”th policy in this manager.

**process\_fn** (*batch: tianshou.data.batch.Batch, buffer: tianshou.data.buffer.base.ReplayBuffer, indice: numpy.ndarray*) → `tianshou.data.batch.Batch`

Dispatch batch data from `obs.agent_id` to every policy’s `process_fn`.

Save original multi-dimensional `rew` in “`save_rew`”, set `rew` to the reward of each agent during their “`process_fn`”, and restore the original reward afterwards.

**exploration\_noise** (*act: Union[numpy.ndarray, tianshou.data.batch.Batch], batch: tianshou.data.batch.Batch*) → `Union[numpy.ndarray, tianshou.data.batch.Batch]`

Add exploration noise from sub-policy onto `act`.

**forward** (*batch: tianshou.data.batch.Batch, state: Optional[Union[dict, tianshou.data.batch.Batch]] = None, \*\*kwargs: Any*) → `tianshou.data.batch.Batch`

Dispatch batch data from `obs.agent_id` to every policy’s `forward`.

**Parameters** `state` – if None, it means all agents have no state. If not None, it should contain keys of “`agent_1`”, “`agent_2`”, ...

**Returns** a Batch with the following contents:

```
{
  "act": actions corresponding to the input
  "state": {
    "agent_1": output state of agent_1's policy for the state
    "agent_2": xxx
    ...
    "agent_n": xxx}
```

(continues on next page)

(continued from previous page)

```

    "out": {
        "agent_1": output of agent_1's policy for the input
        "agent_2": xxx
        ...
        "agent_n": xxx}
    }

```

**learn** (*batch*: *tianshou.data.batch.Batch*, *\*\*kwargs*: *Any*) → Dict[str, Union[float, List[float]]]

Dispatch the data to all policies for learning.

**Returns** a dict with the following contents:

```

{
    "agent_1/item1": item 1 of agent_1's policy.learn output
    "agent_1/item2": item 2 of agent_1's policy.learn output
    "agent_2/xxx": xxx
    ...
    "agent_n/xxx": xxx
}

```

**training**: bool

## 1.12 tianshou.trainer

### 1.12.1 On-policy

```

class tianshou.trainer.OnpolicyTrainer (policy: tianshou.policy.base.BasePolicy,
                                             train_collector: tian-
                                             shou.data.collector.Collector, test_collector:
                                             Optional[tianshou.data.collector.Collector],
                                             max_epoch: int, step_per_epoch: int, re-
                                             peat_per_collect: int, episode_per_test:
                                             int, batch_size: int, step_per_collect: Op-
                                             tional[int] = None, episode_per_collect:
                                             Optional[int] = None, train_fn: Op-
                                             tional[Callable[[int, int], None]] = None, test_fn:
                                             Optional[Callable[[int, Optional[int]], None]]
                                             = None, stop_fn: Optional[Callable[[float],
                                             bool]] = None, save_best_fn: Op-
                                             tional[Callable[[tianshou.policy.base.BasePolicy],
                                             None]] = None, save_checkpoint_fn: Op-
                                             tional[Callable[[int, int, int], None]] =
                                             None, resume_from_log: bool = False, re-
                                             ward_metric: Optional[Callable[[numpy.ndarray],
                                             numpy.ndarray]] = None, logger: tian-
                                             shou.utils.logger.base.BaseLogger = <tian-
                                             shou.utils.logger.base.LazyLogger object>, verbose:
                                             bool = True, test_in_train: bool = True, **kwargs:
                                             Any)

```

Bases: *tianshou.trainer.base.BaseTrainer*

An iterator class for onpolicy trainer procedure.

Returns an iterator that yields a 3-tuple (epoch, stats, info) of train results on every epoch.

The “step” in onpolicy trainer means an environment step (a.k.a. transition).

Example usage:

```
trainer = OnpolicyTrainer(...)
for epoch, epoch_stat, info in trainer:
    print("Epoch:", epoch)
    print(epoch_stat)
    print(info)
    do_something_with_policy()
    query_something_about_policy()
    make_a_plot_with(epoch_stat)
    display(info)
```

- **epoch** int: the epoch number
- **epoch\_stat** dict: a large collection of metrics of the current epoch
- **info** dict: result returned from `gather_info()`

You can even iterate on several trainers at the same time:

```
trainer1 = OnpolicyTrainer(...)
trainer2 = OnpolicyTrainer(...)
for result1, result2, ... in zip(trainer1, trainer2, ...):
    compare_results(result1, result2, ...)
```

## Parameters

- **policy** – an instance of the `BasePolicy` class.
- **train\_collector** (`Collector`) – the collector used for training.
- **test\_collector** (`Collector`) – the collector used for testing. If it’s None, then no testing will be performed.
- **max\_epoch** (`int`) – the maximum number of epochs for training. The training process might be finished before reaching `max_epoch` if `stop_fn` is set.
- **step\_per\_epoch** (`int`) – the number of transitions collected per epoch.
- **repeat\_per\_collect** (`int`) – the number of repeat time for policy learning, for example, set it to 2 means the policy needs to learn each given batch data twice.
- **episode\_per\_test** (`int`) – the number of episodes for one policy evaluation.
- **batch\_size** (`int`) – the batch size of sample data, which is going to feed in the policy network.
- **step\_per\_collect** (`int`) – the number of transitions the collector would collect before the network update, i.e., trainer will collect “step\_per\_collect” transitions and do some policy network update repeatedly in each epoch.
- **episode\_per\_collect** (`int`) – the number of episodes the collector would collect before the network update, i.e., trainer will collect “episode\_per\_collect” episodes and do some policy network update repeatedly in each epoch.
- **train\_fn** (`function`) – a hook called at the beginning of training in each epoch. It can be used to perform custom additional operations, with the signature `f(num_epoch: int, step_idx: int) -> None`.

- **test\_fn** (*function*) – a hook called at the beginning of testing in each epoch. It can be used to perform custom additional operations, with the signature `f(num_epoch: int, step_idx: int) -> None`.
- **save\_best\_fn** (*function*) – a hook called when the undiscounted average mean reward in evaluation phase gets better, with the signature `f(policy: BasePolicy) -> None`. It was `save_fn` previously.
- **save\_checkpoint\_fn** (*function*) – a function to save training process, with the signature `f(epoch: int, env_step: int, gradient_step: int) -> None`; you can save whatever you want.
- **resume\_from\_log** (*bool*) – resume `env_step/gradient_step` and other metadata from existing tensorboard log. Default to `False`.
- **stop\_fn** (*function*) – a function with signature `f(mean_rewards: float) -> bool`, receives the average undiscounted returns of the testing result, returns a boolean which indicates whether reaching the goal.
- **reward\_metric** (*function*) – a function with signature `f(rewards: np.ndarray with shape (num_episode, agent_num)) -> np.ndarray with shape (num_episode,)`, used in multi-agent RL. We need to return a single scalar for each episode's result to monitor training in the multi-agent RL setting. This function specifies what is the desired metric, e.g., the reward of agent 1 or the average reward over all agents.
- **logger** (`BaseLogger`) – A logger that logs statistics during training/testing/updating. Default to a logger that doesn't log anything.
- **verbose** (*bool*) – whether to print the information. Default to `True`.
- **test\_in\_train** (*bool*) – whether to test in the training phase. Default to `True`.

---

**Note:** Only either one of `step_per_collect` and `episode_per_collect` can be specified.

---

**policy\_update\_fn** (*data: Dict[str, Any], result: Optional[Dict[str, Any]] = None*) → None  
Perform one on-policy update.

`tianshou.trainer.onpolicy_trainer` (*\*args, \*\*kwargs*) → `Dict[str, Union[float, str]]`  
Wrapper for `OnpolicyTrainer` run method.

It is identical to `OnpolicyTrainer(...).run()`.

**Returns** See `gather_info()`.

`tianshou.trainer.onpolicy_trainer_iter`  
alias of `tianshou.trainer.onpolicy.OnpolicyTrainer`



### 1.12.2 Off-policy

```
class tianshou.trainer.OffpolicyTrainer (policy:          tianshou.policy.base.BasePolicy,
                                         train_collector:      tianshou.data.collector.Collector,
                                         test_collector:      Optional[tianshou.data.collector.Collector],
                                         max_epoch:          int,
                                         step_per_epoch:      int,
                                         step_per_collect:    int,
                                         episode_per_test:    int,
                                         batch_size:          int,
                                         update_per_step:      Union[int, float] = 1,
                                         train_fn:             Optional[Callable[[int, int], None]] = None,
                                         test_fn:             Optional[Callable[[int, Optional[int]], None]] = None,
                                         stop_fn:             Optional[Callable[[float, bool]] = None,
                                         save_best_fn:         Optional[Callable[[tianshou.policy.base.BasePolicy], None]] = None,
                                         save_checkpoint_fn:   Optional[Callable[[int, int, int], None]] = None,
                                         resume_from_log:      bool = False,
                                         reward_metric:        Optional[Callable[[numpy.ndarray], numpy.ndarray]] = None,
                                         logger:               tianshou.utils.logger.base.BaseLogger = <tianshou.utils.logger.base.LazyLogger object>,
                                         verbose:             bool = True,
                                         test_in_train:        bool = True,
                                         **kwargs: Any)
```

Bases: `tianshou.trainer.base.BaseTrainer`

An iterator class for offpolicy trainer procedure.

Returns an iterator that yields a 3-tuple (epoch, stats, info) of train results on every epoch.

The “step” in offpolicy trainer means an environment step (a.k.a. transition).

Example usage:

```
trainer = OffpolicyTrainer(...)
for epoch, epoch_stat, info in trainer:
    print("Epoch:", epoch)
    print(epoch_stat)
    print(info)
    do_something_with_policy()
    query_something_about_policy()
    make_a_plot_with(epoch_stat)
    display(info)
```

- epoch int: the epoch number
- epoch\_stat dict: a large collection of metrics of the current epoch
- info dict: result returned from `gather_info()`

You can even iterate on several trainers at the same time:

```
trainer1 = OffpolicyTrainer(...)
trainer2 = OffpolicyTrainer(...)
for result1, result2, ... in zip(trainer1, trainer2, ...):
    compare_results(result1, result2, ...)
```

#### Parameters

- **policy** – an instance of the `BasePolicy` class.
- **train\_collector** (`Collector`) – the collector used for training.
- **test\_collector** (`Collector`) – the collector used for testing. If it's `None`, then no testing will be performed.
- **max\_epoch** (`int`) – the maximum number of epochs for training. The training process might be finished before reaching `max_epoch` if `stop_fn` is set.
- **step\_per\_epoch** (`int`) – the number of transitions collected per epoch.
- **step\_per\_collect** (`int`) – the number of transitions the collector would collect before the network update, i.e., trainer will collect “step\_per\_collect” transitions and do some policy network update repeatedly in each epoch.
- **episode\_per\_test** – the number of episodes for one policy evaluation.
- **batch\_size** (`int`) – the batch size of sample data, which is going to feed in the policy network.
- **update\_per\_step** (`int/float`) – the number of times the policy network would be updated per transition after (`step_per_collect`) transitions are collected, e.g., if `update_per_step` set to 0.3, and `step_per_collect` is 256, policy will be updated  $\text{round}(256 * 0.3 = 76.8) = 77$  times after 256 transitions are collected by the collector. Default to 1.
- **train\_fn** (`function`) – a hook called at the beginning of training in each epoch. It can be used to perform custom additional operations, with the signature `f(num_epoch: int, step_idx: int) -> None`.
- **test\_fn** (`function`) – a hook called at the beginning of testing in each epoch. It can be used to perform custom additional operations, with the signature `f(num_epoch: int, step_idx: int) -> None`.
- **save\_best\_fn** (`function`) – a hook called when the undiscounted average mean reward in evaluation phase gets better, with the signature `f(policy: BasePolicy) -> None`. It was `save_fn` previously.
- **save\_checkpoint\_fn** (`function`) – a function to save training process, with the signature `f(epoch: int, env_step: int, gradient_step: int) -> None`; you can save whatever you want.
- **resume\_from\_log** (`bool`) – resume `env_step/gradient_step` and other metadata from existing tensorboard log. Default to `False`.
- **stop\_fn** (`function`) – a function with signature `f(mean_rewards: float) -> bool`, receives the average undiscounted returns of the testing result, returns a boolean which indicates whether reaching the goal.
- **reward\_metric** (`function`) – a function with signature `f(rewards: np.ndarray with shape (num_episode, agent_num)) -> np.ndarray with shape (num_episode,)`, used in multi-agent RL. We need to return a single scalar for each episode's result to monitor training in the multi-agent RL setting. This function specifies what is the desired metric, e.g., the reward of agent 1 or the average reward over all agents.
- **logger** (`BaseLogger`) – A logger that logs statistics during training/testing/updating. Default to a logger that doesn't log anything.
- **verbose** (`bool`) – whether to print the information. Default to `True`.
- **test\_in\_train** (`bool`) – whether to test in the training phase. Default to `True`.

**policy\_update\_fn** (*data: Dict[str, Any], result: Dict[str, Any]*) → None  
 Perform off-policy updates.

**tianshou.trainer.offpolicy\_trainer** (\*args, \*\*kwargs) → Dict[str, Union[float, str]]  
 Wrapper for OffPolicyTrainer run method.

It is identical to `OffpolicyTrainer(...).run()`.

**Returns** See `gather_info()`.

**tianshou.trainer.offpolicy\_trainer\_iter**  
 alias of `tianshou.trainer.offpolicy.OffpolicyTrainer`

### 1.12.3 Offline

**class** `tianshou.trainer.OfflineTrainer` (*policy: tianshou.policy.base.BasePolicy, buffer: tianshou.data.buffer.base.ReplayBuffer, test\_collector: Optional[tianshou.data.collector.Collector], max\_epoch: int, update\_per\_epoch: int, episode\_per\_test: int, batch\_size: int, test\_fn: Optional[Callable[[int, Optional[int]], None]] = None, stop\_fn: Optional[Callable[[float, bool]] = None, save\_best\_fn: Optional[Callable[[tianshou.policy.base.BasePolicy], None]] = None, save\_checkpoint\_fn: Optional[Callable[[int, int, int], None]] = None, resume\_from\_log: bool = False, reward\_metric: Optional[Callable[[numpy.ndarray], numpy.ndarray]] = None, logger: tianshou.utils.logger.base.BaseLogger = <tianshou.utils.logger.base.LazyLogger object>, verbose: bool = True, \*\*kwargs: Any*)

Bases: `tianshou.trainer.base.BaseTrainer`

An iterator class for offline trainer procedure.

Returns an iterator that yields a 3-tuple (epoch, stats, info) of train results on every epoch.

The “step” in offline trainer means a gradient step.

Example usage:

```
trainer = OfflineTrainer(...)
for epoch, epoch_stat, info in trainer:
    print("Epoch:", epoch)
    print(epoch_stat)
    print(info)
    do_something_with_policy()
    query_something_about_policy()
    make_a_plot_with(epoch_stat)
    display(info)
```

- epoch int: the epoch number
- epoch\_stat dict: a large collection of metrics of the current epoch
- info dict: result returned from `gather_info()`

You can even iterate on several trainers at the same time:

```

trainer1 = OfflineTrainer(...)
trainer2 = OfflineTrainer(...)
for result1, result2, ... in zip(trainer1, trainer2, ...):
    compare_results(result1, result2, ...)

```

### Parameters

- **policy** – an instance of the *BasePolicy* class.
- **buffer** – an instance of the *ReplayBuffer* class. This buffer must be populated with experiences for offline RL.
- **test\_collector** (*Collector*) – the collector used for testing. If it's None, then no testing will be performed.
- **max\_epoch** (*int*) – the maximum number of epochs for training. The training process might be finished before reaching max\_epoch if stop\_fn is set.
- **update\_per\_epoch** (*int*) – the number of policy network updates, so-called gradient steps, per epoch.
- **episode\_per\_test** – the number of episodes for one policy evaluation.
- **batch\_size** (*int*) – the batch size of sample data, which is going to feed in the policy network.
- **test\_fn** (*function*) – a hook called at the beginning of testing in each epoch. It can be used to perform custom additional operations, with the signature `f(num_epoch: int, step_idx: int) -> None`.
- **save\_best\_fn** (*function*) – a hook called when the undiscounted average mean reward in evaluation phase gets better, with the signature `f(policy: BasePolicy) -> None`. It was `save_fn` previously.
- **save\_checkpoint\_fn** (*function*) – a function to save training process, with the signature `f(epoch: int, env_step: int, gradient_step: int) -> None`; you can save whatever you want. Because offline-RL doesn't have env\_step, the env\_step is always 0 here.
- **resume\_from\_log** (*bool*) – resume gradient\_step and other metadata from existing tensorboard log. Default to False.
- **stop\_fn** (*function*) – a function with signature `f(mean_rewards: float) -> bool`, receives the average undiscounted returns of the testing result, returns a boolean which indicates whether reaching the goal.
- **reward\_metric** (*function*) – a function with signature `f(rewards: np.ndarray with shape (num_episode, agent_num)) -> np.ndarray with shape (num_episode,)`, used in multi-agent RL. We need to return a single scalar for each episode's result to monitor training in the multi-agent RL setting. This function specifies what is the desired metric, e.g., the reward of agent 1 or the average reward over all agents.
- **logger** (*BaseLogger*) – A logger that logs statistics during updating/testing. Default to a logger that doesn't log anything.
- **verbose** (*bool*) – whether to print the information. Default to True.

**policy\_update\_fn** (*data: Dict[str, Any], result: Optional[Dict[str, Any]] = None*) → None  
 Perform one off-line policy update.

`tianshou.trainer.offline_trainer(*args, **kwargs) → Dict[str, Union[float, str]]`

Wrapper for `offline_trainer` run method.

It is identical to `OfflineTrainer(...).run()`.

**Returns** See `gather_info()`.

`tianshou.trainer.offline_trainer_iter`

alias of `tianshou.trainer.offline.OfflineTrainer`

### 1.12.4 utils

`tianshou.trainer.test_episode(policy: tianshou.policy.base.BasePolicy, collector: tianshou.data.collector.Collector, test_fn: Optional[Callable[[int, Optional[int]], None]], epoch: int, n_episode: int, logger: Optional[tianshou.utils.logger.base.BaseLogger] = None, global_step: Optional[int] = None, reward_metric: Optional[Callable[[numpy.ndarray], numpy.ndarray]] = None) → Dict[str, Any]`

A simple wrapper of testing policy in collector.

`tianshou.trainer.gather_info(start_time: float, train_collector: Optional[tianshou.data.collector.Collector], test_collector: Optional[tianshou.data.collector.Collector], best_reward: float, best_reward_std: float) → Dict[str, Union[float, str]]`

A simple wrapper of gathering information from collectors.

#### Returns

A dictionary with the following keys:

- `train_step` the total collected step of training collector;
- `train_episode` the total collected episode of training collector;
- `train_time/collector` the time for collecting transitions in the training collector;
- `train_time/model` the time for training models;
- `train_speed` the speed of training (env\_step per second);
- `test_step` the total collected step of test collector;
- `test_episode` the total collected episode of test collector;
- `test_time` the time for testing;
- `test_speed` the speed of testing (env\_step per second);
- `best_reward` the best reward over the test results;
- `duration` the total elapsed time.

## 1.13 tianshou.exploration

**class** tianshou.exploration.BaseNoise

Bases: abc.ABC, object

The action noise base class.

**reset** () → None

Reset to the initial state.

**abstract** **\_\_call\_\_** (size: Sequence[int]) → numpy.ndarray

Generate new noise.

**class** tianshou.exploration.GaussianNoise (mu: float = 0.0, sigma: float = 1.0)

Bases: tianshou.exploration.random.BaseNoise

The vanilla Gaussian process, for exploration in DDPG by default.

**\_\_call\_\_** (size: Sequence[int]) → numpy.ndarray

Generate new noise.

**class** tianshou.exploration.OUNoise (mu: float = 0.0, sigma: float = 0.3, theta: float = 0.15, dt: float = 0.01, x0: Optional[Union[float, numpy.ndarray]] = None)

Bases: tianshou.exploration.random.BaseNoise

Class for Ornstein-Uhlenbeck process, as used for exploration in DDPG.

Usage:

```
# init
self.noise = OUNoise()
# generate noise
noise = self.noise(logits.shape, eps)
```

For required parameters, you can refer to the [stackoverflow page](#). However, our experiment result shows that (similar to OpenAI SpinningUp) using vanilla Gaussian process has little difference from using the Ornstein-Uhlenbeck process.

**reset** () → None

Reset to the initial state.

**\_\_call\_\_** (size: Sequence[int], mu: Optional[float] = None) → numpy.ndarray

Generate new noise.

Return an numpy array which size is equal to size.

## 1.14 tianshou.utils

Utils package.

**class** tianshou.utils.MovAvg (size: int = 100)

Bases: object

Class for moving average.

It will automatically exclude the infinity and NaN. Usage:

```

>>> stat = MovAvg(size=66)
>>> stat.add(torch.tensor(5))
5.0
>>> stat.add(float('inf')) # which will not add to stat
5.0
>>> stat.add([6, 7, 8])
6.5
>>> stat.get()
6.5
>>> print(f'{stat.mean():.2f}±{stat.std():.2f}')
6.50±1.12

```

**add** (*data\_array*: Union[numbers.Number, numpy.number, list, numpy.ndarray, torch.Tensor]) → float  
Add a scalar into *MovAvg*.

You can add `torch.Tensor` with only one element, a python scalar, or a list of python scalar.

**get** () → float  
Get the average.

**mean** () → float  
Get the average. Same as *get* ().

**std** () → float  
Get the standard deviation.

**class** tianshou.utils.**RunningMeanStd** (*mean*: Union[float, numpy.ndarray] = 0.0, *std*: Union[float, numpy.ndarray] = 1.0)

Bases: object

Calculates the running mean and std of a data stream.

[https://en.wikipedia.org/wiki/Algorithms\\_for\\_calculating\\_variance#Parallel\\_algorithm](https://en.wikipedia.org/wiki/Algorithms_for_calculating_variance#Parallel_algorithm)

**update** (*data\_array*: numpy.ndarray) → None  
Add a batch of item into RMS with the same shape, modify mean/var/count.

**class** tianshou.utils.**BaseLogger** (*train\_interval*: int = 1000, *test\_interval*: int = 1, *update\_interval*: int = 1000)

Bases: abc.ABC

The base class for any logger which is compatible with trainer.

Try to overwrite `write()` method to use your own writer.

#### Parameters

- **train\_interval** (*int*) – the log interval in `log_train_data()`. Default to 1000.
- **test\_interval** (*int*) – the log interval in `log_test_data()`. Default to 1.
- **update\_interval** (*int*) – the log interval in `log_update_data()`. Default to 1000.

**abstract write** (*step\_type*: str, *step*: int, *data*: Dict[str, Union[int, numbers.Number, numpy.number, numpy.ndarray]]) → None

Specify how the writer is used to log data.

#### Parameters

- **step\_type** (*str*) – namespace which the data dict belongs to.
- **step** (*int*) – stands for the ordinate of the data dict.
- **data** (*dict*) – the data to write with format `{key: value}`.

**log\_train\_data** (*collect\_result: dict, step: int*) → None

Use writer to log statistics generated during training.

**Parameters**

- **collect\_result** – a dict containing information of data collected in training stage, i.e., returns of collector.collect().
- **step** (*int*) – stands for the timestep the collect\_result being logged.

**log\_test\_data** (*collect\_result: dict, step: int*) → None

Use writer to log statistics generated during evaluating.

**Parameters**

- **collect\_result** – a dict containing information of data collected in evaluating stage, i.e., returns of collector.collect().
- **step** (*int*) – stands for the timestep the collect\_result being logged.

**log\_update\_data** (*update\_result: dict, step: int*) → None

Use writer to log statistics generated during updating.

**Parameters**

- **update\_result** – a dict containing information of data collected in updating stage, i.e., returns of policy.update().
- **step** (*int*) – stands for the timestep the collect\_result being logged.

**save\_data** (*epoch: int, env\_step: int, gradient\_step: int, save\_checkpoint\_fn: Optional[Callable[[int, int, int], None]] = None*) → None

Use writer to log metadata when calling save\_checkpoint\_fn in trainer.

**Parameters**

- **epoch** (*int*) – the epoch in trainer.
- **env\_step** (*int*) – the env\_step in trainer.
- **gradient\_step** (*int*) – the gradient\_step in trainer.
- **save\_checkpoint\_fn** (*function*) – a hook defined by user, see trainer documentation for detail.

**restore\_data** () → Tuple[int, int, int]

Return the metadata from existing log.

If it finds nothing or an error occurs during the recover process, it will return the default parameters.

**Returns** epoch, env\_step, gradient\_step.

```
class tianshou.utils.TensorboardLogger (writer: torch.utils.tensorboard.writer.SummaryWriter,  
                                         train_interval: int = 1000, test_interval: int = 1,  
                                         update_interval: int = 1000, save_interval: int = 1,  
                                         write_flush: bool = True)
```

Bases: tianshou.utils.logger.base.BaseLogger

A logger that relies on tensorboard SummaryWriter by default to visualize and log statistics.

**Parameters**

- **writer** (*SummaryWriter*) – the writer to log data.
- **train\_interval** (*int*) – the log interval in log\_train\_data(). Default to 1000.
- **test\_interval** (*int*) – the log interval in log\_test\_data(). Default to 1.



- **update\_interval** (*int*) – the log interval in `log_update_data()`. Default to 1000.
- **save\_interval** (*int*) – the save interval in `save_data()`. Default to 1 (save at the end of each epoch).
- **write\_flush** (*bool*) – whether to flush tensorboard result after each `add_scalar` operation. Default to True.

**write** (*step\_type: str, step: int, data: Dict[str, Union[int, numbers.Number, numpy.number, numpy.ndarray]]*) → None  
Specify how the writer is used to log data.

#### Parameters

- **step\_type** (*str*) – namespace which the data dict belongs to.
- **step** (*int*) – stands for the ordinate of the data dict.
- **data** (*dict*) – the data to write with format {key: value}.

**save\_data** (*epoch: int, env\_step: int, gradient\_step: int, save\_checkpoint\_fn: Optional[Callable[[int, int, int], None]] = None*) → None  
Use writer to log metadata when calling `save_checkpoint_fn` in trainer.

#### Parameters

- **epoch** (*int*) – the epoch in trainer.
- **env\_step** (*int*) – the `env_step` in trainer.
- **gradient\_step** (*int*) – the `gradient_step` in trainer.
- **save\_checkpoint\_fn** (*function*) – a hook defined by user, see trainer documentation for detail.

**restore\_data** () → Tuple[int, int, int]  
Return the metadata from existing log.

If it finds nothing or an error occurs during the recover process, it will return the default parameters.

**Returns** epoch, env\_step, gradient\_step.

**class** tianshou.utils.**BasicLogger** (*\*args: Any, \*\*kwargs: Any*)  
Bases: tianshou.utils.logger.tensorboard.TensorboardLogger  
BasicLogger has changed its name to TensorboardLogger in #427.  
This class is for compatibility.

**class** tianshou.utils.**LazyLogger**  
Bases: tianshou.utils.logger.base.BaseLogger  
A logger that does nothing. Used as the placeholder in trainer.

**write** (*step\_type: str, step: int, data: Dict[str, Union[int, numbers.Number, numpy.number, numpy.ndarray]]*) → None  
The LazyLogger writes nothing.

**class** tianshou.utils.**WandbLogger** (*train\_interval: int = 1000, test\_interval: int = 1, update\_interval: int = 1000, save\_interval: int = 1000, write\_flush: bool = True, project: Optional[str] = None, name: Optional[str] = None, entity: Optional[str] = None, run\_id: Optional[str] = None, config: Optional[argparse.Namespace] = None*)  
Bases: tianshou.utils.logger.base.BaseLogger  
Weights and Biases logger that sends data to <https://wandb.ai/>.

This logger creates three panels with plots: train, test, and update. Make sure to select the correct access for each panel in weights and biases:

Example of usage:

```
logger = WandbLogger()
logger.load(SummaryWriter(log_path))
result = onpolicy_trainer(policy, train_collector, test_collector,
                          logger=logger)
```

### Parameters

- **train\_interval** (*int*) – the log interval in `log_train_data()`. Default to 1000.
- **test\_interval** (*int*) – the log interval in `log_test_data()`. Default to 1.
- **update\_interval** (*int*) – the log interval in `log_update_data()`. Default to 1000.
- **save\_interval** (*int*) – the save interval in `save_data()`. Default to 1 (save at the end of each epoch).
- **write\_flush** (*bool*) – whether to flush tensorboard result after each `add_scalar` operation. Default to True.
- **project** (*str*) – W&B project name. Default to “tianshou”.
- **name** (*str*) – W&B run name. Default to None. If None, random name is assigned.
- **entity** (*str*) – W&B team/organization name. Default to None.
- **run\_id** (*str*) – run id of W&B run to be resumed. Default to None.
- **config** (*argparse.Namespace*) – experiment configurations. Default to None.

**load** (*writer: torch.utils.tensorboard.writer.SummaryWriter*) → None

**write** (*step\_type: str, step: int, data: Dict[str, Union[int, numbers.Number, numpy.number, numpy.ndarray]]*) → None

Specify how the writer is used to log data.

### Parameters

- **step\_type** (*str*) – namespace which the data dict belongs to.
- **step** (*int*) – stands for the ordinate of the data dict.
- **data** (*dict*) – the data to write with format `{key: value}`.

**save\_data** (*epoch: int, env\_step: int, gradient\_step: int, save\_checkpoint\_fn: Optional[Callable[[int, int, int], None]] = None*) → None

Use writer to log metadata when calling `save_checkpoint_fn` in trainer.

### Parameters

- **epoch** (*int*) – the epoch in trainer.
- **env\_step** (*int*) – the env\_step in trainer.
- **gradient\_step** (*int*) – the gradient\_step in trainer.
- **save\_checkpoint\_fn** (*function*) – a hook defined by user, see trainer documentation for detail.

**restore\_data** () → `Tuple[int, int, int]`

Return the metadata from existing log.

If it finds nothing or an error occurs during the recover process, it will return the default parameters.

**Returns** epoch, env\_step, gradient\_step.

tianshou.utils.deprecation(msg: str) → None

Deprecation warning wrapper.

**class** tianshou.utils.MultipleLRSchedulers(\*args: torch.optim.lr\_scheduler.LambdaLR)

Bases: object

A wrapper for multiple learning rate schedulers.

Every time `step()` is called, it calls the `step()` method of each of the schedulers that it contains. Example usage:

```
scheduler1 = ConstantLR(opt1, factor=0.1, total_iters=2)
scheduler2 = ExponentialLR(opt2, gamma=0.9)
scheduler = MultipleLRSchedulers(scheduler1, scheduler2)
policy = PPOPolicy(..., lr_scheduler=scheduler)
```

**step()** → None

Take a step in each of the learning rate schedulers.

**state\_dict()** → List[Dict]

Get state\_dict for each of the learning rate schedulers.

**Returns** A list of state\_dict of learning rate schedulers.

**load\_state\_dict**(state\_dict: List[Dict]) → None

Load states from state\_dict.

**Parameters** **state\_dict** (List[Dict]) – A list of learning rate scheduler state\_dict, in the same order as the schedulers.

## 1.14.1 Pre-defined Networks

### Common

tianshou.utils.net.common.miniblock(input\_size: int, output\_size: int = 0, norm\_layer: Optional[Type[torch.nn.modules.module.Module]] = None, activation: Optional[Type[torch.nn.modules.module.Module]] = None, linear\_layer: Type[torch.nn.modules.linear.Linear] = <class 'torch.nn.modules.linear.Linear'>) → List[torch.nn.modules.module.Module]

Construct a miniblock with given input/output-size, norm layer and activation.

**class** tianshou.utils.net.common.MLP(input\_dim: int, output\_dim: int = 0, hidden\_sizes: Sequence[int] = (), norm\_layer: Optional[Union[Type[torch.nn.modules.module.Module], Sequence[Type[torch.nn.modules.module.Module]]]] = None, activation: Optional[Union[Type[torch.nn.modules.module.Module], Sequence[Type[torch.nn.modules.module.Module]]]] = <class 'torch.nn.modules.activation.ReLU'>, device: Optional[Union[str, int, torch.device]] = None, linear\_layer: Type[torch.nn.modules.linear.Linear] = <class 'torch.nn.modules.linear.Linear'>, flatten\_input: bool = True)

Bases: `torch.nn.modules.module.Module`

Simple MLP backbone.

Create a MLP of size `input_dim * hidden_sizes[0] * hidden_sizes[1] * ... * hidden_sizes[-1] * output_dim`

#### Parameters

- **input\_dim** (*int*) – dimension of the input vector.
- **output\_dim** (*int*) – dimension of the output vector. If set to 0, there is no final linear layer.
- **hidden\_sizes** – shape of MLP passed in as a list, not including `input_dim` and `output_dim`.
- **norm\_layer** – use which normalization before activation, e.g., `nn.LayerNorm` and `nn.BatchNorm1d`. Default to no normalization. You can also pass a list of normalization modules with the same length of `hidden_sizes`, to use different normalization module in different layers. Default to no normalization.
- **activation** – which activation to use after each layer, can be both the same activation for all layers if passed in `nn.Module`, or different activation for different Modules if passed in a list. Default to `nn.ReLU`.
- **device** – which device to create this model on. Default to `None`.
- **linear\_layer** – use this module as linear layer. Default to `nn.Linear`.
- **flatten\_input** (*bool*) – whether to flatten input data. Default to `True`.

**forward** (*obs: Union[numpy.ndarray, torch.Tensor]*) → `torch.Tensor`

Defines the computation performed at every call.

Should be overridden by all subclasses.

---

**Note:** Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

---

**training:** `bool`

```
class tianshou.utils.net.common.Net (state_shape: Union[int, Sequence[int]], ac-
tion_shape: Union[int, Sequence[int]] = 0, hid-
den_sizes: Sequence[int] = (), norm_layer: Op-
tional[Type[torch.nn.modules.module.Module]]
= None, activation: Op-
tional[Type[torch.nn.modules.module.Module]] = <class
'torch.nn.modules.activation.ReLU'>, device: Union[str,
int, torch.device] = 'cpu', softmax: bool = False, concat:
bool = False, num_atoms: int = 1, dueling_param:
Optional[Tuple[Dict[str, Any], Dict[str, Any]]] = None,
linear_layer: Type[torch.nn.modules.linear.Linear] =
<class 'torch.nn.modules.linear.Linear'>)
```

Bases: `torch.nn.modules.module.Module`

Wrapper of MLP to support more specific DRL usage.

For advanced usage (how to customize the network), please refer to [Build the Network](#).

#### Parameters

- **state\_shape** – int or a sequence of int of the shape of state.
- **action\_shape** – int or a sequence of int of the shape of action.
- **hidden\_sizes** – shape of MLP passed in as a list.
- **norm\_layer** – use which normalization before activation, e.g., `nn.LayerNorm` and `nn.BatchNorm1d`. Default to no normalization. You can also pass a list of normalization modules with the same length of `hidden_sizes`, to use different normalization module in different layers. Default to no normalization.
- **activation** – which activation to use after each layer, can be both the same activation for all layers if passed in `nn.Module`, or different activation for different Modules if passed in a list. Default to `nn.ReLU`.
- **device** – specify the device when the network actually runs. Default to “cpu”.
- **softmax** (*bool*) – whether to apply a softmax layer over the last layer’s output.
- **concat** (*bool*) – whether the input shape is concatenated by `state_shape` and `action_shape`. If it is True, `action_shape` is not the output shape, but affects the input shape only.
- **num\_atoms** (*int*) – in order to expand to the net of distributional RL. Default to 1 (not use).
- **dueling\_param** (*bool*) – whether to use dueling network to calculate Q values (for Dueling DQN). If you want to use dueling option, you should pass a tuple of two dict (first for Q and second for V) stating self-defined arguments as stated in `class:~tianshou.utils.net.common.MLP`. Default to None.
- **linear\_layer** – use this module as linear layer. Default to `nn.Linear`.

See also:

Please refer to [MLP](#) for more detailed explanation on the usage of activation, `norm_layer`, etc.

You can also refer to [Actor](#), [Critic](#), etc, to see how it’s suggested be used.

**forward** (*obs: Union[numpy.ndarray, torch.Tensor], state: Optional[Any] = None, info: Dict[str, Any] = {}*) → Tuple[torch.Tensor, Any]  
Mapping: obs -> flatten (inside MLP)-> logits.

**training:** bool

```
class tianshou.utils.net.common.Recurrent (layer_num: int, state_shape: Union[int,
Sequence[int]], action_shape: Union[int,
Sequence[int]], device: Union[str, int,
torch.device] = 'cpu', hidden_layer_size:
int = 128)
```

Bases: `torch.nn.modules.module.Module`

Simple Recurrent network based on LSTM.

For advanced usage (how to customize the network), please refer to [Build the Network](#).

**forward** (*obs: Union[numpy.ndarray, torch.Tensor], state: Optional[Dict[str, torch.Tensor]] = None, info: Dict[str, Any] = {}*) → Tuple[torch.Tensor, Dict[str, torch.Tensor]]  
Mapping: obs -> flatten -> logits.

In the evaluation mode, *obs* should be with shape `[bsz, dim]`; in the training mode, *obs* should be with shape `[bsz, len, dim]`. See the code and comment for more detail.

**training:** bool

```
class tianshou.utils.net.common.ActorCritic (actor: torch.nn.modules.module.Module,  
                                             critic: torch.nn.modules.module.Module)
```

Bases: torch.nn.modules.module.Module

An actor-critic network for parsing parameters.

Using `actor_critic.parameters()` instead of `set.union` or `list+list` to avoid issue #449.

**Parameters**

- **actor** (*nn.Module*) – the actor network.
- **critic** (*nn.Module*) – the critic network.

**training:** bool

```
class tianshou.utils.net.common.DataParallelNet (net: torch.nn.modules.module.Module)
```

Bases: torch.nn.modules.module.Module

DataParallel wrapper for training agent with multi-GPU.

This class does only the conversion of input data type, from numpy array to torch's Tensor. If the input is a nested dictionary, the user should create a similar class to do the same thing.

**Parameters** **net** (*nn.Module*) – the network to be distributed in different GPUs.

**forward** (*obs: Union[numpy.ndarray, torch.Tensor], \*args: Any, \*\*kwargs: Any*) → Tuple[Any, Any]

Defines the computation performed at every call.

Should be overridden by all subclasses.

---

**Note:** Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

---

**training:** bool

```
class tianshou.utils.net.common.EnsembleLinear (ensemble_size: int, in_feature: int,  
                                              out_feature: int, bias: bool = True)
```

Bases: torch.nn.modules.module.Module

Linear Layer of Ensemble network.

**Parameters**

- **ensemble\_size** (*int*) – Number of subnets in the ensemble.
- **inp\_feature** (*int*) – dimension of the input vector.
- **out\_feature** (*int*) – dimension of the output vector.
- **bias** (*bool*) – whether to include an additive bias, default to be True.

**training:** bool

**forward** (*x: torch.Tensor*) → torch.Tensor

Defines the computation performed at every call.

Should be overridden by all subclasses.

---

**Note:** Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

---

## Discrete

```
class tianshou.utils.net.discrete.Actor (preprocess_net: torch.nn.modules.module.Module,
                                         action_shape: Sequence[int], hidden_sizes: Sequence[int] = (), softmax_output: bool = True, preprocess_net_output_dim: Optional[int] = None,
                                         device: Union[str, int, torch.device] = 'cpu')
```

Bases: torch.nn.modules.module.Module

Simple actor network.

Will create an actor operated in discrete action space with structure of preprocess\_net → action\_shape.

### Parameters

- **preprocess\_net** – a self-defined preprocess\_net which output a flattened hidden state.
- **action\_shape** – a sequence of int for the shape of action.
- **hidden\_sizes** – a sequence of int for constructing the MLP after preprocess\_net. Default to empty sequence (where the MLP now contains only a single linear layer).
- **softmax\_output** (*bool*) – whether to apply a softmax layer over the last layer's output.
- **preprocess\_net\_output\_dim** (*int*) – the output dimension of preprocess\_net.

For advanced usage (how to customize the network), please refer to [Build the Network](#).

### See also:

Please refer to [Net](#) as an instance of how preprocess\_net is suggested to be defined.

```
forward (obs: Union[numpy.ndarray, torch.Tensor], state: Optional[Any] = None, info: Dict[str, Any] = {}) → Tuple[torch.Tensor, Any]
Mapping: s → Q(s, *).
```

```
training: bool
```

```
class tianshou.utils.net.discrete.Critic (preprocess_net: torch.nn.modules.module.Module,
                                         hidden_sizes: Sequence[int] = (), last_size: int = 1, preprocess_net_output_dim: Optional[int] = None, device: Union[str, int, torch.device] = 'cpu')
```

Bases: torch.nn.modules.module.Module

Simple critic network. Will create an actor operated in discrete action space with structure of preprocess\_net → 1(q value).

### Parameters

- **preprocess\_net** – a self-defined preprocess\_net which output a flattened hidden state.
- **hidden\_sizes** – a sequence of int for constructing the MLP after preprocess\_net. Default to empty sequence (where the MLP now contains only a single linear layer).
- **last\_size** (*int*) – the output dimension of Critic network. Default to 1.
- **preprocess\_net\_output\_dim** (*int*) – the output dimension of preprocess\_net.

For advanced usage (how to customize the network), please refer to [Build the Network](#).

### See also:

Please refer to [Net](#) as an instance of how preprocess\_net is suggested to be defined.

**forward** (*obs*: Union[numpy.ndarray, torch.Tensor], *\*\*kwargs*: Any) → torch.Tensor  
Mapping:  $s \rightarrow V(s)$ .

**training**: bool

**class** tianshou.utils.net.discrete.CosineEmbeddingNetwork (*num\_cosines*: int, *embedding\_dim*: int)

Bases: torch.nn.modules.module.Module

Cosine embedding network for IQN. Convert a scalar in [0, 1] to a list of n-dim vectors.

#### Parameters

- **num\_cosines** – the number of cosines used for the embedding.
- **embedding\_dim** – the dimension of the embedding/output.

---

**Note:** From [https://github.com/ku2482/fqf-iqn-qrdqn.pytorch/blob/master/fqf\\_iqn\\_qrdqn/network.py](https://github.com/ku2482/fqf-iqn-qrdqn.pytorch/blob/master/fqf_iqn_qrdqn/network.py).

---

**forward** (*taus*: torch.Tensor) → torch.Tensor  
Defines the computation performed at every call.

Should be overridden by all subclasses.

---

**Note:** Although the recipe for forward pass needs to be defined within this function, one should call the Module instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

---

**training**: bool

**class** tianshou.utils.net.discrete.ImplicitQuantileNetwork (*preprocess\_net*:  
torch.nn.modules.module.Module,  
*action\_shape*: Sequence[int],  
*hidden\_sizes*: Sequence[int],  
*num\_cosines*: int = 64,  
*preprocess\_net\_output\_dim*: Optional[int] = None,  
*device*: Union[str, torch.device] = 'cpu')

Bases: tianshou.utils.net.discrete.Critic

Implicit Quantile Network.

#### Parameters

- **preprocess\_net** – a self-defined preprocess\_net which output a flattened hidden state.
- **action\_shape** (*int*) – a sequence of int for the shape of action.
- **hidden\_sizes** – a sequence of int for constructing the MLP after preprocess\_net. Default to empty sequence (where the MLP now contains only a single linear layer).
- **num\_cosines** (*int*) – the number of cosines to use for cosine embedding. Default to 64.
- **preprocess\_net\_output\_dim** (*int*) – the output dimension of preprocess\_net.



---

**Note:** Although this class inherits Critic, it is actually a quantile Q-Network with output shape (batch\_size, action\_dim, sample\_size).

The second item of the first return value is tau vector.

---

**forward** (*obs*: Union[numpy.ndarray, torch.Tensor], *sample\_size*: int, *\*\*kwargs*: Any) → Tuple[Any, torch.Tensor]  
 Mapping:  $s \rightarrow Q(s, *)$ .

**training**: bool

**class** tianshou.utils.net.discrete.FractionProposalNetwork (*num\_fractions*: int, *embedding\_dim*: int)

Bases: torch.nn.modules.module.Module

Fraction proposal network for FQF.

#### Parameters

- **num\_fractions** – the number of factions to propose.
- **embedding\_dim** – the dimension of the embedding/input.

---

**Note:** Adapted from [https://github.com/ku2482/fqf-ign-qrdqn.pytorch/blob/master/fqf\\_ign\\_qrdqn/network.py](https://github.com/ku2482/fqf-ign-qrdqn.pytorch/blob/master/fqf_ign_qrdqn/network.py).

---

**forward** (*obs\_embeddings*: torch.Tensor) → Tuple[torch.Tensor, torch.Tensor, torch.Tensor]

Defines the computation performed at every call.

Should be overridden by all subclasses.

---

**Note:** Although the recipe for forward pass needs to be defined within this function, one should call the Module instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

---

**training**: bool

**class** tianshou.utils.net.discrete.FullQuantileFunction (*preprocess\_net*: torch.nn.modules.module.Module, *action\_shape*: Sequence[int], *hidden\_sizes*: Sequence[int] = (), *num\_cosines*: int = 64, *preprocess\_net\_output\_dim*: Optional[int] = None, *device*: Union[str, int, torch.device] = 'cpu')

Bases: tianshou.utils.net.discrete.ImplicitQuantileNetwork

Full(y parameterized) Quantile Function.

#### Parameters

- **preprocess\_net** – a self-defined preprocess\_net which output a flattened hidden state.
- **action\_shape** (*int*) – a sequence of int for the shape of action.

- **hidden\_sizes** – a sequence of int for constructing the MLP after preprocess\_net. Default to empty sequence (where the MLP now contains only a single linear layer).
- **num\_cosines** (*int*) – the number of cosines to use for cosine embedding. Default to 64.
- **preprocess\_net\_output\_dim** (*int*) – the output dimension of preprocess\_net.

---

**Note:** The first return value is a tuple of (quantiles, fractions, quantiles\_tau), where fractions is a Batch(taus, tau\_hats, entropies).

---

**forward** (*obs*: Union[numpy.ndarray, torch.Tensor], *propose\_model*: tianshou.utils.net.discrete.FractionProposalNetwork, *fractions*: Optional[tianshou.data.batch.Batch] = None, *\*\*kwargs*: Any) → Tuple[Any, torch.Tensor]  
Mapping:  $s \rightarrow Q(s, *)$ .

**training:** bool

**class** tianshou.utils.net.discrete.NoisyLinear (*in\_features*: int, *out\_features*: int, *noisy\_std*: float = 0.5)

Bases: torch.nn.modules.module.Module

Implementation of Noisy Networks. arXiv:1706.10295.

#### Parameters

- **in\_features** (*int*) – the number of input features.
- **out\_features** (*int*) – the number of output features.
- **noisy\_std** (*float*) – initial standard deviation of noisy linear layers.

---

**Note:** Adapted from [https://github.com/ku2482/fqf-ign-qrdqn.pytorch/blob/master/fqf\\_ign\\_qrdqn/network.py](https://github.com/ku2482/fqf-ign-qrdqn.pytorch/blob/master/fqf_ign_qrdqn/network.py).

---

**reset** () → None

**f** (*x*: torch.Tensor) → torch.Tensor

**sample** () → None

**forward** (*x*: torch.Tensor) → torch.Tensor

Defines the computation performed at every call.

Should be overridden by all subclasses.

---

**Note:** Although the recipe for forward pass needs to be defined within this function, one should call the Module instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

---

**training:** bool

tianshou.utils.net.discrete.sample\_noise (*model*: torch.nn.modules.module.Module) → bool

Sample the random noises of NoisyLinear modules in the model.

**Parameters** **model** – a PyTorch module which may have NoisyLinear submodules.

**Returns** True if model has at least one NoisyLinear submodule; otherwise, False.

```
class tianshou.utils.net.discrete.IntrinsicCuriosityModule (feature_net:
                                                                torch.nn.modules.module.Module,
                                                                feature_dim:      int,
                                                                action_dim:      int,
                                                                hidden_sizes:    Sequence[int] = (),
                                                                device:          Union[str,
                                                                torch.device] = 'cpu')
```

Bases: torch.nn.modules.module.Module

Implementation of Intrinsic Curiosity Module. arXiv:1705.05363.

#### Parameters

- **feature\_net** (*torch.nn.Module*) – a self-defined feature\_net which output a flattened hidden state.
- **feature\_dim** (*int*) – input dimension of the feature net.
- **action\_dim** (*int*) – dimension of the action space.
- **hidden\_sizes** – hidden layer sizes for forward and inverse models.
- **device** – device for the module.

**forward** (*s1: Union[numpy.ndarray, torch.Tensor], act: Union[numpy.ndarray, torch.Tensor], s2: Union[numpy.ndarray, torch.Tensor], \*\*kwargs: Any*) → Tuple[torch.Tensor, torch.Tensor]  
Mapping: s1, act, s2 -> mse\_loss, act\_hat.

**training:** bool

## Continuous

```
class tianshou.utils.net.continuous.Actor (preprocess_net:
                                                                torch.nn.modules.module.Module,
                                                                action_shape:    Sequence[int],
                                                                hidden_sizes:    Sequence[int] = (),
                                                                max_action:      float = 1.0,
                                                                device:          Union[str, int, torch.device] = 'cpu',
                                                                preprocess_net_output_dim: Optional[int] = None)
```

Bases: torch.nn.modules.module.Module

Simple actor network. Will create an actor operated in continuous action space with structure of preprocess\_net → action\_shape.

#### Parameters

- **preprocess\_net** – a self-defined preprocess\_net which output a flattened hidden state.
- **action\_shape** – a sequence of int for the shape of action.
- **hidden\_sizes** – a sequence of int for constructing the MLP after preprocess\_net. Default to empty sequence (where the MLP now contains only a single linear layer).
- **max\_action** (*float*) – the scale for the final action logits. Default to 1.
- **preprocess\_net\_output\_dim** (*int*) – the output dimension of preprocess\_net.

For advanced usage (how to customize the network), please refer to [Build the Network](#).

#### See also:

Please refer to [Net](#) as an instance of how preprocess\_net is suggested to be defined.

```
forward (obs: Union[numpy.ndarray, torch.Tensor], state: Optional[Any] = None, info: Dict[str, Any] = {}) → Tuple[torch.Tensor, Any]  
Mapping: obs -> logits -> action.
```

```
training: bool
```

```
class tianshou.utils.net.continuous.Critic (preprocess_net:  
    torch.nn.modules.module.Module,      hid-  
    den_sizes: Sequence[int] = (), de-  
    vice: Union[str, int, torch.device]  
    = 'cpu', preprocess_net_output_dim:  
    Optional[int] = None, linear_layer:  
    Type[torch.nn.modules.linear.Linear] =  
    <class 'torch.nn.modules.linear.Linear'>,  
    flatten_input: bool = True)
```

Bases: torch.nn.modules.module.Module

Simple critic network. Will create an actor operated in continuous action space with structure of preprocess\_net → 1(q value).

#### Parameters

- **preprocess\_net** – a self-defined preprocess\_net which output a flattened hidden state.
- **hidden\_sizes** – a sequence of int for constructing the MLP after preprocess\_net. Default to empty sequence (where the MLP now contains only a single linear layer).
- **preprocess\_net\_output\_dim** (*int*) – the output dimension of preprocess\_net.
- **linear\_layer** – use this module as linear layer. Default to nn.Linear.
- **flatten\_input** (*bool*) – whether to flatten input data for the last layer. Default to True.

For advanced usage (how to customize the network), please refer to [Build the Network](#).

#### See also:

Please refer to [Net](#) as an instance of how preprocess\_net is suggested to be defined.

```
forward (obs: Union[numpy.ndarray, torch.Tensor], act: Optional[Union[numpy.ndarray,  
    torch.Tensor]] = None, info: Dict[str, Any] = {}) → torch.Tensor  
Mapping: (s, a) -> logits -> Q(s, a).
```

```
training: bool
```

```
class tianshou.utils.net.continuous.ActorProb (preprocess_net:  
    torch.nn.modules.module.Module, ac-  
    tion_shape: Sequence[int], hidden_sizes:  
    Sequence[int] = (), max_action: float =  
    1.0, device: Union[str, int, torch.device]  
    = 'cpu', unbounded: bool = False,  
    conditioned_sigma: bool = False, pre-  
    process_net_output_dim: Optional[int] =  
    None)
```

Bases: torch.nn.modules.module.Module

Simple actor network (output with a Gauss distribution).

#### Parameters

- **preprocess\_net** – a self-defined preprocess\_net which output a flattened hidden state.
- **action\_shape** – a sequence of int for the shape of action.

- **hidden\_sizes** – a sequence of int for constructing the MLP after preprocess\_net. Default to empty sequence (where the MLP now contains only a single linear layer).
- **max\_action** (*float*) – the scale for the final action logits. Default to 1.
- **unbounded** (*bool*) – whether to apply tanh activation on final logits. Default to False.
- **conditioned\_sigma** (*bool*) – True when sigma is calculated from the input, False when sigma is an independent parameter. Default to False.
- **preprocess\_net\_output\_dim** (*int*) – the output dimension of preprocess\_net.

For advanced usage (how to customize the network), please refer to [Build the Network](#).

See also:

Please refer to [Net](#) as an instance of how preprocess\_net is suggested to be defined.

**forward** (*obs: Union[numpy.ndarray, torch.Tensor], state: Optional[Any] = None, info: Dict[str, Any] = {}*) → Tuple[Tuple[torch.Tensor, torch.Tensor], Any]  
Mapping: obs → logits → (mu, sigma).

**training: bool**

```
class tianshou.utils.net.continuous.RecurrentActorProb (layer_num: int, state_shape:
                                                         Sequence[int], action_shape: Sequence[int],
                                                         hidden_layer_size: int = 128, max_action: float = 1.0, device: Union[str, int, torch.device] = 'cpu',
                                                         unbounded: bool = False, conditioned_sigma: bool = False)
```

Bases: torch.nn.modules.module.Module

Recurrent version of ActorProb.

For advanced usage (how to customize the network), please refer to [Build the Network](#).

**forward** (*obs: Union[numpy.ndarray, torch.Tensor], state: Optional[Dict[str, torch.Tensor]] = None, info: Dict[str, Any] = {}*) → Tuple[Tuple[torch.Tensor, torch.Tensor], Dict[str, torch.Tensor]]  
Almost the same as [Recurrent](#).

**training: bool**

```
class tianshou.utils.net.continuous.RecurrentCritic (layer_num: int, state_shape:
                                                       Sequence[int], action_shape: Sequence[int] = [0], device: Union[str, int, torch.device] = 'cpu', hidden_layer_size: int = 128)
```

Bases: torch.nn.modules.module.Module

Recurrent version of Critic.

For advanced usage (how to customize the network), please refer to [Build the Network](#).

**forward** (*obs: Union[numpy.ndarray, torch.Tensor], act: Optional[Union[numpy.ndarray, torch.Tensor]] = None, info: Dict[str, Any] = {}*) → torch.Tensor  
Almost the same as [Recurrent](#).

**training: bool**

```
class tianshou.utils.net.continuous.Perturbation (preprocess_net:
                                                torch.nn.modules.module.Module,
                                                max_action: float, device: Union[str,
                                                int, torch.device] = 'cpu', phi: float =
                                                0.05)
```

Bases: torch.nn.modules.module.Module

Implementation of perturbation network in BCQ algorithm. Given a state and action, it can generate perturbed action.

#### Parameters

- **preprocess\_net** (*torch.nn.Module*) – a self-defined preprocess\_net which output a flattened hidden state.
- **max\_action** (*float*) – the maximum value of each dimension of action.
- **int, torch.device] device** (*Union[str,)* – which device to create this model on. Default to cpu.
- **phi** (*float*) – max perturbation parameter for BCQ. Default to 0.05.

For advanced usage (how to customize the network), please refer to [Build the Network](#).

#### See also:

You can refer to *examples/offline/offline\_bcq.py* to see how to use it.

**forward** (*state: torch.Tensor, action: torch.Tensor*) → torch.Tensor

Defines the computation performed at every call.

Should be overridden by all subclasses.

---

**Note:** Although the recipe for forward pass needs to be defined within this function, one should call the Module instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

---

**training:** bool

```
class tianshou.utils.net.continuous.VAE (encoder: torch.nn.modules.module.Module, de-
                                          coder: torch.nn.modules.module.Module, hid-
                                          den_dim: int, latent_dim: int, max_action: float,
                                          device: Union[str, torch.device] = 'cpu')
```

Bases: torch.nn.modules.module.Module

Implementation of VAE. It models the distribution of action. Given a state, it can generate actions similar to those in batch. It is used in BCQ algorithm.

#### Parameters

- **encoder** (*torch.nn.Module*) – the encoder in VAE. Its input\_dim must be state\_dim + action\_dim, and output\_dim must be hidden\_dim.
- **decoder** (*torch.nn.Module*) – the decoder in VAE. Its input\_dim must be state\_dim + latent\_dim, and output\_dim must be action\_dim.
- **hidden\_dim** (*int*) – the size of the last linear-layer in encoder.
- **latent\_dim** (*int*) – the size of latent layer.
- **max\_action** (*float*) – the maximum value of each dimension of action.

- **torch.device] device** (*Union[str]*) – which device to create this model on. Default to “cpu”.

For advanced usage (how to customize the network), please refer to *Build the Network*.

#### See also:

You can refer to *examples/offline/offline\_bcq.py* to see how to use it.

**forward** (*state: torch.Tensor, action: torch.Tensor*) → *Tuple[torch.Tensor, torch.Tensor, torch.Tensor]*

Defines the computation performed at every call.

Should be overridden by all subclasses.

---

**Note:** Although the recipe for forward pass needs to be defined within this function, one should call the *Module* instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

---

**decode** (*state: torch.Tensor, latent\_z: Optional[torch.Tensor] = None*) → *torch.Tensor*

**training:** *bool*

## 1.15 Contributing to Tianshou

### 1.15.1 Install Develop Version

To install Tianshou in an “editable” mode, run

```
$ pip install -e ".[dev]"
```

in the main directory. This installation is removable by

```
$ python setup.py develop --uninstall
```

### 1.15.2 PEP8 Code Style Check and Code Formatter

We follow PEP8 python code style with flake8. To check, in the main directory, run:

```
$ make lint
```

We use isort and yapf to format all codes. To format, in the main directory, run:

```
$ make format
```

To check if formatted correctly, in the main directory, run:

```
$ make check-codestyle
```

### 1.15.3 Type Check

We use `mypy` to check the type annotations. To check, in the main directory, run:

```
$ make mypy
```

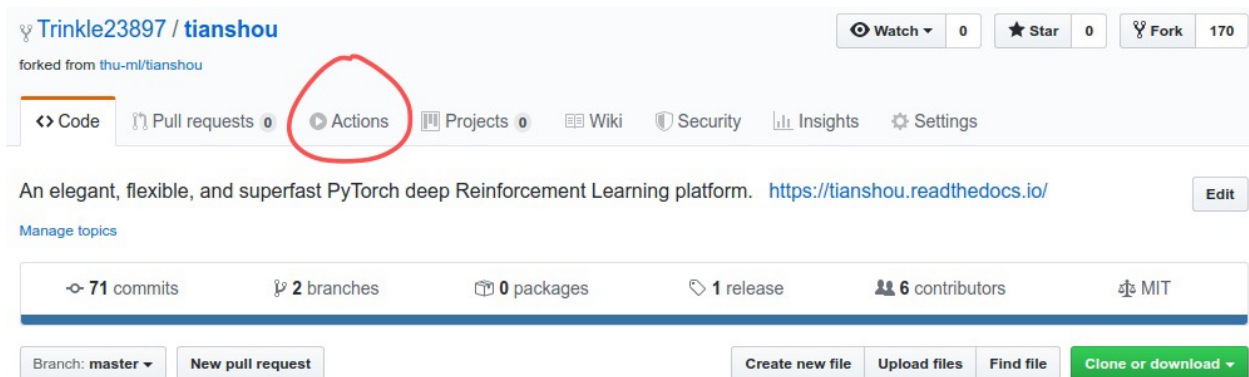
### 1.15.4 Test Locally

This command will run automatic tests in the main directory

```
$ make pytest
```

### 1.15.5 Test by GitHub Actions

1. Click the `Actions` button in your own repo:



2. Click the green button:



#### Workflows aren't being run on this forked repository

Because this repository contained workflow files when it was forked, we have disabled them from running on this fork. Make sure you understand the configured workflows and their expected usage before enabling Actions on this repository.



**I understand my workflows, go ahead and run them**

[View the workflows directory](#)

3. You will see `Actions Enabled` on the top of html page.
4. When you push a new commit to your own repo (e.g. `git push`), it will automatically run the test in this page:



The screenshot shows the GitHub Actions interface. On the left, there's a sidebar with 'Workflows' and 'All workflows' tabs. The 'All workflows' tab is active. Below it, there's a search bar 'Filter workflows' and a list of workflows. One workflow is visible: 'update contributing.md' by 'Trinkle23897'. It shows the event 'Commit d3fb4f9 pushed by Trinkle23897' and the branch 'master'. The status is 'In progress' and it was updated '14 seconds ago'.

### 1.15.6 Documentation

Documentations are written under the `docs/` directory as ReStructuredText (`.rst`) files. `index.rst` is the main page. A Tutorial on ReStructuredText can be found [here](#).

API References are automatically generated by [Sphinx](#) according to the outlines under `docs/api/` and should be modified when any code changes.

To compile documentation into webpage, run

```
$ make doc
```

The generated webpage is in `docs/_build` and can be viewed with browser (<http://0.0.0.0:8000/>).

Chinese documentation is in <https://tianshou.readthedocs.io/zh/latest/>.

### 1.15.7 Documentation Generation Test

We have the following three documentation tests:

1. `pydocstyle`: test all docstring under `tianshou/`;
2. `doc8`: test ReStructuredText format;
3. `sphinx` test: test if there is any error/warning when generating front-end html documentation.

To check, in the main directory, run:

```
$ make check-docstyle
```

## 1.16 Contributor

We always welcome contributions to help make Tianshou better. Below are an incomplete list of our contributors (find more on [this page](#)).

- Jiayi Weng ([Trinkle23897](#))
- Alexis Duburcq ([duburcq](#))
- Kaichao You ([youkaichao](#))
- Huayu Chen ([ChenDRAG](#))



## INDICES AND TABLES

- `genindex`
- `modindex`
- `search`



## BIBLIOGRAPHY

- [MKS+15] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, Andreas Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015. URL: <https://doi.org/10.1038/nature14236>, doi:10.1038/nature14236.
- [LHP+16] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*. 2016. URL: <http://arxiv.org/abs/1509.02971>.
- [SWD+17] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *CoRR*, 2017. URL: <http://arxiv.org/abs/1707.06347>, arXiv:1707.06347.



## PYTHON MODULE INDEX

### t

`tianshou.exploration`, [122](#)  
`tianshou.utils`, [122](#)  
`tianshou.utils.net.common`, [127](#)  
`tianshou.utils.net.continuous`, [135](#)  
`tianshou.utils.net.discrete`, [131](#)





## Symbols

`__call__()` (*tianshou.exploration.BaseNoise method*), 122  
`__call__()` (*tianshou.exploration.GaussianNoise method*), 122  
`__call__()` (*tianshou.exploration.OUNoise method*), 122  
`__getitem__()` (*tianshou.data.Batch method*), 55  
`__getitem__()` (*tianshou.data.PrioritizedReplayBuffer method*), 60  
`__getitem__()` (*tianshou.data.ReplayBuffer method*), 59  
`__len__()` (*tianshou.data.Batch method*), 56  
`__len__()` (*tianshou.data.ReplayBuffer method*), 57  
`__len__()` (*tianshou.data.ReplayBufferManager method*), 60  
`__len__()` (*tianshou.env.BaseVectorEnv method*), 66  
`__setitem__()` (*tianshou.data.Batch method*), 55

## A

A2CPolicy (*class in tianshou.policy*), 86  
 Actor (*class in tianshou.utils.net.continuous*), 135  
 Actor (*class in tianshou.utils.net.discrete*), 131  
 actor (*tianshou.policy.CQLPolicy attribute*), 105  
 actor (*tianshou.policy.DiscreteSACPolicy attribute*), 100  
 actor (*tianshou.policy.REDQPolicy attribute*), 98  
 actor (*tianshou.policy.SACPolicy attribute*), 96  
 actor (*tianshou.policy.TD3Policy attribute*), 94  
 actor\_optim (*tianshou.policy.CQLPolicy attribute*), 105  
 actor\_optim (*tianshou.policy.DiscreteSACPolicy attribute*), 100  
 actor\_optim (*tianshou.policy.REDQPolicy attribute*), 98  
 actor\_optim (*tianshou.policy.SACPolicy attribute*), 96  
 actor\_optim (*tianshou.policy.TD3Policy attribute*), 94  
 actor\_pred() (*tianshou.policy.CQLPolicy method*), 104

ActorCritic (*class in tianshou.utils.net.common*), 129  
 ActorProb (*class in tianshou.utils.net.continuous*), 136  
 add() (*tianshou.data.CachedReplayBuffer method*), 62  
 add() (*tianshou.data.PrioritizedReplayBuffer method*), 59  
 add() (*tianshou.data.ReplayBuffer method*), 58  
 add() (*tianshou.data.ReplayBufferManager method*), 61  
 add() (*tianshou.utils.MovAvg method*), 123  
 AsyncCollector (*class in tianshou.data*), 64

## B

BaseLogger (*class in tianshou.utils*), 123  
 BaseNoise (*class in tianshou.exploration*), 122  
 BasePolicy (*class in tianshou.policy*), 71  
 BaseVectorEnv (*class in tianshou.env*), 65  
 BasicLogger (*class in tianshou.utils*), 125  
 Batch (*class in tianshou.data*), 55  
 BCQPolicy (*class in tianshou.policy*), 101

## C

C51Policy (*class in tianshou.policy*), 78  
 CachedReplayBuffer (*class in tianshou.data*), 62  
 calc\_actor\_loss() (*tianshou.policy.CQLPolicy method*), 104  
 calc\_pi\_values() (*tianshou.policy.CQLPolicy method*), 104  
 calc\_random\_values() (*tianshou.policy.CQLPolicy method*), 104  
 cat() (*tianshou.data.Batch static method*), 55  
 cat\_() (*tianshou.data.Batch method*), 55  
 close() (*tianshou.env.BaseVectorEnv method*), 67  
 close() (*tianshou.env.worker.EnvWorker method*), 69  
 close\_env() (*tianshou.env.worker.DummyEnvWorker method*), 70  
 close\_env() (*tianshou.env.worker.EnvWorker method*), 69  
 close\_env() (*tianshou.env.worker.RayEnvWorker method*), 71  
 close\_env() (*tianshou.env.worker.SubprocEnvWorker method*), 70

- `collect()` (*tianshou.data.AsyncCollector method*), 65  
`collect()` (*tianshou.data.Collector method*), 64  
`Collector` (*class in tianshou.data*), 63  
`compute_episodic_return()` (*tianshou.policy.BasePolicy static method*), 74  
`compute_nstep_return()` (*tianshou.policy.BasePolicy static method*), 75  
`compute_q_value()` (*tianshou.policy.C51Policy method*), 78  
`compute_q_value()` (*tianshou.policy.DQNPolicy method*), 76  
`compute_q_value()` (*tianshou.policy.QRDQNPolicy method*), 80  
`CosineEmbeddingNetwork` (*class in tianshou.utils.net.discrete*), 132  
`CQLPolicy` (*class in tianshou.policy*), 103  
`Critic` (*class in tianshou.utils.net.continuous*), 136  
`Critic` (*class in tianshou.utils.net.discrete*), 131  
`critic` (*tianshou.policy.CQLPolicy attribute*), 105  
`critic` (*tianshou.policy.DiscreteSACPolicy attribute*), 100  
`critic` (*tianshou.policy.REDQPolicy attribute*), 98  
`critic` (*tianshou.policy.SACPolicy attribute*), 96  
`critic` (*tianshou.policy.TD3Policy attribute*), 94  
`critic_optim` (*tianshou.policy.CQLPolicy attribute*), 105  
`critic_optim` (*tianshou.policy.DiscreteSACPolicy attribute*), 100  
`critic_optim` (*tianshou.policy.REDQPolicy attribute*), 98  
`critic_optim` (*tianshou.policy.SACPolicy attribute*), 96  
`critic_optim` (*tianshou.policy.TD3Policy attribute*), 94
- ## D
- `DataParallelNet` (*class in tianshou.utils.net.common*), 130  
`DDPGPolicy` (*class in tianshou.policy*), 91  
`decode()` (*tianshou.utils.net.continuous.VAE method*), 139  
`deprecation()` (*in module tianshou.utils*), 127  
`disc()` (*tianshou.policy.GAILPolicy method*), 110  
`DiscreteBCQPolicy` (*class in tianshou.policy*), 105  
`DiscreteCQLPolicy` (*class in tianshou.policy*), 106  
`DiscreteCRRPolicy` (*class in tianshou.policy*), 107  
`DiscreteSACPolicy` (*class in tianshou.policy*), 98  
`DQNPolicy` (*class in tianshou.policy*), 76  
`DummyEnvWorker` (*class in tianshou.env.worker*), 69  
`DummyVectorEnv` (*class in tianshou.env*), 68
- ## E
- `empty()` (*tianshou.data.Batch static method*), 56  
`empty_()` (*tianshou.data.Batch method*), 56  
`EnsembleLinear` (*class in tianshou.utils.net.common*), 130  
`EnvWorker` (*class in tianshou.env.worker*), 69  
`exploration_noise()` (*tianshou.policy.BasePolicy method*), 72  
`exploration_noise()` (*tianshou.policy.DDPGPolicy method*), 93  
`exploration_noise()` (*tianshou.policy.DiscreteSACPolicy method*), 100  
`exploration_noise()` (*tianshou.policy.DQNPolicy method*), 77  
`exploration_noise()` (*tianshou.policy.ICMPolicy method*), 112  
`exploration_noise()` (*tianshou.policy.MultiAgentPolicyManager method*), 113
- ## F
- `f()` (*tianshou.utils.net.discrete.NoisyLinear method*), 134  
`forward()` (*tianshou.policy.BasePolicy method*), 72  
`forward()` (*tianshou.policy.BCQPolicy method*), 102  
`forward()` (*tianshou.policy.DDPGPolicy method*), 92  
`forward()` (*tianshou.policy.DiscreteBCQPolicy method*), 106  
`forward()` (*tianshou.policy.DiscreteSACPolicy method*), 99  
`forward()` (*tianshou.policy.DQNPolicy method*), 76  
`forward()` (*tianshou.policy.FQFPolicy method*), 83  
`forward()` (*tianshou.policy.ICMPolicy method*), 112  
`forward()` (*tianshou.policy.ImitationPolicy method*), 100  
`forward()` (*tianshou.policy.IQNPolicy method*), 81  
`forward()` (*tianshou.policy.MultiAgentPolicyManager method*), 113  
`forward()` (*tianshou.policy.PGPolicy method*), 84  
`forward()` (*tianshou.policy.PSRLPolicy method*), 111  
`forward()` (*tianshou.policy.RandomPolicy method*), 75  
`forward()` (*tianshou.policy.REDQPolicy method*), 98  
`forward()` (*tianshou.policy.SACPolicy method*), 96  
`forward()` (*tianshou.utils.net.common.DataParallelNet method*), 130  
`forward()` (*tianshou.utils.net.common.EnsembleLinear method*), 130  
`forward()` (*tianshou.utils.net.common.MLP method*), 128  
`forward()` (*tianshou.utils.net.common.Net method*), 129  
`forward()` (*tianshou.utils.net.common.Recurrent method*), 129  
`forward()` (*tianshou.utils.net.continuous.Actor method*), 135



- load\_hdf5() (*tianshou.data.ReplayBuffer* class method), 58
- load\_state\_dict() (*tianshou.utils.MultipleLRSchedulers* method), 127
- log\_test\_data() (*tianshou.utils.BaseLogger* method), 124
- log\_train\_data() (*tianshou.utils.BaseLogger* method), 123
- log\_update\_data() (*tianshou.utils.BaseLogger* method), 124
- ## M
- map\_action() (*tianshou.policy.BasePolicy* method), 73
- map\_action\_inverse() (*tianshou.policy.BasePolicy* method), 73
- mean() (*tianshou.utils.MovAvg* method), 123
- miniblock() (in module *tianshou.utils.net.common*), 127
- MLP (class in *tianshou.utils.net.common*), 127
- module
- tianshou.exploration*, 122
  - tianshou.utils*, 122
  - tianshou.utils.net.common*, 127
  - tianshou.utils.net.continuous*, 135
  - tianshou.utils.net.discrete*, 131
- MovAvg (class in *tianshou.utils*), 122
- MultiAgentPolicyManager (class in *tianshou.policy*), 113
- MultipleLRSchedulers (class in *tianshou.utils*), 127
- ## N
- Net (class in *tianshou.utils.net.common*), 128
- next() (*tianshou.data.ReplayBuffer* method), 58
- next() (*tianshou.data.ReplayBufferManager* method), 60
- NoisyLinear (class in *tianshou.utils.net.discrete*), 134
- NPGPolicy (class in *tianshou.policy*), 85
- ## O
- offline\_trainer() (in module *tianshou.trainer*), 120
- offline\_trainer\_iter (in module *tianshou.trainer*), 121
- OfflineTrainer (class in *tianshou.trainer*), 119
- offpolicy\_trainer() (in module *tianshou.trainer*), 119
- offpolicy\_trainer\_iter (in module *tianshou.trainer*), 119
- OffpolicyTrainer (class in *tianshou.trainer*), 117
- onpolicy\_trainer() (in module *tianshou.trainer*), 116
- onpolicy\_trainer\_iter (in module *tianshou.trainer*), 116
- OnpolicyTrainer (class in *tianshou.trainer*), 114
- OUNoise (class in *tianshou.exploration*), 122
- ## P
- Perturbation (class in *tianshou.utils.net.continuous*), 137
- PGPolicy (class in *tianshou.policy*), 84
- policy\_update\_fn() (*tianshou.trainer.OfflineTrainer* method), 120
- policy\_update\_fn() (*tianshou.trainer.OffpolicyTrainer* method), 118
- policy\_update\_fn() (*tianshou.trainer.OnpolicyTrainer* method), 116
- post\_process\_fn() (*tianshou.policy.BasePolicy* method), 73
- post\_process\_fn() (*tianshou.policy.ICMPolicy* method), 112
- PPOPPolicy (class in *tianshou.policy*), 89
- prev() (*tianshou.data.ReplayBuffer* method), 58
- prev() (*tianshou.data.ReplayBufferManager* method), 60
- PrioritizedReplayBuffer (class in *tianshou.data*), 59
- PrioritizedReplayBufferManager (class in *tianshou.data*), 61
- PrioritizedVectorReplayBuffer (class in *tianshou.data*), 62
- process\_fn() (*tianshou.policy.A2CPolicy* method), 87
- process\_fn() (*tianshou.policy.BasePolicy* method), 73
- process\_fn() (*tianshou.policy.CQLPolicy* method), 104
- process\_fn() (*tianshou.policy.DDPGPolicy* method), 92
- process\_fn() (*tianshou.policy.DQNPolicy* method), 76
- process\_fn() (*tianshou.policy.GAILPolicy* method), 110
- process\_fn() (*tianshou.policy.ICMPolicy* method), 112
- process\_fn() (*tianshou.policy.MultiAgentPolicyManager* method), 113
- process\_fn() (*tianshou.policy.NPGPolicy* method), 86
- process\_fn() (*tianshou.policy.PGPolicy* method), 84
- process\_fn() (*tianshou.policy.PPOPPolicy* method), 90
- PSRLPolicy (class in *tianshou.policy*), 111



## Q

QRDQNPoly (class in *tianshou.policy*), 79

## R

RainbowPolicy (class in *tianshou.policy*), 79

RandomPolicy (class in *tianshou.policy*), 75

RayEnvWorker (class in *tianshou.env.worker*), 70

RayVectorEnv (class in *tianshou.env*), 68

Recurrent (class in *tianshou.utils.net.common*), 129

RecurrentActorProb (class in *tianshou.utils.net.continuous*), 137

RecurrentCritic (class in *tianshou.utils.net.continuous*), 137

recv () (*tianshou.env.worker.EnvWorker* method), 69

recv () (*tianshou.env.worker.RayEnvWorker* method), 71

recv () (*tianshou.env.worker.SubprocEnvWorker* method), 70

REDQPolicy (class in *tianshou.policy*), 96

render () (*tianshou.env.BaseVectorEnv* method), 67

render () (*tianshou.env.worker.DummyEnvWorker* method), 70

render () (*tianshou.env.worker.EnvWorker* method), 69

render () (*tianshou.env.worker.RayEnvWorker* method), 71

render () (*tianshou.env.worker.SubprocEnvWorker* method), 70

replace\_policy () (*tianshou.policy.MultiAgentPolicyManager* method), 113

ReplayBuffer (class in *tianshou.data*), 57

ReplayBufferManager (class in *tianshou.data*), 60

reset () (*tianshou.data.Collector* method), 63

reset () (*tianshou.data.ReplayBuffer* method), 58

reset () (*tianshou.data.ReplayBufferManager* method), 60

reset () (*tianshou.env.BaseVectorEnv* method), 67

reset () (*tianshou.env.worker.DummyEnvWorker* method), 69

reset () (*tianshou.env.worker.EnvWorker* method), 69

reset () (*tianshou.env.worker.RayEnvWorker* method), 70

reset () (*tianshou.exploration.BaseNoise* method), 122

reset () (*tianshou.exploration.OUNoise* method), 122

reset () (*tianshou.utils.net.discrete.NoisyLinear* method), 134

reset\_buffer () (*tianshou.data.Collector* method), 63

reset\_env () (*tianshou.data.AsyncCollector* method), 64

reset\_env () (*tianshou.data.Collector* method), 63

reset\_stat () (*tianshou.data.Collector* method), 63

restore\_data () (*tianshou.utils.BaseLogger* method), 124

restore\_data () (*tianshou.utils.TensorboardLogger* method), 125

restore\_data () (*tianshou.utils.WandbLogger* method), 126

RunningMeanStd (class in *tianshou.utils*), 123

## S

SACPolicy (class in *tianshou.policy*), 94

sample () (*tianshou.data.ReplayBuffer* method), 58

sample () (*tianshou.utils.net.discrete.NoisyLinear* method), 134

sample\_indices () (*tianshou.data.PrioritizedReplayBuffer* method), 59

sample\_indices () (*tianshou.data.ReplayBuffer* method), 58

sample\_indices () (*tianshou.data.ReplayBufferManager* method), 61

sample\_noise () (in module *tianshou.utils.net.discrete*), 134

save\_data () (*tianshou.utils.BaseLogger* method), 124

save\_data () (*tianshou.utils.TensorboardLogger* method), 125

save\_data () (*tianshou.utils.WandbLogger* method), 126

save\_hdf5 () (*tianshou.data.ReplayBuffer* method), 58

seed () (*tianshou.env.BaseVectorEnv* method), 67

seed () (*tianshou.env.worker.DummyEnvWorker* method), 70

seed () (*tianshou.env.worker.EnvWorker* method), 69

seed () (*tianshou.env.worker.RayEnvWorker* method), 71

seed () (*tianshou.env.worker.SubprocEnvWorker* method), 70

send () (*tianshou.env.worker.DummyEnvWorker* method), 70

send () (*tianshou.env.worker.EnvWorker* method), 69

send () (*tianshou.env.worker.RayEnvWorker* method), 71

send () (*tianshou.env.worker.SubprocEnvWorker* method), 70

set\_agent\_id () (*tianshou.policy.BasePolicy* method), 72

set\_batch () (*tianshou.data.ReplayBuffer* method), 58

set\_batch () (*tianshou.data.ReplayBufferManager* method), 60

set\_beta () (*tianshou.data.PrioritizedReplayBuffer* method), 60

set\_beta () (*tianshou.data.PrioritizedVectorReplayBuffer* method), 62

`set_env_attr()` (*tianshou.env.BaseVectorEnv method*), 67  
`set_env_attr()` (*tianshou.env.worker.DummyEnvWorker method*), 69  
`set_env_attr()` (*tianshou.env.worker.EnvWorker method*), 69  
`set_env_attr()` (*tianshou.env.worker.RayEnvWorker method*), 70  
`set_env_attr()` (*tianshou.env.worker.SubprocEnvWorker method*), 70  
`set_eps()` (*tianshou.policy.DQNPolicy method*), 76  
`set_eps()` (*tianshou.policy.ICMPolicy method*), 112  
`set_exp_noise()` (*tianshou.policy.DDPGPolicy method*), 92  
`shape()` (*tianshou.data.Batch property*), 57  
`ShmemVectorEnv` (*class in tianshou.env*), 68  
`soft_update()` (*tianshou.policy.BasePolicy method*), 72  
`split()` (*tianshou.data.Batch method*), 57  
`stack()` (*tianshou.data.Batch static method*), 55  
`stack_()` (*tianshou.data.Batch method*), 55  
`state_dict()` (*tianshou.utils.MultipleLRSchedulers method*), 127  
`std()` (*tianshou.utils.MovAvg method*), 123  
`step()` (*tianshou.env.BaseVectorEnv method*), 67  
`step()` (*tianshou.env.worker.EnvWorker method*), 69  
`step()` (*tianshou.utils.MultipleLRSchedulers method*), 127  
`SubprocEnvWorker` (*class in tianshou.env.worker*), 70  
`SubprocVectorEnv` (*class in tianshou.env*), 68  
`sync_weight()` (*tianshou.policy.BCQPolicy method*), 102  
`sync_weight()` (*tianshou.policy.CQLPolicy method*), 104  
`sync_weight()` (*tianshou.policy.DDPGPolicy method*), 92  
`sync_weight()` (*tianshou.policy.DiscreteCRRPolicy method*), 108  
`sync_weight()` (*tianshou.policy.DQNPolicy method*), 76  
`sync_weight()` (*tianshou.policy.REDQPolicy method*), 98  
`sync_weight()` (*tianshou.policy.SACPolicy method*), 96  
`sync_weight()` (*tianshou.policy.TD3Policy method*), 94  
  
**T**  
`TD3Policy` (*class in tianshou.policy*), 93  
`TensorboardLogger` (*class in tianshou.utils*), 124  
`test_episode()` (*in module tianshou.trainer*), 121  
`tianshou.exploration` (*module*), 122  
`tianshou.utils` (*module*), 122  
`tianshou.utils.net.common` (*module*), 127  
`tianshou.utils.net.continuous` (*module*), 135  
`tianshou.utils.net.discrete` (*module*), 131  
`to_numpy()` (*tianshou.data.Batch method*), 55  
`to_torch()` (*tianshou.data.Batch method*), 55  
`train()` (*tianshou.policy.BCQPolicy method*), 102  
`train()` (*tianshou.policy.CQLPolicy method*), 104  
`train()` (*tianshou.policy.DDPGPolicy method*), 92  
`train()` (*tianshou.policy.DiscreteBCQPolicy method*), 105  
`train()` (*tianshou.policy.DQNPolicy method*), 76  
`train()` (*tianshou.policy.ICMPolicy method*), 112  
`train()` (*tianshou.policy.REDQPolicy method*), 98  
`train()` (*tianshou.policy.SACPolicy method*), 96  
`train()` (*tianshou.policy.TD3Policy method*), 94  
`training` (*tianshou.policy.A2CPolicy attribute*), 88  
`training` (*tianshou.policy.BasePolicy attribute*), 75  
`training` (*tianshou.policy.BCQPolicy attribute*), 103  
`training` (*tianshou.policy.C51Policy attribute*), 78  
`training` (*tianshou.policy.CQLPolicy attribute*), 105  
`training` (*tianshou.policy.DDPGPolicy attribute*), 93  
`training` (*tianshou.policy.DiscreteBCQPolicy attribute*), 106  
`training` (*tianshou.policy.DiscreteCQLPolicy attribute*), 107  
`training` (*tianshou.policy.DiscreteCRRPolicy attribute*), 108  
`training` (*tianshou.policy.DiscreteSACPolicy attribute*), 100  
`training` (*tianshou.policy.DQNPolicy attribute*), 78  
`training` (*tianshou.policy.FQFPolicy attribute*), 84  
`training` (*tianshou.policy.GAILPolicy attribute*), 110  
`training` (*tianshou.policy.ICMPolicy attribute*), 113  
`training` (*tianshou.policy.ImitationPolicy attribute*), 101  
`training` (*tianshou.policy.IQNPolicy attribute*), 82  
`training` (*tianshou.policy.MultiAgentPolicyManager attribute*), 114  
`training` (*tianshou.policy.NPGPolicy attribute*), 86  
`training` (*tianshou.policy.PGPolicy attribute*), 85  
`training` (*tianshou.policy.PPOPolicy attribute*), 91  
`training` (*tianshou.policy.PSRLPolicy attribute*), 111  
`training` (*tianshou.policy.QRDQNPolicy attribute*), 80  
`training` (*tianshou.policy.RainbowPolicy attribute*), 79  
`training` (*tianshou.policy.RandomPolicy attribute*), 75

training (*tianshou.policy.REDQPolicy* attribute), 98  
 training (*tianshou.policy.SACPolicy* attribute), 96  
 training (*tianshou.policy.TD3Policy* attribute), 94  
 training (*tianshou.policy.TRPOPolicy* attribute), 89  
 training (*tianshou.utils.net.common.ActorCritic* attribute), 130  
 training (*tianshou.utils.net.common.DataParallelNet* attribute), 130  
 training (*tianshou.utils.net.common.EnsembleLinear* attribute), 130  
 training (*tianshou.utils.net.common.MLP* attribute), 128  
 training (*tianshou.utils.net.common.Net* attribute), 129  
 training (*tianshou.utils.net.common.Recurrent* attribute), 129  
 training (*tianshou.utils.net.continuous.Actor* attribute), 136  
 training (*tianshou.utils.net.continuous.ActorProb* attribute), 137  
 training (*tianshou.utils.net.continuous.Critic* attribute), 136  
 training (*tianshou.utils.net.continuous.Perturbation* attribute), 138  
 training (*tianshou.utils.net.continuous.RecurrentActorProb* attribute), 137  
 training (*tianshou.utils.net.continuous.RecurrentCritic* attribute), 137  
 training (*tianshou.utils.net.continuous.VAE* attribute), 139  
 training (*tianshou.utils.net.discrete.Actor* attribute), 131  
 training (*tianshou.utils.net.discrete.CosineEmbeddingNetwork* attribute), 132  
 training (*tianshou.utils.net.discrete.Critic* attribute), 132  
 training (*tianshou.utils.net.discrete.FractionProposalNetwork* attribute), 133  
 training (*tianshou.utils.net.discrete.FullQuantileFunction* attribute), 134  
 training (*tianshou.utils.net.discrete.ImplicitQuantileNetwork* attribute), 133  
 training (*tianshou.utils.net.discrete.IntrinsicCuriosityModule* attribute), 135  
 training (*tianshou.utils.net.discrete.NoisyLinear* attribute), 134  
 TRPOPolicy (class in *tianshou.policy*), 88  
 update () (*tianshou.data.Batch* method), 56  
 update () (*tianshou.data.PrioritizedReplayBuffer* method), 59  
 update () (*tianshou.data.ReplayBuffer* method), 58  
 update () (*tianshou.data.ReplayBufferManager* method), 60  
 update () (*tianshou.policy.BasePolicy* method), 73  
 update () (*tianshou.utils.RunningMeanStd* method), 123  
 update\_weight () (*tianshou.data.PrioritizedReplayBuffer* method), 60  
**V**  
 VAE (class in *tianshou.utils.net.continuous*), 138  
 value\_mask () (*tianshou.policy.BasePolicy* static method), 74  
 VectorReplayBuffer (class in *tianshou.data*), 61  
**W**  
 wait () (*tianshou.env.worker.DummyEnvWorker* static method), 69  
 wait () (*tianshou.env.worker.EnvWorker* static method), 69  
 wait () (*tianshou.env.worker.RayEnvWorker* static method), 70  
 wait () (*tianshou.env.worker.SubprocEnvWorker* static method), 70  
 WandbLogger (class in *tianshou.utils*), 125  
 write () (*tianshou.utils.BaseLogger* method), 123  
 write () (*tianshou.utils.LazyLogger* method), 125  
 write () (*tianshou.utils.TensorboardLogger* method), 125  
 write () (*tianshou.utils.WandbLogger* method), 126

## U

unfinished\_index () (*tianshou.data.ReplayBuffer* method), 58  
 unfinished\_index () (*tianshou.data.ReplayBufferManager* method), 60