
Tianshou

Release 0.2.4

Tianshou contributors

Jul 10, 2020

TUTORIALS

1 Installation	3
2 Indices and tables	57
Bibliography	59
Python Module Index	61
Index	63

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:

- *PGPolicy* Policy Gradient
- *DQNPolicy* Deep Q-Network
- *DQNPolicy* Double DQN with n-step returns
- *A2CPolicy* Advantage Actor-Critic
- *DDPGPolicy* Deep Deterministic Policy Gradient
- *PPOPolicy* Proximal Policy Optimization
- *TD3Policy* Twin Delayed DDPG
- *SACPolicy* Soft Actor-Critic
- *ImitationPolicy* Imitation Learning
- *PrioritizedReplayBuffer* Prioritized Experience Replay
- *compute_episodic_return()* Generalized Advantage Estimator

Here is Tianshou's other features:

- Elegant framework, using only ~2000 lines of code
- Support parallel environment sampling for all algorithms: *Parallel Sampling*
- Support recurrent state representation in actor network and critic network (RNN-style training for POMDP): *RNN-style Training*
- Support any type of environment state (e.g. a dict, a self-defined class, ...): *User-defined Environment and Different State Representation*
- Support customized training process: *Customize Training Process*
- Support n-step returns estimation *compute_nstep_return()* for all Q-learning based algorithms

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

INSTALLATION

Tianshou is currently hosted on [PyPI](#). You can simply install Tianshou with the following command (with Python \geq 3.6):

```
pip3 install tianshou
```

You can also install with the newest version through GitHub:

```
pip3 install git+https://github.com/thu-ml/tianshou.git@master
```

If you use Anaconda or Miniconda, you can install Tianshou through the following command lines:

```
# create a new virtualenv and install pip, change the env name if you like
conda create -n myenv pip
# activate the environment
conda activate myenv
# install tianshou
pip install tianshou
```

After installation, open your python console and type

```
import tianshou as ts
print(ts.__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/.

1.1 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.

1.1.1 Make an Environment

First of all, you have to make an environment for your agent to interact with. 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 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, depending on the probability distribution on the action.

1.1.2 Setup Multi-environment Wrapper

It is available if you want the original `gym.Env`:

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

Tianshou supports parallel sampling for all algorithms. It provides three types of vectorized environment wrapper: *VectorEnv*, *SubprocVectorEnv*, and *RayVectorEnv*. It can be used as follows:

```
train_envs = ts.env.VectorEnv([lambda: gym.make('CartPole-v0') for _ in range(8)])
test_envs = ts.env.VectorEnv([lambda: gym.make('CartPole-v0') for _ in range(100)])
```

Here, we set up 8 environments in `train_envs` and 100 environments in `test_envs`.

For the demonstration, here we use the second block of codes.

1.1.3 Build the Network

Tianshou supports any user-defined PyTorch networks and optimizers but with the limitation of input and output API. Here is an example code:

```
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))
        ])
    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
```

(continues on next page)

(continued from previous page)

```

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 have a try with those pre-defined 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`, and intermediate result during the 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 [SWD+17], 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).

1.1.4 Setup Policy

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

```

policy = ts.policy.DQNPoly(policy, net, optim,
    discount_factor=0.9, estimation_step=3,
    use_target_network=True, target_update_freq=320)

```

1.1.5 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.

```

train_collector = ts.data.Collector(policy, train_envs, ts.data.
    ↳ReplayBuffer(size=20000))
test_collector = ts.data.Collector(policy, test_envs)

```

1.1.6 Train Policy with a Trainer

Tianshou provides *onpolicy_trainer* and *offpolicy_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=1000, collect_per_step=10,
    episode_per_test=100, batch_size=64,
    train_fn=lambda e: policy.set_eps(0.1),
    test_fn=lambda e: policy.set_eps(0.05),
    stop_fn=lambda x: x >= env.spec.reward_threshold,

```

(continues on next page)

(continued from previous page)

```
writer=None)
print(f'Finished training! Use {result["duration"]}')

```

The meaning of each parameter is as follows:

- `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 step for updating policy network in one epoch;
- `collect_per_step`: The number of frames the collector would collect before the network update. For example, the code above means “collect 10 frames 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 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 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.
- `writer`: See below.

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

```
from torch.utils.tensorboard import SummaryWriter
writer = SummaryWriter('log/dqn')

```

Pass the writer 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.

1.1.7 Save/Load Policy

Since the policy inherits the `torch.nn.Module` class, 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'))
```

1.1.8 Watch the Agent's Performance

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

```
collector = ts.data.Collector(policy, env)
collector.collect(n_episode=1, render=1 / 35)
collector.close()
```

1.1.9 Train a Policy with Customized Codes

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

No problem! Tianshou supports user-defined training code. Here is the usage:

```
# pre-collect 5000 frames with random action before training
policy.set_eps(1)
train_collector.collect(n_step=5000)

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['rew'] >= env.spec.reward_threshold or i % 1000 == 0:
        policy.set_eps(0.05)
        result = test_collector.collect(n_episode=100)
        if result['rew'] >= env.spec.reward_threshold:
            print(f'Finished training! Test mean returns: {result["rew"]}')
            break
        else:
            # back to training eps
            policy.set_eps(0.1)

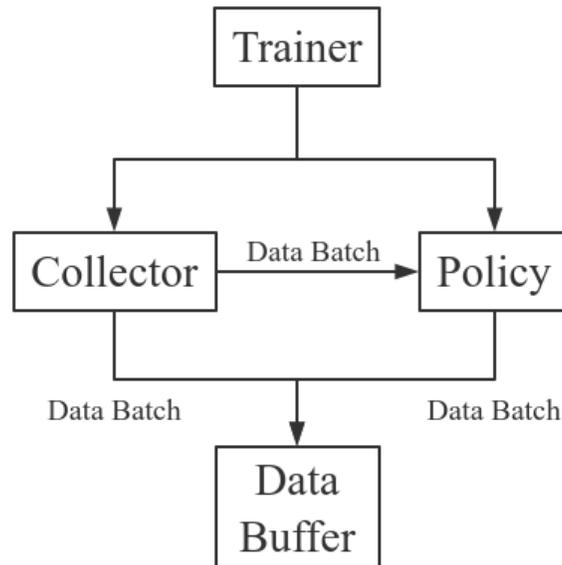
    # train policy with a sampled batch data
    losses = policy.learn(train_collector.sample(batch_size=64))
```

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

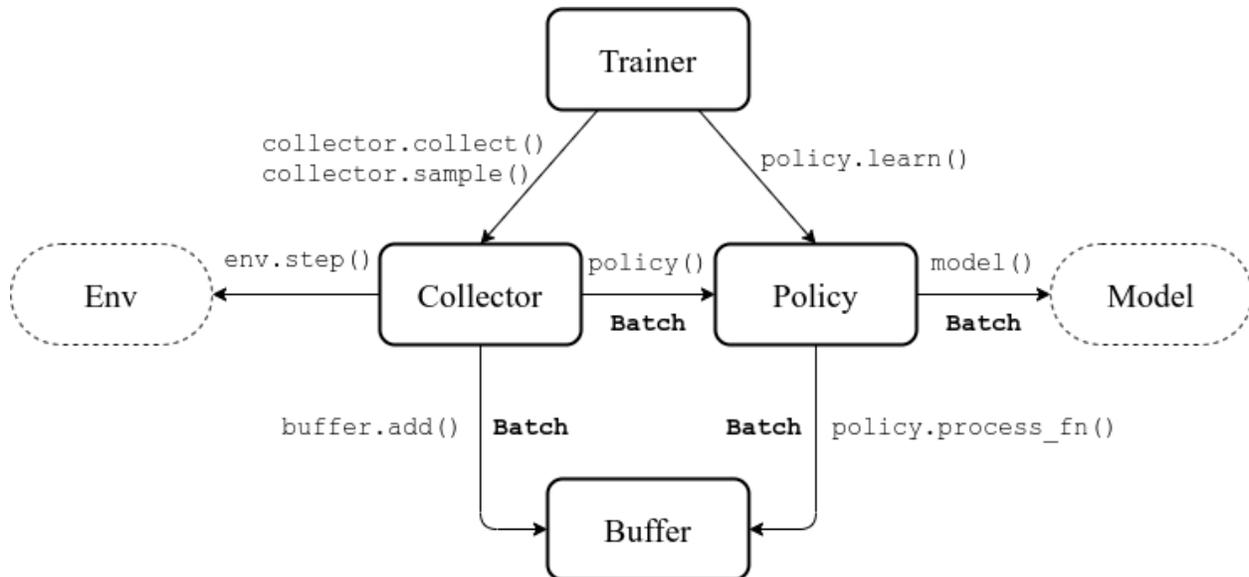
References

1.2 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:



1.2.1 Data 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. Here is the usage:

```
>>> import numpy as np
>>> from tianshou.data import Batch
>>> 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',
)
```

In short, you can define a *Batch* with any key-value pair.

For Numpy arrays, only data types with `np.object`, `bool`, and number are supported. For strings or other data types, however, they can be held in `np.object` arrays.

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 ;

Batch object can be initialized by a wide variety of arguments, ranging from the key/value pairs or dictionary, to list and Numpy arrays of `dict` or *Batch* instances where each element is considered as an individual sample and get stacked together:

```
>>> data = Batch([{'a': {'b': [0.0, "info"]}}])
>>> print(data[0])
Batch(
  a: Batch(
    b: array([0.0, 'info'], dtype=object),
  ),
)
```

Batch has the same API as a native Python `dict`. In this regard, one can access stored data using string key, or iterate over stored data:

```
>>> data = Batch(a=4, b=[5, 5])
>>> print(data["a"])
4
>>> for key, value in data.items():
>>>     print(f"{key}: {value}")
```

(continues on next page)

(continued from previous page)

```
a: 4
b: [5, 5]
```

Batch also partially reproduces the Numpy API for arrays. It also supports the advanced slicing method, such as `batch[:, i]`, if the index is valid. You can access or iterate over the individual samples, if any:

```
>>> data = Batch(a=np.array([[0.0, 2.0], [1.0, 3.0]]), b=[[5, -5]])
>>> print(data[0])
Batch(
  a: array([0., 2.])
  b: array([ 5, -5]),
)
>>> for sample in data:
>>>     print(sample.a)
[0., 2.]

>>> print(data.shape)
[1, 2]
>>> data[:, 1] += 1
>>> print(data)
Batch(
  a: array([[0., 3.],
            [1., 4.]])
  b: array([[ 5, -4]])
)
```

Similarly, one can also perform simple algebra on it, and stack, split or concatenate multiple instances:

```
>>> 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.]])
)
>>> print(np.mean(data))
Batch(
  b: 0.0,
  a: array([0.5, 2.5]),
)
>>> data_split = list(data.split(1, False))
>>> print(list(data.split(1, 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.],
```

(continues on next page)

(continued from previous page)

```

        [1., 3.]],
    )

```

Note that stacking of inconsistent data is also supported. In which case, `None` is added in list or `np.ndarray` of objects, `0` otherwise.

```

>>> 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),
)

```

Method `empty_sets` elements to `0` or `None` for `np.object`.

```

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

```

`shape()` and `__len__()` methods are also provided to respectively get the shape and the length of a `Batch` instance. It mimics the Numpy API for Numpy arrays, which means that getting the length of a scalar `Batch` raises an exception.

```

>>> 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()

```

Convenience helpers are available to convert in-place the stored data into Numpy arrays or Torch tensors.

Finally, note that `Batch` is serializable and therefore Pickle compatible. This is especially important for distributed sampling.

`tianshou.data.Batch.shape`

Return `self.shape`.

1.2.2 Data Buffer

`ReplayBuffer` stores data generated from interaction between the policy and environment. It stores basically 7 types of data, as mentioned in `Batch`, based on `numpy.ndarray`. Here is the usage:

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

>>> buf2 = ReplayBuffer(size=10)
>>> for i in range(15):
...     buf2.add(obs=i, act=i, rew=i, done=i, 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)
array([ 0.,  1.,  2.,  5.,  6.,  7.,  8.,  9., 10., 11., 12., 13., 14.,
        0.,  0.,  0.,  0.,  0.,  0.,  0.])

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

`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
...     buf.add(obs={'id': i}, act=i, rew=i, done=done,
...             obs_next={'id': i + 1})
>>> print(buf) # you can see obs_next is not saved in buf
ReplayBuffer(
  act: array([ 9., 10., 11., 12., 13., 14., 15., 7., 8.]),
  done: array([0., 1., 0., 0., 0., 0., 1., 0., 0.]),
  info: Batch(),
  obs: Batch(
    id: array([ 9., 10., 11., 12., 13., 14., 15., 7., 8.]),
  ),
  policy: Batch(),
  rew: array([ 9., 10., 11., 12., 13., 14., 15., 7., 8.]),
)
>>> index = np.arange(len(buf))
>>> print(buf.get(index, 'obs').id)
```

(continues on next page)

(continued from previous page)

```

[[ 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()
0.0
>>> # we can get obs_next through __getitem__, even if it doesn't exist
>>> print(buf[:].obs_next.id)
[[ 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.]
 [ 7.  7.  7.  8.]
 [ 7.  7.  8.  9.]]

```

param int size the size of replay buffer.

param int stack_num the frame-stack sampling argument, should be greater than 1, defaults to 0 (no stacking).

param bool ignore_obs_next whether to store obs_next, defaults to `False`.

param bool sample_avail the parameter indicating sampling only available index when using frame-stack sampling method, defaults to `False`. This feature is not supported in Prioritized Replay Buffer currently.

Tianshou provides other type of data buffer such as *ListReplayBuffer* (based on list), *PrioritizedReplayBuffer* (based on Segment Tree and `numpy.ndarray`). Check out *ReplayBuffer* for more detail.

1.2.3 Policy

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

A policy class typically has four 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.

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

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

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

```
# pseudocode, cannot work
s = env.reset()
buffer = Buffer(size=10000)
agent = DQN()
for i in range(int(1e6)):
    a = agent.compute_action(s)
    s_, r, d, _ = env.step(a)
    buffer.store(s, a, s_, r, d)
    s = s_
    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, indice):
        buffer_len = len(buffer)
        batch_2 = buffer[(indice + 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[(indice + 2) % buffer_len]
        # in short, buffer.obs[i] is equal to buffer[i].obs, but the former is more_
        ↪efficient.
        Q = self(batch_2, eps=0) # shape: [batchsize, action_shape]
        maxQ = Q.max(dim=-1)
        batch.returns = batch.rew \
            + self._gamma * buffer.rew[(indice + 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](#).

1.2.4 Collector

The *Collector* enables the policy to interact with different types of environments conveniently. In short, *Collector* has two main methods:

- *collect()*: let the policy perform (at least) a specified number of step *n_step* or episode *n_episode* and store the data in the replay buffer;
- *sample()*: sample a data batch from replay buffer; it will call *process_fn()* before returning the final batch data.

Why do we mention **at least** here? For a single environment, the collector will finish exactly *n_step* or *n_episode*. However, for multiple environments, we could not directly store the collected data into the replay buffer, since it breaks the principle of storing data chronologically.

The solution is to add some cache buffers inside the collector. Once collecting a **full episode of trajectory**, it will move the stored data from the cache buffer to the main buffer. To satisfy this condition, the collector will interact with environments that may exceed the given step number or episode number.

The general explanation is listed in *A High-level Explanation*. Other usages of collector are listed in *Collector* documentation.

1.2.5 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 two types of trainer: *onpolicy_trainer()* and *offpolicy_trainer()*, corresponding to on-policy algorithms (such as Policy Gradient) and off-policy algorithms (such as DQN). Please check out *tianshou.trainer* for the usage.

There will be more types of trainers, for instance, multi-agent trainer.

1.2.6 A High-level Explanation

We give a high-level explanation through the pseudocode used in section *Policy*:

```
# pseudocode, cannot work                                     # methods in tianshou
s = env.reset()                                              #
buffer = Buffer(size=10000)                                   # buffer = tianshou.
↪ data.ReplayBuffer(size=10000)
agent = DQN()                                                # policy.__init__(...)
for i in range(int(1e6)):                                     # done in trainer
    a = agent.compute_action(s)                               # policy(batch, ...)
    s_, r, d, _ = env.step(a)                                 # collector.collect(..
↪.)
    buffer.store(s, a, s_, r, d)                              # collector.collect(..
↪.)
    s = s_                                                    # collector.collect(..
↪.)
    if i % 1000 == 0:                                        # done in trainer
        b_s, b_a, b_s_, b_r, b_d = buffer.get(size=64)      # collector.
↪ sample(batch_size)
        # compute 2-step returns. How?
        b_ret = compute_2_step_return(buffer, b_r, b_d, ...) # policy.process_
↪ fn(batch, buffer, indice)
```

(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
↪...)

```

1.2.7 Conclusion

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

1.3 Train a model-free RL agent within 30s

This page summarizes some hyper-parameter tuning experience and code-level trick when training a model-free DRL agent.

You can also contribute to this page with your own tricks :)

1.3.1 Avoid batch-size = 1

In the traditional RL training loop, we always use the policy to interact with only one environment for collecting data. That means most of the time the network use batch-size = 1. Quite inefficient! Here is an example of showing how inefficient it is:

```

import torch, time
from torch import nn

class Net(nn.Module):
    def __init__(self):
        super().__init__()
        self.model = nn.Sequential(
            nn.Linear(3, 128), nn.ReLU(inplace=True),
            nn.Linear(128, 128), nn.ReLU(inplace=True),
            nn.Linear(128, 1))
    def forward(self, s):
        return self.model(s)

net = Net()
cnt = 1000
div = 128
a = torch.randn([128, 3])

t = time.time()
for i in range(cnt):
    b = net(a)
t1 = (time.time() - t) / cnt
print(t1)
t = time.time()
for i in range(cnt):
    for a_ in a.split(a.shape[0] // div):
        b = net(a_)
t2 = (time.time() - t) / cnt
print(t2)
print(t2 / t1)

```

The first test uses batch-size 128, and the second test uses batch-size = 1 for 128 times. In our test, the first is 70-80 times faster than the second.

So how could we avoid the case of batch-size = 1? The answer is synchronize sampling: we create multiple independent environments and sample simultaneously. It is similar to A2C, but other algorithms can also use this method. In our experiments, sampling from more environments benefits not only the sample speed but also the converge speed of neural network (we guess it lowers the sample bias).

By the way, A2C is better than A3C in some cases: A3C needs to act independently and sync the gradient to master, but, in a single node, using A3C to act with batch-size = 1 is quite resource-consuming.

1.3.2 Algorithm specific tricks

Here is about the experience of hyper-parameter tuning on CartPole and Pendulum:

- *DQNPolicy*: use estimation_step greater than 1 and target network, also with a suitable size of replay buffer;
- *PGPolicy*: TBD
- *A2CPolicy*: TBD
- *PPOPolicy*: TBD
- *DDPGPolicy*, *TD3Policy*, and *SACPolicy*: We found two tricks. The first is to ignore the done flag. The second is to normalize reward to a standard normal distribution (it is against the theoretical analysis, but indeed works very well). The two tricks work amazingly on Mujoco tasks, typically with a faster converge speed (1M -> 200K).
- On-policy algorithms: increase the repeat-time (to 2 or 4 for trivial benchmark, 10 for mujoco) of the given batch in each training update will make the algorithm more stable.

1.3.3 Code-level optimization

Tianshou has many short-but-efficient lines of code. For example, when we want to compute $V(s)$ and $V(s')$ by the same network, the best way is to concatenate s and s' together instead of computing the value function using twice of network forward.

1.3.4 Finally

With fast-speed sampling, we could use large batch-size and large learning rate for faster convergence.

RL algorithms are seed-sensitive. Try more seeds and pick the best. But for our demo, we just used seed = 0 and found it work surprisingly well on policy gradient, so we did not try other seed.

1.4 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*.

1.4.1 Build Policy Network

See *Build the Network*.

1.4.2 Build New Policy

See *BasePolicy*.

1.4.3 Customize Training Process

See *Train a Policy with Customized Codes*.

1.4.4 Parallel Sampling

Use *VectorEnv* or *SubprocVectorEnv*.

```
env_fns = [  
    lambda: MyTestEnv(size=2),  
    lambda: MyTestEnv(size=3),  
    lambda: MyTestEnv(size=4),  
    lambda: MyTestEnv(size=5),  
]  
venv = SubprocVectorEnv(env_fns)
```

where `env_fns` is a list of callable env hooker. The above code can be written in for-loop as well:

```
env_fns = [lambda x=i: MyTestEnv(size=x) for i in [2, 3, 4, 5]]  
venv = SubprocVectorEnv(env_fns)
```

1.4.5 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.

This function receives typically 7 keys, as listed in *Batch*, and returns the modified part within a dict or a Batch. For example, you can write your hook as:

```
import numpy as np  
from collections import deque  
class MyProcessor:  
    def __init__(self, size=100):
```

(continues on next page)

(continued from previous page)

```

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 'rew' not in kwargs:
        # means that it is called after env.reset(), it can only process the obs
        return {} # 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'])
    # you can also return with {'rew': kwargs['rew']}

```

And finally,

```

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

```

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

1.4.6 RNN-style Training

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

First, add an argument `stack_num` to `ReplayBuffer`:

```

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

```

Then, change the network to recurrent-style, for example, class `Recurrent` in [code snippet 1](#), or `RecurrentActor` and `RecurrentCritic` in [code snippet 2](#).

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`.

1.4.7 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)` -> None
- `render(mode)` -> None
- `close()` -> None
- `observation_space`
- `action_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]),
 '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 be stored in `ReplayBuffer` as:

```
>>> from tianshou.data import ReplayBuffer
>>> b = ReplayBuffer(size=3)
>>> b.add(obs=e.reset(), act=0, rew=0, done=0)
>>> print(b)
ReplayBuffer(
  act: array([0, 0, 0]),
  done: array([0, 0, 0]),
  info: Batch(),
  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]),
  ),
```

(continues on next page)

(continued from previous page)

```

    policy: Batch(),
    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, indice = 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
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(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(a):
    ...
    return copy.deepcopy(self.graph), reward, done, {}

```

1.5 tianshou.data

```

class tianshou.data.Batch(batch_dict: Optional[Union[dict, Batch, Tuple[Union[dict, Batch]],
List[Union[dict, Batch]], numpy.ndarray]] = None, copy: bool = False,
**kwargs)

```

Bases: object

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. Here is the usage:

```

>>> import numpy as np
>>> from tianshou.data import Batch
>>> 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',
)

```

In short, you can define a *Batch* with any key-value pair.

For Numpy arrays, only data types with `np.object`, `bool`, and `number` are supported. For strings or other data types, however, they can be held in `np.object` arrays.

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 ;

Batch object can be initialized by a wide variety of arguments, ranging from the key/value pairs or dictionary, to list and Numpy arrays of `dict` or *Batch* instances where each element is considered as an individual sample and get stacked together:

```

>>> data = Batch([{'a': {'b': [0.0, "info"]}}])
>>> print(data[0])

```

(continues on next page)

(continued from previous page)

```
Batch(
  a: Batch(
    b: array([0.0, 'info'], dtype=object),
  ),
)
```

Batch has the same API as a native Python dict. In this regard, one can access stored data using string key, or iterate over stored data:

```
>>> data = Batch(a=4, b=[5, 5])
>>> print(data["a"])
4
>>> for key, value in data.items():
>>>     print(f"{key}: {value}")
a: 4
b: [5, 5]
```

Batch also partially reproduces the Numpy API for arrays. It also supports the advanced slicing method, such as `batch[:, i]`, if the index is valid. You can access or iterate over the individual samples, if any:

```
>>> data = Batch(a=np.array([[0.0, 2.0], [1.0, 3.0]]), b=[[5, -5]])
>>> print(data[0])
Batch(
  a: array([0., 2.])
  b: array([ 5, -5]),
)
>>> for sample in data:
>>>     print(sample.a)
[0., 2.]

>>> print(data.shape)
[1, 2]
>>> data[:, 1] += 1
>>> print(data)
Batch(
  a: array([[0., 3.],
            [1., 4.]])
  b: array([[ 5, -4]]),
)
```

Similarly, one can also perform simple algebra on it, and stack, split or concatenate multiple instances:

```
>>> 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.]])
)
>>> print(np.mean(data))
Batch(
  b: 0.0,
  a: array([0.5, 2.5]),
)
```

(continues on next page)

(continued from previous page)

```

>>> data_split = list(data.split(1, False))
>>> print(list(data.split(1, 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.])),
)

```

Note that stacking of inconsistent data is also supported. In which case, `None` is added in `list` or `np.ndarray` of objects, `0` otherwise.

```

>>> 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),
)

```

Method `empty_sets` elements to `0` or `None` for `np.object`.

```

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

```

`shape()` and `__len__()` methods are also provided to respectively get the shape and the length of a `Batch` instance. It mimics the Numpy API for Numpy arrays, which means that getting the length of a scalar `Batch` raises an exception.

```

>>> data = Batch(a=[5., 4.], b=np.zeros((2, 3, 4)))
>>> data.shape

```

(continues on next page)

(continued from previous page)

```
[2]
>>> len(data)
2
>>> data[0].shape
[]
>>> len(data[0])
TypeError: Object of type 'Batch' has no len()
```

Convenience helpers are available to convert in-place the stored data into Numpy arrays or Torch tensors.

Finally, note that `Batch` is serializable and therefore Pickle compatible. This is especially important for distributed sampling.

`__getitem__` (*index*: Union[str, slice, int, numpy.integer, numpy.ndarray, List[int]]) → tianshou.data.batch.Batch
Return self[index].

`__len__` () → int
Return len(self).

`append` (*batch*: tianshou.data.batch.Batch) → None

`static cat` (*batches*: List[Batch]) → tianshou.data.batch.Batch
Concatenate a `Batch` object into a single new batch.

`cat_` (*batch*: tianshou.data.batch.Batch) → None
Concatenate a `Batch` object into current batch.

`static empty` (*batch*: tianshou.data.batch.Batch, *index*: Union[str, slice, int, numpy.integer, 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`.

`empty_` (*index*: Union[str, slice, int, numpy.integer, numpy.ndarray, List[int]] = None) → tianshou.data.batch.Batch
Return an empty a `Batch` object with 0 or None filled. If *index* is specified, it will only reset the specific indexed-data.

`get` (*k*: str, *d*: Optional[Any] = None) → Union[tianshou.data.batch.Batch, Any]
Return self[k] if k in self else d. d defaults to None.

`items` () → List[Tuple[str, Any]]
Return self.items().

`keys` () → List[str]
Return self.keys().

`property shape`
Return self.shape.

`split` (*size*: Optional[int] = None, *shuffle*: bool = True) → Iterator[tianshou.data.batch.Batch]
Split whole data into multiple small batches.

Parameters

- **size** (*int*) – if it is None, it does not split the data batch; otherwise it will divide the data batch with the given size. Default to None.
- **shuffle** (*bool*) – randomly shuffle the entire data batch if it is True, otherwise remain in the same. Default to True.

`static stack` (*batches*: List[Batch], *axis*: int = 0) → tianshou.data.batch.Batch
Stack a `Batch` object into a single new batch.

stack_ (*batches: List[Union[dict, Batch]], axis: int = 0*) → None

Stack a *Batch* object *i* into current batch.

to_numpy () → None

Change all torch.Tensor to numpy.ndarray. This is an in-place operation.

to_torch (*dtype: Optional[torch.dtype] = None, device: Union[str, int, torch.device] = 'cpu'*) → None

Change all numpy.ndarray to torch.Tensor. This is an in-place operation.

values () → List[Any]

Return self.values().

```
class tianshou.data.Collector (policy: tianshou.policy.base.BasePolicy, env: Union[gym.core.Env, tianshou.env.vecenv.BaseVectorEnv], buffer: Optional[Union[tianshou.data.buffer.ReplayBuffer, List[tianshou.data.buffer.ReplayBuffer]]] = None, preprocess_fn: Callable[[Any], Union[dict, tianshou.data.batch.Batch]] = None, stat_size: Optional[int] = 100, action_noise: Optional[tianshou.exploration.random.BaseNoise] = None, **kwargs)
```

Bases: object

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

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, or a list of *ReplayBuffer*. If set to None, it will automatically assign a small-size *ReplayBuffer*.
- **preprocess_fn** (*function*) – a function called before the data has been added to the buffer, see issue #42, defaults to None.
- **stat_size** (*int*) – for the moving average of recording speed, defaults to 100.
- **action_noise** (*BaseNoise*) – add a noise to continuous action. Normally a policy already has a noise param for exploration in training phase, so this is recommended to use in test collector for some purpose.

The *preprocess_fn* is a function called before the data has been added to the buffer with batch format, which receives up to 7 keys as listed in *Batch*. It will receive with only *obs* when the collector resets the environment. It returns either a dict or a *Batch* with the modified keys and values. Examples are in “test/base/test_collector.py”.

Example:

```
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
envs = VectorEnv([lambda: gym.make('CartPole-v0') for _ in range(3)])
buffers = [ReplayBuffer(size=5000) for _ in range(3)]
# you can also pass a list of replay buffer to collector, for multi-env
# collector = Collector(policy, envs, buffer=buffers)
collector = Collector(policy, envs, buffer=replay_buffer)
```

(continues on next page)

(continued from previous page)

```

# collect at least 3 episodes
collector.collect(n_episode=3)
# collect 1 episode for the first env, 3 for the third env
collector.collect(n_episode=[1, 0, 3])
# collect at least 2 steps
collector.collect(n_step=2)
# collect episodes with visual rendering (the render argument is the
# sleep time between rendering consecutive frames)
collector.collect(n_episode=1, render=0.03)

# sample data with a given number of batch-size:
batch_data = collector.sample(batch_size=64)
# policy.learn(batch_data) # btw, vanilla policy gradient only
# supports on-policy training, so here we pick all data in the buffer
batch_data = collector.sample(batch_size=0)
policy.learn(batch_data)
# on-policy algorithms use the collected data only once, so here we
# clear the buffer
collector.reset_buffer()

```

For the scenario of collecting data from multiple environments to a single buffer, the cache buffers will turn on automatically. It may return the data more than the given limitation.

Note: Please make sure the given environment has a time limitation.

close () → None

Close the environment(s).

collect (*n_step*: int = 0, *n_episode*: Union[int, List[int]] = 0, *random*: bool = False, *render*: Optional[float] = None, *log_fn*: Optional[Callable[[dict], None]] = None) → Dict[str, float]
Collect a specified number of step or episode.

Parameters

- **n_step** (*int*) – how many steps you want to collect.
- **n_episode** (*int or list*) – how many episodes you want to collect (in each environment).
- **random** (*bool*) – whether to use random policy for collecting data, defaults to False.
- **render** (*float*) – the sleep time between rendering consecutive frames, defaults to None (no rendering).
- **log_fn** (*function*) – a function which receives env info, typically for tensorboard logging.

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* the collected number of episodes.
- *n/st* the collected number of steps.

- `v/st` the speed of steps per second.
- `v/ep` the speed of episode per second.
- `rew` the mean reward over collected episodes.
- `len` the mean length over collected episodes.

get_env_num () → int

Return the number of environments the collector have.

render (**kwargs) → None

Render all the environment(s).

reset () → None

Reset all related variables in the collector.

reset_buffer () → None

Reset the main data buffer.

reset_env () → None

Reset all of the environment(s)' states and reset all of the cache buffers (if need).

sample (batch_size: int) → `tianshou.data.batch.Batch`

Sample a data batch from the internal replay buffer. It will call `process_fn` () before returning the final batch data.

Parameters `batch_size` (int) – 0 means it will extract all the data from the buffer, otherwise it will extract the data with the given `batch_size`.

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

Reset all the seed(s) of the given environment(s).

class `tianshou.data.ListReplayBuffer` (**kwargs)

Bases: `tianshou.data.buffer.ReplayBuffer`

The function of `ListReplayBuffer` is almost the same as `ReplayBuffer`. The only difference is that `ListReplayBuffer` is based on `list`. Therefore, it does not support advanced indexing, which means you cannot sample a batch of data out of it. It is typically used for storing data.

See also:

Please refer to `ReplayBuffer` for more detailed explanation.

reset () → None

Clear all the data in replay buffer.

sample (batch_size: int) → `Tuple[tianshou.data.batch.Batch, numpy.ndarray]`

Get a random sample from buffer with size equal to `batch_size`. Return all the data in the buffer if `batch_size` is 0.

Returns Sample data and its corresponding index inside the buffer.

class `tianshou.data.PrioritizedReplayBuffer` (size: int, alpha: float, beta: float, mode: str = 'weight', replace: bool = False, **kwargs)

Bases: `tianshou.data.buffer.ReplayBuffer`

Prioritized replay buffer implementation.

Parameters

- **alpha** (float) – the prioritization exponent.
- **beta** (float) – the importance sample soft coefficient.
- **mode** (str) – defaults to `weight`.

- **replace** (*bool*) – whether to sample with replacement

See also:

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

getitem (*index: Union[slice, int, numpy.integer, numpy.ndarray]*) → `tianshou.data.batch.Batch`
 Return a data batch: `self[index]`. If `stack_num` is set to be `> 0`, return the stacked obs and `obs_next` with shape `[batch, len, ...]`.

add (*obs: Union[dict, numpy.ndarray], act: Union[numpy.ndarray, float], rew: Union[int, float], done: bool, obs_next: Optional[Union[dict, numpy.ndarray]] = None, info: dict = {}, policy: Optional[Union[dict, tianshou.data.batch.Batch]] = {}, weight: float = 1.0, **kwargs*) → `None`
 Add a batch of data into replay buffer.

property replace

reset () → `None`

Clear all the data in replay buffer.

sample (*batch_size: int*) → `Tuple[tianshou.data.batch.Batch, numpy.ndarray]`

Get a random sample from buffer with priority probability. Return all the data in the buffer if `batch_size` is 0.

Returns Sample data and its corresponding index inside the buffer.

update_weight (*indice: Union[slice, numpy.ndarray], new_weight: numpy.ndarray*) → `None`

Update priority weight by indice in this buffer.

Parameters

- **indice** (*np.ndarray*) – indice you want to update weight
- **new_weight** (*np.ndarray*) – new priority weight you want to update

class `tianshou.data.ReplayBuffer` (*size: int, stack_num: Optional[int] = 0, ignore_obs_next: bool = False, sample_avail: bool = False, **kwargs*)

Bases: `object`

`ReplayBuffer` stores data generated from interaction between the policy and environment. It stores basically 7 types of data, as mentioned in `Batch`, based on `numpy.ndarray`. Here is the usage:

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

>>> buf2 = ReplayBuffer(size=10)
>>> for i in range(15):
...     buf2.add(obs=i, act=i, rew=i, done=i, 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.])
```

(continues on next page)

(continued from previous page)

```

>>> # move buf2's result into buf (meanwhile keep it chronologically)
>>> buf.update(buf2)
array([[ 0.,  1.,  2.,  5.,  6.,  7.,  8.,  9., 10., 11., 12., 13., 14.,
         0.,  0.,  0.,  0.,  0.,  0.,  0.]])

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

```

`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
...     buf.add(obs={'id': i}, act=i, rew=i, done=done,
...             obs_next={'id': i + 1})
>>> print(buf) # you can see obs_next is not saved in buf
ReplayBuffer(
  act: array([ 9., 10., 11., 12., 13., 14., 15.,  7.,  8.]),
  done: array([0.,  1.,  0.,  0.,  0.,  0.,  1.,  0.,  0.]),
  info: Batch(),
  obs: Batch(
    id: array([ 9., 10., 11., 12., 13., 14., 15.,  7.,  8.]),
  ),
  policy: Batch(),
  rew: array([ 9., 10., 11., 12., 13., 14., 15.,  7.,  8.]),
)
>>> 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()
0.0
>>> # we can get obs_next through __getitem__, even if it doesn't exist
>>> print(buf[:].obs_next.id)
[[ 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.]
 [ 7.  7.  7.  8.]
 [ 7.  7.  8.  9.]]

```

Parameters

- **size** (*int*) – the size of replay buffer.
- **stack_num** (*int*) – the frame-stack sampling argument, should be greater than 1, defaults to 0 (no stacking).
- **ignore_obs_next** (*bool*) – whether to store obs_next, defaults to `False`.
- **sample_avail** (*bool*) – the parameter indicating sampling only available index when using frame-stack sampling method, defaults to `False`. This feature is not supported in Prioritized Replay Buffer currently.

`__getitem__` (*index: Union[slice, int, numpy.integer, numpy.ndarray]*) → `tianshou.data.batch.Batch`
 Return a data batch: `self[index]`. If `stack_num` is set to be `> 0`, return the stacked obs and obs_next with shape `[batch, len, ...]`.

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

add (*obs: Union[dict, tianshou.data.batch.Batch, numpy.ndarray], act: Union[numpy.ndarray, float], rew: Union[int, float], done: bool, obs_next: Optional[Union[dict, tianshou.data.batch.Batch, numpy.ndarray]] = None, info: dict = {}, policy: Optional[Union[dict, tianshou.data.batch.Batch]] = {}, **kwargs*) → `None`
 Add a batch of data into replay buffer.

get (*indice: Union[slice, int, numpy.integer, numpy.ndarray], key: str, stack_num: Optional[int] = None*) → `Union[tianshou.data.batch.Batch, numpy.ndarray]`
 Return the stacked result, e.g. `[s_{t-3}, s_{t-2}, s_{t-1}, s_t]`, where `s` is `self.key`, `t` is `indice`. The `stack_num` (here equals to 4) is given from buffer initialization procedure.

reset () → `None`
 Clear all the data in replay buffer.

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

Returns Sample data and its corresponding index inside the buffer.

update (*buffer: tianshou.data.buffer.ReplayBuffer*) → `None`
 Move the data from the given buffer to self.

`tianshou.data.to_numpy` (*x: Union[torch.Tensor, dict, tianshou.data.batch.Batch, numpy.ndarray]*) → `Union[dict, tianshou.data.batch.Batch, numpy.ndarray]`
 Return an object without `torch.Tensor`.

`tianshou.data.to_torch` (*x: Union[torch.Tensor, dict, tianshou.data.batch.Batch, numpy.ndarray], dtype: Optional[torch.dtype] = None, device: Union[str, int, torch.device] = 'cpu'*) → `Union[dict, tianshou.data.batch.Batch, torch.Tensor]`
 Return an object without `np.ndarray`.

`tianshou.data.to_torch_as` (*x: Union[torch.Tensor, dict, tianshou.data.batch.Batch, numpy.ndarray], y: torch.Tensor*) → `Union[dict, tianshou.data.batch.Batch, torch.Tensor]`
 Return an object without `np.ndarray`. Same as `to_torch(x, dtype=y.dtype, device=y.device)`.

1.6 tianshou.env

class `tianshou.env.BaseVectorEnv` (*env_fns: List[Callable[], gym.core.Env]*)

Bases: `abc.ABC`, `gym.core.Env`

Base class for vectorized environments wrapper. Usage:

```
env_num = 8
envs = VectorEnv([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
```

__len__ () → int

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

abstract close () → None

Close all of the environments.

Environments will automatically `close()` themselves when garbage collected or when the program exits.

abstract render (**kwargs) → None

Render all of the environments.

abstract reset (*id: Optional[Union[int, List[int]]] = None*)

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

abstract seed (*seed: Optional[Union[int, List[int]]] = None*) → 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”.

abstract step (*action: numpy.ndarray*) → Tuple[numpy.ndarray, numpy.ndarray, numpy.ndarray, numpy.ndarray]

Run one timestep of all the environments’ dynamics. 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 (`obs`, `rew`, `done`, `info`).

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)

class `tianshou.env.RayVectorEnv` (*env_fns: List[Callable[], gym.core.Env]*)

Bases: `tianshou.env.vecenv.BaseVectorEnv`

Vectorized environment wrapper based on `ray`. However, according to our test, it is about two times slower than `SubprocVectorEnv`.

See also:

Please refer to `BaseVectorEnv` for more detailed explanation.

close () → List[Any]

Close all of the environments.

Environments will automatically `close()` themselves when garbage collected or when the program exits.

render (***kwargs*) → List[Any]

Render all of the environments.

reset (*id: Optional[Union[int, List[int]]] = None*) → `numpy.ndarray`

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

seed (*seed: Optional[Union[int, List[int]]] = None*) → 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”.

step (*action: numpy.ndarray*) → Tuple[`numpy.ndarray`, `numpy.ndarray`, `numpy.ndarray`, `numpy.ndarray`]

Run one timestep of all the environments’ dynamics. 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 (`obs`, `rew`, `done`, `info`).

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)

class `tianshou.env.SubprocVectorEnv` (*env_fns: List[Callable[], gym.core.Env]*)

Bases: `tianshou.env.vecenv.BaseVectorEnv`

Vectorized environment wrapper based on subprocess.

See also:

Please refer to `BaseVectorEnv` for more detailed explanation.

close () → List[Any]

Close all of the environments.

Environments will automatically close() themselves when garbage collected or when the program exits.

render (**kwargs) → List[Any]

Render all of the environments.

reset (id: Optional[Union[int, List[int]]] = None) → numpy.ndarray

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

seed (seed: Optional[Union[int, List[int]]] = None) → 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".

step (action: numpy.ndarray) → Tuple[numpy.ndarray, numpy.ndarray, numpy.ndarray, numpy.ndarray]

Run one timestep of all the environments' dynamics. 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 (obs, rew, done, info).

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)

class tianshou.env.**VectorEnv** (env_fns: List[Callable[], gym.core.Env])

Bases: tianshou.env.vecenv.BaseVectorEnv

Dummy vectorized environment wrapper, implemented in for-loop.

See also:

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

close () → List[Any]

Close all of the environments.

Environments will automatically close() themselves when garbage collected or when the program exits.

render (**kwargs) → List[Any]

Render all of the environments.

reset (id: Optional[Union[int, List[int]]] = None) → numpy.ndarray

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

seed (*seed*: *Optional[Union[int, List[int]]] = None*) → 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”.

step (*action*: *numpy.ndarray*) → Tuple[*numpy.ndarray*, *numpy.ndarray*, *numpy.ndarray*, *numpy.ndarray*]

Run one timestep of all the environments’ dynamics. 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 (obs, rew, done, info).

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)

1.7 tianshou.policy

```
class tianshou.policy.A2CPolicy (actor: torch.nn.modules.module.Module,
                                critic: torch.nn.modules.module.Module, optim: torch.optim.optimizer.Optimizer,
                                dist_fn: torch.distributions.distribution.Distribution = <class torch.distributions.categorical.Categorical>,
                                discount_factor: float = 0.99, vf_coef: float = 0.5, ent_coef: float = 0.01,
                                max_grad_norm: Optional[float] = None, gae_lambda: float = 0.95,
                                reward_normalization: bool = False, **kwargs)
```

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** (*torch.distributions.Distribution*) – for computing the action, defaults to `torch.distributions.Categorical`.
- **discount_factor** (*float*) – in [0, 1], defaults to 0.99.
- **vf_coef** (*float*) – weight for value loss, defaults to 0.5.
- **ent_coef** (*float*) – weight for entropy loss, defaults to 0.01.

- `max_grad_norm` (*float*) – clipping gradients in back propagation, defaults to `None`.
- `gae_lambda` (*float*) – in `[0, 1]`, param for Generalized Advantage Estimation, defaults to `0.95`.

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*) → `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*) → `Dict[str, List[float]]`
 Update policy with a given batch of data.

Returns A dict which includes loss and its corresponding label.

process_fn (*batch*: `tianshou.data.batch.Batch`, *buffer*: `tianshou.data.buffer.ReplayBuffer`, *indice*: `numpy.ndarray`) → `tianshou.data.batch.Batch`
 Compute the discounted returns for each frame:

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

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

class `tianshou.policy.BasePolicy` (***kwargs*)
 Bases: `abc.ABC`, `torch.nn.modules.module.Module`

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

A policy class typically has four 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.

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'))
```

static compute_episodic_return (*batch*: `tianshou.data.batch.Batch`, *v_s*: `Optional[Union[numpy.ndarray, torch.Tensor]] = None`, *gamma*: `float = 0.99`, *gae_lambda*: `float = 0.95`) → `tianshou.data.batch.Batch`

Compute returns over given full-length episodes, including the implementation of Generalized Advantage Estimator (arXiv:1506.02438).

Parameters

- **batch** (`Batch`) – a data batch which contains several full-episode data chronologically.
- **v_s** (`numpy.ndarray`) – the value function of all next states $V(s')$.
- **gamma** (`float`) – the discount factor, should be in $[0, 1]$, defaults to 0.99.
- **gae_lambda** (`float`) – the parameter for Generalized Advantage Estimation, should be in $[0, 1]$, defaults to 0.95.

Returns a `Batch`. The result will be stored in `batch.returns`.

static compute_nstep_return (*batch*: `tianshou.data.batch.Batch`, *buffer*: `tianshou.data.buffer.ReplayBuffer`, *indice*: `numpy.ndarray`, *target_q_fn*: `Callable[[tianshou.data.buffer.ReplayBuffer, numpy.ndarray], torch.Tensor]`, *gamma*: `float = 0.99`, *n_step*: `int = 1`, *rew_norm*: `bool = False`) → `numpy.ndarray`

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 γ 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`) – a data buffer which contains several full-episode data chronologically.
- **indice** (`numpy.ndarray`) – sampled timestep.
- **target_q_fn** (`function`) – a function receives $t + n - 1$ step's data and compute target Q value.
- **gamma** (`float`) – the discount factor, should be in $[0, 1]$, defaults to 0.99.
- **n_step** (`int`) – the number of estimation step, should be an int greater than 0, defaults to 1.
- **rew_norm** (`bool`) – normalize the reward to Normal(0, 1), defaults to `False`.

Returns a Batch. The result will be stored in batch.returns as a torch.Tensor with shape (bsz,).

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

After version $\geq 0.2.3$, the keyword “policy” is reserved and the corresponding data will be stored into the replay buffer in numpy. For instance,

```
# some code
return Batch(..., policy=Batch(log_prob=dist.log_prob(act)))
# and in the sampled data batch, you can directly call
# batch.policy.log_prob to get your data, although it is stored in
# np.ndarray.
```

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

Update policy with a given batch of data.

Returns A dict which includes loss and its corresponding label.

```
process_fn (batch: tianshou.data.batch.Batch, buffer: tianshou.data.buffer.ReplayBuffer, indice: numpy.ndarray) → tianshou.data.batch.Batch
```

Pre-process the data from the provided replay buffer. Check out *Policy* for more information.

```
class tianshou.policy.DDPGPolicy (actor: torch.nn.modules.module.Module, actor_optim: torch.optim.optimizer.Optimizer, critic: torch.nn.modules.module.Module, critic_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>, action_range: Optional[Tuple[float, float]] = None, reward_normalization: bool = False, ignore_done: bool = False, estimation_step: int = 1, **kwargs)
```

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, defaults to 0.005.
- **gamma** (*float*) – discount factor, in [0, 1], defaults to 0.99.
- **exploration_noise** (*BaseNoise*) – the exploration noise, add to the action, defaults to `GaussianNoise(sigma=0.1)`.
- **action_range** (*(float, float)*) – the action range (minimum, maximum).
- **reward_normalization** (*bool*) – normalize the reward to `Normal(0, 1)`, defaults to `False`.
- **ignore_done** (*bool*) – ignore the done flag while training the policy, defaults to `False`.
- **estimation_step** (*int*) – greater than 1, the number of steps to look ahead.

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, model: str = 'actor', input: str = 'obs', exploring: bool = True, **kwargs*) → `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*) → `Dict[str, float]`
 Update policy with a given batch of data.

Returns A dict which includes loss and its corresponding label.

process_fn (*batch: tianshou.data.batch.Batch, buffer: tianshou.data.buffer.ReplayBuffer, indice: numpy.ndarray*) → `tianshou.data.batch.Batch`
 Pre-process the data from the provided replay buffer. Check out `Policy` for more information.

set_exp_noise (*noise: Optional[tianshou.exploration.random.BaseNoise]*) → `None`
 Set the exploration noise.

sync_weight () → `None`
 Soft-update the weight for the target network.

train (*mode=True*) → `torch.nn.modules.module.Module`
 Set the module in training mode, except for the target network.

class `tianshou.policy.DQNPolicy` (*model: torch.nn.modules.module.Module, optim: torch.optim.optimizer.Optimizer, discount_factor: float = 0.99, estimation_step: int = 1, target_update_freq: Optional[int] = 0, **kwargs*)

Bases: `tianshou.policy.base.BasePolicy`

Implementation of Deep Q Network. arXiv:1312.5602 Implementation of Double Q-Learning. arXiv:1509.06461

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*) – greater than 1, the number of steps to look ahead.
- **target_update_freq** (*int*) – the target network update frequency (0 if you do not use the target network).

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, model: str = 'model', input: str = 'obs', eps: Optional[float] = None, **kwargs*) → *tianshou.data.batch.Batch*
 Compute action over the given batch data.

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*) → *Dict[str, float]*
 Update policy with a given batch of data.

Returns A dict which includes loss and its corresponding label.

process_fn (*batch: tianshou.data.batch.Batch, buffer: tianshou.data.buffer.ReplayBuffer, indice: numpy.ndarray*) → *tianshou.data.batch.Batch*
 Compute the n-step return for Q-learning targets:

$$G_t = \sum_{i=t}^{t+n-1} \gamma^{i-t} (1 - d_i) r_i + \gamma^n (1 - d_{t+n}) \max_a Q_{old}(s_{t+n}, \arg \max_a (Q_{new}(s_{t+n}, a)))$$

, where γ is the discount factor, $\gamma \in [0, 1]$, d_t is the done flag of step t . If there is no target network, the Q_{old} is equal to Q_{new} .

set_eps (*eps: float*) → *None*
 Set the eps for epsilon-greedy exploration.

sync_weight () → *None*
 Synchronize the weight for the target network.

train (*mode=True*) → *torch.nn.modules.module.Module*
 Set the module in training mode, except for the target network.

class *tianshou.policy.ImitationPolicy* (*model: torch.nn.modules.module.Module, optim: torch.optim.optimizer.Optimizer, mode: str = 'continuous', **kwargs*)

Bases: *tianshou.policy.base.BasePolicy*

Implementation of vanilla imitation learning (for continuous action space).

Parameters

- **model** (*torch.nn.Module*) – a model following the rules in *BasePolicy*. (s -> a)
- **optim** (*torch.optim.Optimizer*) – for optimizing the model.
- **mode** (*str*) – indicate the imitation type (“continuous” or “discrete” action space), defaults to “continuous”.

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*) → *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=...)
```

After version >= 0.2.3, the keyword “policy” is reserved and the corresponding data will be stored into the replay buffer in *numpy*. For instance,

```
# some code
return Batch(..., policy=Batch(log_prob=dist.log_prob(act)))
# and in the sampled data batch, you can directly call
# batch.policy.log_prob to get your data, although it is stored in
# np.ndarray.
```

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

Update policy with a given batch of data.

Returns A dict which includes loss and its corresponding label.

```
class tianshou.policy.PGPolicy(model: torch.nn.modules.module.Module, optim: torch.optim.optimizer.Optimizer, dist_fn: torch.distributions.distribution.Distribution = <class 'torch.distributions.categorical.Categorical'>, discount_factor: float = 0.99, reward_normalization: bool = False, **kwargs)
```

Bases: *tianshou.policy.base.BasePolicy*

Implementation of Vanilla Policy Gradient.

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** (*torch.distributions.Distribution*) – for computing the action.

- **discount_factor** (*float*) – in $[0, 1]$.

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*) → *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*) → *Dict[str, List[float]]*
 Update policy with a given batch of data.

Returns A dict which includes loss and its corresponding label.

process_fn (*batch: tianshou.data.batch.Batch, buffer: tianshou.data.buffer.ReplayBuffer, indice: numpy.ndarray*) → *tianshou.data.batch.Batch*
 Compute the discounted returns for each frame:

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

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

```
class tianshou.policy.PPOPolicy (actor: torch.nn.modules.module.Module,
                                critic: torch.nn.modules.module.Module, op-
                                tim: torch.optim.optimizer.Optimizer, dist_fn:
                                torch.distributions.distribution.Distribution, discount_factor:
                                float = 0.99, max_grad_norm: Optional[float] = None,
                                eps_clip: float = 0.2, vf_coef: float = 0.5, ent_coef: float
                                = 0.01, action_range: Optional[Tuple[float, float]] = None,
                                gae_lambda: float = 0.95, dual_clip: Optional[float] = None,
                                value_clip: bool = True, reward_normalization: bool = True,
                                **kwargs)
```

Bases: *tianshou.policy.modelfree.pg.PGPPolicy*

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** (*torch.distributions.Distribution*) – for computing the action.

- **discount_factor** (*float*) – in $[0, 1]$, defaults to 0.99.
- **max_grad_norm** (*float*) – clipping gradients in back propagation, defaults to `None`.
- **eps_clip** (*float*) – ϵ in L_{CLIP} in the original paper, defaults to 0.2.
- **vf_coef** (*float*) – weight for value loss, defaults to 0.5.
- **ent_coef** (*float*) – weight for entropy loss, defaults to 0.01.
- **action_range** (*float, float*) – the action range (minimum, maximum).
- **gae_lambda** (*float*) – in $[0, 1]$, param for Generalized Advantage Estimation, defaults to 0.95.
- **dual_clip** (*float*) – a parameter c mentioned in arXiv:1912.09729 Equ. 5, where $c > 1$ is a constant indicating the lower bound, defaults 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, defaults to `True`.
- **reward_normalization** (*bool*) – normalize the returns to $\text{Normal}(0, 1)$, defaults to `True`.

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*) → `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*) → `Dict[str, List[float]]`
 Update policy with a given batch of data.

Returns A dict which includes loss and its corresponding label.

process_fn (*batch: tianshou.data.batch.Batch, buffer: tianshou.data.buffer.ReplayBuffer, indice: numpy.ndarray*) → `tianshou.data.batch.Batch`
 Compute the discounted returns for each frame:

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

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

```
class tianshou.policy.SACPolicy(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, alpha: Tuple[float, torch.Tensor, torch.optim.optimizer.Optimizer] = 0.2, action_range: Optional[Tuple[float, float]] = None, reward_normalization: bool = False, ignore_done: bool = False, estimation_step: int = 1, exploration_noise: Optional[tianshou.exploration.random.BaseNoise] = None, **kwargs)
```

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, defaults to 0.005.
- **gamma** (*float*) – discount factor, in [0, 1], defaults to 0.99.
- **exploration_noise** (*BaseNoise*) – the noise intensity, add to the action, defaults to 0.1.
- **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.
- **action_range** (*(float, float)*) – the action range (minimum, maximum).
- **reward_normalization** (*bool*) – normalize the reward to Normal(0, 1), defaults to False.
- **ignore_done** (*bool*) – ignore the done flag while training the policy, defaults to False.
- **exploration_noise** – add a noise to action for exploration. This is useful when solving hard-exploration problem.

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', exploring: bool = True, **kwargs) → 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*) → Dict[str, float]

Update policy with a given batch of data.

Returns A dict which includes loss and its corresponding label.

sync_weight () → None

Soft-update the weight for the target network.

train (*mode*=*True*) → torch.nn.modules.module.Module

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

```
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, action_range: Optional[Tuple[float, float]] = None, reward_normalization: bool = False, ignore_done: bool = False, estimation_step: int = 1, **kwargs)
```

Bases: `tianshou.policy.modelfree.ddpg.DDPGPoly`

Implementation of Twin Delayed Deep Deterministic Policy Gradient, 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, defaults to 0.005.
- **gamma** (*float*) – discount factor, in [0, 1], defaults to 0.99.
- **exploration_noise** (*float*) – the exploration noise, add to the action, defaults 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.
- **action_range** (*(float, float)*) – the action range (minimum, maximum).
- **reward_normalization** (*bool*) – normalize the reward to Normal(0, 1), defaults to False.
- **ignore_done** (*bool*) – ignore the done flag while training the policy, defaults to False.

See also:

Please refer to *BasePolicy* for more detailed explanation.

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

Update policy with a given batch of data.

Returns A dict which includes loss and its corresponding label.

sync_weight () → None

Soft-update the weight for the target network.

train (*mode=True*) → torch.nn.modules.module.Module

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

1.8 tianshou.trainer

`tianshou.trainer.gather_info` (*start_time: float, train_c: tianshou.data.collector.Collector, test_c: tianshou.data.collector.Collector, best_reward: 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 frames in the training collector;
- `train_time/model` the time for training models;
- `train_speed` the speed of training (frames 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 (frames per second);
- `best_reward` the best reward over the test results;
- `duration` the total elapsed time.

```
tianshou.trainer.offpolicy_trainer (policy: tianshou.policy.base.BasePolicy, train_collector:
tianshou.data.collector.Collector, test_collector:
tianshou.data.collector.Collector, max_epoch:
int, step_per_epoch: int, collect_per_step: int,
episode_per_test: Union[int, List[int]], batch_size:
int, update_per_step: int = 1, train_fn: Op-
tional[Callable[[int], None]] = None, test_fn: Op-
tional[Callable[[int], None]] = None, stop_fn: Op-
tional[Callable[[float], bool]] = None, save_fn:
Optional[Callable[[tianshou.policy.base.BasePolicy],
None]] = None, log_fn: Op-
tional[Callable[[dict], None]] = None, writer: Op-
tional[torch.utils.tensorboard.writer.SummaryWriter]
= None, log_interval: int = 1, verbose: bool = True,
**kwargs) → Dict[str, Union[float, str]]
```

A wrapper for off-policy trainer procedure.

Parameters

- **policy** – an instance of the *BasePolicy* class.
- **train_collector** (*Collector*) – the collector used for training.
- **test_collector** (*Collector*) – the collector used for testing.
- **max_epoch** (*int*) – the maximum of epochs for training. The training process might be finished before reaching the `max_epoch`.
- **step_per_epoch** (*int*) – the number of step for updating policy network in one epoch.
- **collect_per_step** (*int*) – the number of frames the collector would collect before the network update. In other words, collect some frames and do some policy network update.
- **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*) – the number of times the policy network would be updated after frames be collected. In other words, collect some frames and do some policy network update.
- **train_fn** (*function*) – a function receives the current number of epoch index and performs some operations at the beginning of training in this epoch.
- **test_fn** (*function*) – a function receives the current number of epoch index and performs some operations at the beginning of testing in this epoch.
- **save_fn** (*function*) – a function for saving policy when the undiscounted average mean reward in evaluation phase gets better.
- **stop_fn** (*function*) – a function receives the average undiscounted returns of the testing result, return a boolean which indicates whether reaching the goal.
- **log_fn** (*function*) – a function receives env info for logging.
- **writer** (*torch.utils.tensorboard.SummaryWriter*) – a TensorBoard SummaryWriter.
- **log_interval** (*int*) – the log interval of the writer.
- **verbose** (*bool*) – whether to print the information.

Returns See `gather_info()`.

```

tianshou.trainer.onpolicy_trainer (policy: tianshou.policy.base.BasePolicy, train_collector:
    tianshou.data.collector.Collector, test_collector:
    tianshou.data.collector.Collector, max_epoch: int,
    step_per_epoch: int, collect_per_step: int, re-
    peat_per_collect: int, episode_per_test: Union[int,
    List[int]], batch_size: int, train_fn: Op-
    tional[Callable[[int], None]] = None, test_fn: Op-
    tional[Callable[[int], None]] = None, stop_fn: Op-
    tional[Callable[[float], bool]] = None, save_fn: Op-
    tional[Callable[[tianshou.policy.base.BasePolicy],
    None]] = None, log_fn: Op-
    tional[Callable[[dict], None]] = None, writer: Op-
    tional[torch.utils.tensorboard.writer.SummaryWriter] =
    None, log_interval: int = 1, verbose: bool = True,
    **kwargs) → Dict[str, Union[float, str]]

```

A wrapper for on-policy trainer procedure.

Parameters

- **policy** – an instance of the *BasePolicy* class.
- **train_collector** (*Collector*) – the collector used for training.
- **test_collector** (*Collector*) – the collector used for testing.
- **max_epoch** (*int*) – the maximum of epochs for training. The training process might be finished before reaching the `max_epoch`.
- **step_per_epoch** (*int*) – the number of step for updating policy network in one epoch.
- **collect_per_step** (*int*) – the number of frames the collector would collect before the network update. In other words, collect some frames and do one policy network update.
- **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 or list of ints*) – 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.
- **train_fn** (*function*) – a function receives the current number of epoch index and performs some operations at the beginning of training in this epoch.
- **test_fn** (*function*) – a function receives the current number of epoch index and performs some operations at the beginning of testing in this epoch.
- **save_fn** (*function*) – a function for saving policy when the undiscounted average mean reward in evaluation phase gets better.
- **stop_fn** (*function*) – a function receives the average undiscounted returns of the testing result, return a boolean which indicates whether reaching the goal.
- **log_fn** (*function*) – a function receives env info for logging.
- **writer** (*torch.utils.tensorboard.SummaryWriter*) – a TensorBoard SummaryWriter.
- **log_interval** (*int*) – the log interval of the writer.
- **verbose** (*bool*) – whether to print the information.

Returns See `gather_info()`.

`tianshou.trainer.test_episode` (*policy: tianshou.policy.base.BasePolicy, collector: tianshou.data.collector.Collector, test_fn: Callable[[int], None], epoch: int, n_episode: Union[int, List[int]]*) → Dict[str, float]

A simple wrapper of testing policy in collector.

1.9 tianshou.exploration

class `tianshou.exploration.BaseNoise` (***kwargs*)
 Bases: `abc.ABC`, `object`

The action noise base class.

abstract `__call__` (***kwargs*) → `numpy.ndarray`
 Generate new noise.

reset (***kwargs*) → `None`
 Reset to the initial state.

class `tianshou.exploration.GaussianNoise` (*mu: float = 0.0, sigma: float = 1.0*)
 Bases: `tianshou.exploration.random.BaseNoise`

Class for vanilla gaussian process, used for exploration in DDPG by default.

`__call__` (*size: tuple*) → `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.

`__call__` (*size: tuple, mu: Optional[float] = None*) → `numpy.ndarray`
 Generate new noise. Return a `numpy.ndarray` which size is equal to `size`.

reset () → `None`
 Reset to the initial state.

1.10 tianshou.utils

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 (*x: Union[float, 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.net.common.Net (*layer_num, state_shape, action_shape=0, device='cpu', softmax=False, concat=False*)

Bases: torch.nn.modules.module.Module

Simple MLP backbone. For advanced usage (how to customize the network), please refer to [Build the Network](#).

Parameters **concat** – 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.

forward (*s, state=None, info={}*)

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.

class tianshou.utils.net.common.Recurrent (*layer_num, state_shape, action_shape, device='cpu'*)

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 (*s, state=None, info={}*)

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.

class `tianshou.utils.net.discrete.Actor` (*preprocess_net, action_shape*)

Bases: `torch.nn.modules.module.Module`

For advanced usage (how to customize the network), please refer to [Build the Network](#).

forward (*s, state=None, info={}*)

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.

class `tianshou.utils.net.discrete.Critic` (*preprocess_net*)

Bases: `torch.nn.modules.module.Module`

For advanced usage (how to customize the network), please refer to [Build the Network](#).

forward (*s, **kwargs*)

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.

class `tianshou.utils.net.discrete.DQN` (*h, w, action_shape, device='cpu'*)

Bases: `torch.nn.modules.module.Module`

For advanced usage (how to customize the network), please refer to [Build the Network](#).

forward (*x, state=None, info={}*)

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.

class `tianshou.utils.net.continuous.Actor` (*preprocess_net, action_shape, max_action, device='cpu'*)

Bases: `torch.nn.modules.module.Module`

For advanced usage (how to customize the network), please refer to [Build the Network](#).

forward (*s, state=None, info={}*)

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.

```
class tianshou.utils.net.continuous.ActorProb (preprocess_net,          action_shape,
                                             max_action,      device='cpu',    un-
                                             bounded=False)
```

Bases: `torch.nn.modules.module.Module`

For advanced usage (how to customize the network), please refer to [Build the Network](#).

forward (*s*, *state=None*, ***kwargs*)

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.

```
class tianshou.utils.net.continuous.Critic (preprocess_net, device='cpu')
```

Bases: `torch.nn.modules.module.Module`

For advanced usage (how to customize the network), please refer to [Build the Network](#).

forward (*s*, *a=None*, ***kwargs*)

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.

```
class tianshou.utils.net.continuous.RecurrentActorProb (layer_num, state_shape, ac-
                                                         tion_shape, max_action, de-
                                                         vice='cpu')
```

Bases: `torch.nn.modules.module.Module`

For advanced usage (how to customize the network), please refer to [Build the Network](#).

forward (*s*, ***kwargs*)

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.

```
class tianshou.utils.net.continuous.RecurrentCritic (layer_num, state_shape, ac-
                                                         tion_shape=0, device='cpu')
```

Bases: `torch.nn.modules.module.Module`

For advanced usage (how to customize the network), please refer to *Build the Network*.

forward (*s*, *a=None*)

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.

1.11 Contributing to Tianshou

1.11.1 Install Develop Version

To install Tianshou in an “editable” mode, run

```
pip3 install -e ".[dev]"
```

in the main directory. This installation is removable by

```
python3 setup.py develop --uninstall
```

1.11.2 PEP8 Code Style Check

We follow PEP8 python code style. To check, in the main directory, run:

```
flake8 . --count --show-source --statistics
```

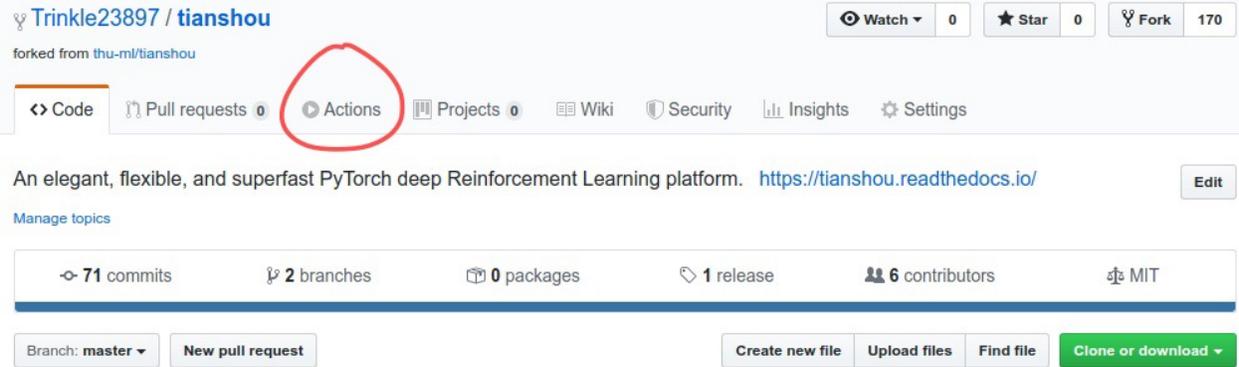
1.11.3 Test Locally

This command will run automatic tests in the main directory

```
pytest test --cov tianshou -s --durations 0 -v
```

1.11.4 Test by GitHub Actions

1. Click the `Actions` button in your own repo:



Trinkle23897 / tianshou
forked from thu-ml/tianshou

<> Code Pull requests 0 Actions Projects 0 Wiki Security Insights Settings

An elegant, flexible, and superfast PyTorch deep Reinforcement Learning platform. <https://tianshou.readthedocs.io/> Edit

Manage topics

71 commits 2 branches 0 packages 1 release 6 contributors MIT

Branch: master New pull request Create new file Upload files Find file Clone or download

2. Click the green button:



Workflows aren't being run on this forked repository

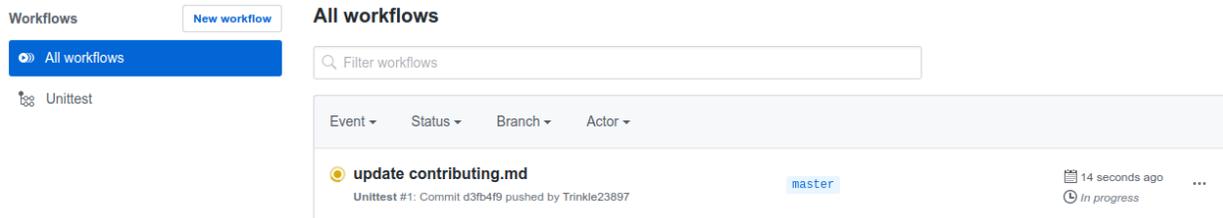
Because this repository contained workflow files when it was forked, we have disabled them from running on this fork. Make sure you understand the configured workflows and their expected usage before enabling Actions on this repository.



[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:



Workflows New workflow

All workflows

Filter workflows

Event	Status	Branch	Actor
update contributing.md	In progress	master	Trinkle23897

Unittest

Unittest #1: Commit d3fb4f9 pushed by Trinkle23897 14 seconds ago

1.11.5 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 webpages, run

```
make html
```

under the `docs/` directory. The generated webpages are in `docs/_build` and can be viewed with browsers.

1.11.6 Chinese Documentation

Chinese documentation is in <https://tianshou.readthedocs.io/zh/latest/>

1.12 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](#))
- Minghao Zhang ([Mehooz](#))
- Alexis Duburcq ([duburcqa](#))

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.data`, 22
`tianshou.env`, 32
`tianshou.exploration`, 49
`tianshou.policy`, 35
`tianshou.trainer`, 46
`tianshou.utils`, 50
`tianshou.utils.net.common`, 50
`tianshou.utils.net.continuous`, 51
`tianshou.utils.net.discrete`, 51

Symbols

`__call__()` (*tianshou.exploration.BaseNoise method*), 49
`__call__()` (*tianshou.exploration.GaussianNoise method*), 49
`__call__()` (*tianshou.exploration.OUNoise method*), 49
`__getitem__()` (*tianshou.data.Batch method*), 25
`__getitem__()` (*tianshou.data.PrioritizedReplayBuffer method*), 29
`__getitem__()` (*tianshou.data.ReplayBuffer method*), 31
`__len__()` (*tianshou.data.Batch method*), 25
`__len__()` (*tianshou.data.ReplayBuffer method*), 31
`__len__()` (*tianshou.env.BaseVectorEnv method*), 32

A

A2CPolicy (*class in tianshou.policy*), 35
Actor (*class in tianshou.utils.net.continuous*), 51
Actor (*class in tianshou.utils.net.discrete*), 51
ActorProb (*class in tianshou.utils.net.continuous*), 52
add() (*tianshou.data.PrioritizedReplayBuffer method*), 29
add() (*tianshou.data.ReplayBuffer method*), 31
add() (*tianshou.utils.MovAvg method*), 50
append() (*tianshou.data.Batch method*), 25

B

BaseNoise (*class in tianshou.exploration*), 49
BasePolicy (*class in tianshou.policy*), 36
BaseVectorEnv (*class in tianshou.env*), 32
Batch (*class in tianshou.data*), 22

C

cat() (*tianshou.data.Batch static method*), 25
cat_() (*tianshou.data.Batch method*), 25
close() (*tianshou.data.Collector method*), 27
close() (*tianshou.env.BaseVectorEnv method*), 32
close() (*tianshou.env.RayVectorEnv method*), 33
close() (*tianshou.env.SubprocVectorEnv method*), 33
close() (*tianshou.env.VectorEnv method*), 34

collect() (*tianshou.data.Collector method*), 27
Collector (*class in tianshou.data*), 26
compute_episodic_return() (*tianshou.policy.BasePolicy static method*), 37
compute_nstep_return() (*tianshou.policy.BasePolicy static method*), 37
Critic (*class in tianshou.utils.net.continuous*), 52
Critic (*class in tianshou.utils.net.discrete*), 51

D

DDPGPolicy (*class in tianshou.policy*), 38
DQN (*class in tianshou.utils.net.discrete*), 51
DQNPolicy (*class in tianshou.policy*), 39

E

empty() (*tianshou.data.Batch static method*), 25
empty_() (*tianshou.data.Batch method*), 25

F

forward() (*tianshou.policy.A2CPolicy method*), 36
forward() (*tianshou.policy.BasePolicy method*), 38
forward() (*tianshou.policy.DDPGPolicy method*), 39
forward() (*tianshou.policy.DQNPolicy method*), 40
forward() (*tianshou.policy.ImitationPolicy method*), 41
forward() (*tianshou.policy.PGPolicy method*), 42
forward() (*tianshou.policy.PPOPolicy method*), 43
forward() (*tianshou.policy.SACPolicy method*), 44
forward() (*tianshou.utils.net.common.Net method*), 50
forward() (*tianshou.utils.net.common.Recurrent method*), 50
forward() (*tianshou.utils.net.continuous.Actor method*), 51
forward() (*tianshou.utils.net.continuous.ActorProb method*), 52
forward() (*tianshou.utils.net.continuous.Critic method*), 52
forward() (*tianshou.utils.net.continuous.RecurrentActorProb method*), 52
forward() (*tianshou.utils.net.continuous.RecurrentCritic method*), 53

- forward() (*tianshou.utils.net.discrete.Actor method*), 51
- forward() (*tianshou.utils.net.discrete.Critic method*), 51
- forward() (*tianshou.utils.net.discrete.DQN method*), 51
- ## G
- gather_info() (*in module tianshou.trainer*), 46
- GaussianNoise (*class in tianshou.exploration*), 49
- get() (*tianshou.data.Batch method*), 25
- get() (*tianshou.data.ReplayBuffer method*), 31
- get() (*tianshou.utils.MovAvg method*), 50
- get_env_num() (*tianshou.data.Collector method*), 28
- ## I
- ImitationPolicy (*class in tianshou.policy*), 40
- items() (*tianshou.data.Batch method*), 25
- ## K
- keys() (*tianshou.data.Batch method*), 25
- ## L
- learn() (*tianshou.policy.A2CPolicy method*), 36
- learn() (*tianshou.policy.BasePolicy method*), 38
- learn() (*tianshou.policy.DDPGPolicy method*), 39
- learn() (*tianshou.policy.DQNPolicy method*), 40
- learn() (*tianshou.policy.ImitationPolicy method*), 41
- learn() (*tianshou.policy.PGPolicy method*), 42
- learn() (*tianshou.policy.PPOPolicy method*), 43
- learn() (*tianshou.policy.SACPolicy method*), 45
- learn() (*tianshou.policy.TD3Policy method*), 46
- ListReplayBuffer (*class in tianshou.data*), 28
- ## M
- mean() (*tianshou.utils.MovAvg method*), 50
- module
- tianshou.data, 22
 - tianshou.env, 32
 - tianshou.exploration, 49
 - tianshou.policy, 35
 - tianshou.trainer, 46
 - tianshou.utils, 50
 - tianshou.utils.net.common, 50
 - tianshou.utils.net.continuous, 51
 - tianshou.utils.net.discrete, 51
- MovAvg (*class in tianshou.utils*), 50
- ## N
- Net (*class in tianshou.utils.net.common*), 50
- ## O
- offpolicy_trainer() (*in module tianshou.trainer*), 46
- onpolicy_trainer() (*in module tianshou.trainer*), 48
- OUNoise (*class in tianshou.exploration*), 49
- ## P
- PGPolicy (*class in tianshou.policy*), 41
- PPOPolicy (*class in tianshou.policy*), 42
- PrioritizedReplayBuffer (*class in tianshou.data*), 28
- process_fn() (*tianshou.policy.A2CPolicy method*), 36
- process_fn() (*tianshou.policy.BasePolicy method*), 38
- process_fn() (*tianshou.policy.DDPGPolicy method*), 39
- process_fn() (*tianshou.policy.DQNPolicy method*), 40
- process_fn() (*tianshou.policy.PGPolicy method*), 42
- process_fn() (*tianshou.policy.PPOPolicy method*), 43
- ## R
- RayVectorEnv (*class in tianshou.env*), 33
- Recurrent (*class in tianshou.utils.net.common*), 50
- RecurrentActorProb (*class in tianshou.utils.net.continuous*), 52
- RecurrentCritic (*class in tianshou.utils.net.continuous*), 52
- render() (*tianshou.data.Collector method*), 28
- render() (*tianshou.env.BaseVectorEnv method*), 32
- render() (*tianshou.env.RayVectorEnv method*), 33
- render() (*tianshou.env.SubprocVectorEnv method*), 34
- render() (*tianshou.env.VectorEnv method*), 34
- replace() (*tianshou.data.PrioritizedReplayBuffer property*), 29
- ReplayBuffer (*class in tianshou.data*), 29
- reset() (*tianshou.data.Collector method*), 28
- reset() (*tianshou.data.ListReplayBuffer method*), 28
- reset() (*tianshou.data.PrioritizedReplayBuffer method*), 29
- reset() (*tianshou.data.ReplayBuffer method*), 31
- reset() (*tianshou.env.BaseVectorEnv method*), 32
- reset() (*tianshou.env.RayVectorEnv method*), 33
- reset() (*tianshou.env.SubprocVectorEnv method*), 34
- reset() (*tianshou.env.VectorEnv method*), 34
- reset() (*tianshou.exploration.BaseNoise method*), 49
- reset() (*tianshou.exploration.OUNoise method*), 49
- reset_buffer() (*tianshou.data.Collector method*), 28
- reset_env() (*tianshou.data.Collector method*), 28
- ## S
- SACPolicy (*class in tianshou.policy*), 43
- sample() (*tianshou.data.Collector method*), 28

`sample()` (*tianshou.data.ListReplayBuffer method*), 28
`sample()` (*tianshou.data.PrioritizedReplayBuffer method*), 29
`sample()` (*tianshou.data.ReplayBuffer method*), 31
`seed()` (*tianshou.data.Collector method*), 28
`seed()` (*tianshou.env.BaseVectorEnv method*), 32
`seed()` (*tianshou.env.RayVectorEnv method*), 33
`seed()` (*tianshou.env.SubprocVectorEnv method*), 34
`seed()` (*tianshou.env.VectorEnv method*), 34
`set_eps()` (*tianshou.policy.DQNPolicy method*), 40
`set_exp_noise()` (*tianshou.policy.DDPGPolicy method*), 39
`shape()` (*tianshou.data.Batch property*), 25
`split()` (*tianshou.data.Batch method*), 25
`stack()` (*tianshou.data.Batch static method*), 25
`stack_()` (*tianshou.data.Batch method*), 25
`std()` (*tianshou.utils.MovAvg method*), 50
`step()` (*tianshou.env.BaseVectorEnv method*), 32
`step()` (*tianshou.env.RayVectorEnv method*), 33
`step()` (*tianshou.env.SubprocVectorEnv method*), 34
`step()` (*tianshou.env.VectorEnv method*), 35
`SubprocVectorEnv` (*class in tianshou.env*), 33
`sync_weight()` (*tianshou.policy.DDPGPolicy method*), 39
`sync_weight()` (*tianshou.policy.DQNPolicy method*), 40
`sync_weight()` (*tianshou.policy.SACPolicy method*), 45
`sync_weight()` (*tianshou.policy.TD3Policy method*), 46

T

`TD3Policy` (*class in tianshou.policy*), 45
`test_episode()` (*in module tianshou.trainer*), 49
`tianshou.data`
 module, 22
`tianshou.env`
 module, 32
`tianshou.exploration`
 module, 49
`tianshou.policy`
 module, 35
`tianshou.trainer`
 module, 46
`tianshou.utils`
 module, 50
`tianshou.utils.net.common`
 module, 50
`tianshou.utils.net.continuous`
 module, 51
`tianshou.utils.net.discrete`
 module, 51
`to_numpy()` (*in module tianshou.data*), 31
`to_numpy()` (*tianshou.data.Batch method*), 26

`to_torch()` (*in module tianshou.data*), 31
`to_torch()` (*tianshou.data.Batch method*), 26
`to_torch_as()` (*in module tianshou.data*), 31
`train()` (*tianshou.policy.DDPGPolicy method*), 39
`train()` (*tianshou.policy.DQNPolicy method*), 40
`train()` (*tianshou.policy.SACPolicy method*), 45
`train()` (*tianshou.policy.TD3Policy method*), 46

U

`update()` (*tianshou.data.ReplayBuffer method*), 31
`update_weight()` (*tianshou.data.PrioritizedReplayBuffer method*), 29

V

`values()` (*tianshou.data.Batch method*), 26
`VectorEnv` (*class in tianshou.env*), 34