

---

# **Stable Baselines3 Documentation**

***Release 1.6.2***

**Stable Baselines3 Contributors**

**Mar 12, 2023**



# USER GUIDE

<b>1</b>	<b>Installation</b>	<b>3</b>
1.1	Prerequisites . . . . .	3
1.2	Bleeding-edge version . . . . .	3
1.3	Development version . . . . .	3
<b>2</b>	<b>RL Algorithms</b>	<b>5</b>
<b>3</b>	<b>Examples</b>	<b>7</b>
3.1	TQC . . . . .	7
3.2	QR-DQN . . . . .	7
3.3	MaskablePPO . . . . .	7
3.4	TRPO . . . . .	8
3.5	ARS . . . . .	8
3.6	RecurrentPPO . . . . .	8
<b>4</b>	<b>ARS</b>	<b>9</b>
4.1	Notes . . . . .	9
4.2	Can I use? . . . . .	9
4.3	Example . . . . .	10
4.4	Results . . . . .	10
4.5	Parameters . . . . .	11
4.6	ARS Policies . . . . .	16
<b>5</b>	<b>Maskable PPO</b>	<b>17</b>
5.1	Notes . . . . .	17
5.2	Can I use? . . . . .	17
5.3	Example . . . . .	18
5.4	Results . . . . .	19
5.5	Parameters . . . . .	24
5.6	MaskablePPO Policies . . . . .	28
<b>6</b>	<b>Recurrent PPO</b>	<b>33</b>
6.1	Notes . . . . .	33
6.2	Can I use? . . . . .	33
6.3	Example . . . . .	34
6.4	Results . . . . .	34
6.5	Parameters . . . . .	35
6.6	RecurrentPPO Policies . . . . .	40
<b>7</b>	<b>QR-DQN</b>	<b>47</b>
7.1	Notes . . . . .	47

7.2	Can I use? . . . . .	47
7.3	Example . . . . .	48
7.4	Results . . . . .	48
7.5	Parameters . . . . .	49
7.6	QR-DQN Policies . . . . .	54
<b>8</b>	<b>TQC</b>	<b>57</b>
8.1	Notes . . . . .	57
8.2	Can I use? . . . . .	57
8.3	Example . . . . .	58
8.4	Results . . . . .	58
8.5	Comments . . . . .	59
8.6	Parameters . . . . .	59
8.7	TQC Policies . . . . .	65
<b>9</b>	<b>TRPO</b>	<b>69</b>
9.1	Notes . . . . .	69
9.2	Can I use? . . . . .	69
9.3	Example . . . . .	70
9.4	Results . . . . .	70
9.5	Parameters . . . . .	71
9.6	TRPO Policies . . . . .	76
<b>10</b>	<b>Utils</b>	<b>81</b>
<b>11</b>	<b>Gym Wrappers</b>	<b>83</b>
11.1	TimeFeatureWrapper . . . . .	83
<b>12</b>	<b>Changelog</b>	<b>85</b>
12.1	Release 1.6.2 (2022-10-10) . . . . .	85
12.2	Release 1.6.1 (2022-09-29) . . . . .	85
12.3	Release 1.6.0 (2022-07-11) . . . . .	86
12.4	Release 1.5.0 (2022-03-25) . . . . .	87
12.5	Release 1.4.0 (2022-01-19) . . . . .	87
12.6	Release 1.3.0 (2021-10-23) . . . . .	88
12.7	Release 1.2.0 (2021-09-08) . . . . .	89
12.8	Release 1.1.0 (2021-07-01) . . . . .	89
12.9	Release 1.0 (2021-03-17) . . . . .	90
12.10	Pre-Release 0.11.1 (2021-02-27) . . . . .	90
12.11	Pre-Release 0.11.0 (2021-02-27) . . . . .	91
12.12	Pre-Release 0.10.0 (2020-10-28) . . . . .	91
12.13	Maintainers . . . . .	92
12.14	Contributors: . . . . .	92
<b>13</b>	<b>Citing Stable Baselines3</b>	<b>93</b>
<b>14</b>	<b>Contributing</b>	<b>95</b>
<b>15</b>	<b>Indices and tables</b>	<b>97</b>
	<b>Python Module Index</b>	<b>99</b>
	<b>Index</b>	<b>101</b>

Contrib package for Stable Baselines3 (SB3) - Experimental code.

Github repository: <https://github.com/Stable-Baselines-Team/stable-baselines3-contrib>

SB3 repository: <https://github.com/DLR-RM/stable-baselines3>

RL Baselines3 Zoo (collection of pre-trained agents): <https://github.com/DLR-RM/rl-baselines3-zoo>

RL Baselines3 Zoo also offers a simple interface to train, evaluate agents and do hyperparameter tuning.



## INSTALLATION

### 1.1 Prerequisites

Please read [Stable-Baselines3 installation guide](#) first.

#### 1.1.1 Stable Release

To install Stable Baselines3 contrib with pip, execute:

```
pip install sb3-contrib
```

### 1.2 Bleeding-edge version

```
pip install git+https://github.com/Stable-Baselines-Team/stable-baselines3-contrib/
```

### 1.3 Development version

To contribute to Stable-Baselines3, with support for running tests and building the documentation.

```
git clone https://github.com/Stable-Baselines-Team/stable-baselines3-contrib/ && cd .  
↪ stable-baselines3-contrib  
pip install -e .
```





## RL ALGORITHMS

This table displays the rl algorithms that are implemented in the Stable Baselines3 contrib project, along with some useful characteristics: support for discrete/continuous actions, multiprocessing.

Name	Box	Discrete	MultiDiscrete	MultiBinary	Multi Processing
ARS	✓				✓
MaskablePPO		✓	✓	✓	✓
QR-DQN		✓			✓
RecurrentPPO	✓	✓	✓	✓	✓
TQC	✓				✓
TRPO	✓	✓	✓	✓	✓

**Note:** Tuple observation spaces are not supported by any environment, however, single-level Dict spaces are

Actions `gym.spaces`:

- **Box:** A N-dimensional box that contains every point in the action space.
- **Discrete:** A list of possible actions, where each timestep only one of the actions can be used.
- **MultiDiscrete:** A list of possible actions, where each timestep only one action of each discrete set can be used.
- **MultiBinary:** A list of possible actions, where each timestep any of the actions can be used in any combination.



## EXAMPLES

### 3.1 TQC

Train a Truncated Quantile Critics (TQC) agent on the Pendulum environment.

```
from sb3_contrib import TQC

model = TQC("MlpPolicy", "Pendulum-v1", top_quantiles_to_drop_per_net=2, verbose=1)
model.learn(total_timesteps=10_000, log_interval=4)
model.save("tqc_pendulum")
```

### 3.2 QR-DQN

Train a Quantile Regression DQN (QR-DQN) agent on the CartPole environment.

```
from sb3_contrib import QRDQN

policy_kwargs = dict(n_quantiles=50)
model = QRDQN("MlpPolicy", "CartPole-v1", policy_kwargs=policy_kwargs, verbose=1)
model.learn(total_timesteps=10_000, log_interval=4)
model.save("qrdqn_cartpole")
```

### 3.3 MaskablePPO

Train a PPO with invalid action masking agent on a toy environment.

```
from sb3_contrib import MaskablePPO
from sb3_contrib.common.envs import InvalidActionEnvDiscrete

env = InvalidActionEnvDiscrete(dim=80, n_invalid_actions=60)
model = MaskablePPO("MlpPolicy", env, verbose=1)
model.learn(5000)
model.save("maskable_toy_env")
```

## 3.4 TRPO

Train a Trust Region Policy Optimization (TRPO) agent on the Pendulum environment.

```
from sb3_contrib import TRPO

model = TRPO("MlpPolicy", "Pendulum-v1", gamma=0.9, verbose=1)
model.learn(total_timesteps=100_000, log_interval=4)
model.save("trpo_pendulum")
```

## 3.5 ARS

Train an agent using Augmented Random Search (ARS) agent on the Pendulum environment

```
from sb3_contrib import ARS

model = ARS("LinearPolicy", "Pendulum-v1", verbose=1)
model.learn(total_timesteps=10000, log_interval=4)
model.save("ars_pendulum")
```

## 3.6 RecurrentPPO

Train a PPO agent with a recurrent policy on the CartPole environment.

---

**Note:** It is particularly important to pass the `lstm_states` and `episode_start` argument to the `predict()` method, so the cell and hidden states of the LSTM are correctly updated.

---

```
import numpy as np

from sb3_contrib import RecurrentPPO

model = RecurrentPPO("MlpLstmPolicy", "CartPole-v1", verbose=1)
model.learn(5000)

env = model.get_env()
obs = env.reset()
# cell and hidden state of the LSTM
lstm_states = None
num_envs = 1
# Episode start signals are used to reset the lstm states
episode_starts = np.ones((num_envs,), dtype=bool)
while True:
    action, lstm_states = model.predict(obs, state=lstm_states, episode_start=episode_
    ↪starts, deterministic=True)
    obs, rewards, dones, info = env.step(action)
    episode_starts = dones
    env.render()
```

## ARS

Augmented Random Search (ARS) is a simple reinforcement algorithm that uses a direct random search over policy parameters. It can be surprisingly effective compared to more sophisticated algorithms. In the [original paper](#) the authors showed that linear policies trained with ARS were competitive with deep reinforcement learning for the MuJuCo locomotion tasks.

SB3s implementation allows for linear policies without bias or squashing function, it also allows for training MLP policies, which include linear policies with bias and squashing functions as a special case.

Normally one wants to train ARS with several seeds to properly evaluate.

**Warning:** ARS multi-processing is different from the classic Stable-Baselines3 multi-processing: it runs  $n$  environments in parallel but asynchronously. This asynchronous multi-processing is considered experimental and does not fully support callbacks: the `on_step()` event is called artificially after the evaluation episodes are over.

### Available Policies

<i>LinearPolicy</i>	alias of <code>ARSLinearPolicy</code>
<i>MlpPolicy</i>	alias of <code>ARSPolicy</code>

## 4.1 Notes

- Original paper: <https://arxiv.org/abs/1803.07055>
- Original Implementation: <https://github.com/modestyachts/ARS>

## 4.2 Can I use?

- Recurrent policies:
- Multi processing: ✓ (cf. example)
- Gym spaces:

Space	Action	Observation
Discrete	✓	✓
Box	✓	✓
MultiDiscrete		✓
MultiBinary		✓
Dict		

## 4.3 Example

```
from sb3_contrib import ARS

# Policy can be LinearPolicy or MlpPolicy
model = ARS("LinearPolicy", "Pendulum-v1", verbose=1)
model.learn(total_timesteps=10000, log_interval=4)
model.save("ars_pendulum")
```

With experimental asynchronous multi-processing:

```
from sb3_contrib import ARS
from sb3_contrib.common.vec_env import AsyncEval

from stable_baselines3.common.env_util import make_vec_env

env_id = "CartPole-v1"
n_envs = 2

model = ARS("LinearPolicy", env_id, n_delta=2, n_top=1, verbose=1)
# Create env for asynchronous evaluation (run in different processes)
async_eval = AsyncEval([lambda: make_vec_env(env_id) for _ in range(n_envs)], model.
    ↪ policy)

model.learn(total_timesteps=200_000, log_interval=4, async_eval=async_eval)
```

## 4.4 Results

Replicating results from the original paper, which used the Mujoco benchmarks. Same parameters from the original paper, using 8 seeds.

Environments	ARS
HalfCheetah	4398 +/- 320
Swimmer	241 +/- 51
Hopper	3320 +/- 120

### 4.4.1 How to replicate the results?

Clone RL-Zoo and checkout the branch `feat/ars`

```
git clone https://github.com/DLR-RM/rl-baselines3-zoo
cd rl-baselines3-zoo/
git checkout feat/ars
```

Run the benchmark. The following code snippet trains 8 seeds in parallel

```
for ENV_ID in Swimmer-v3 HalfCheetah-v3 Hopper-v3
do
    for SEED_NUM in {1..8}
    do
        SEED=$RANDOM
        python train.py --algo ars --env $ENV_ID --eval-episodes 10 --eval-freq 10000 -n_
↪200000000 --seed $SEED &
        sleep 1
    done
done
wait
```

Plot the results:

```
python scripts/all_plots.py -a ars -e HalfCheetah Swimmer Hopper -f logs/ -o logs/ars_
↪results -max 200000000
python scripts/plot_from_file.py -i logs/ars_results.pkl -l ARS
```

## 4.5 Parameters

```
class sb3_contrib.ars.ARS(policy, env, n_delta=8, n_top=None, learning_rate=0.02, delta_std=0.05,
                           zero_policy=True, alive_bonus_offset=0, n_eval_episodes=1,
                           policy_kwargs=None, tensorboard_log=None, seed=None, verbose=0,
                           device='cpu', _init_setup_model=True)
```

Augmented Random Search: <https://arxiv.org/abs/1803.07055>

Original implementation: <https://github.com/modestyachts/ARS> C++/Cuda Implementation: <https://github.com/google-research/tiny-differentiable-simulator/> 150 LOC Numpy Implementation: [https://github.com/alexis-jacq/numpy\\_ARS/blob/master/asr.py](https://github.com/alexis-jacq/numpy_ARS/blob/master/asr.py)

### Parameters

- **policy** (Union[str, Type[ARSPolicy]]) – The policy to train, can be an instance of ARSPolicy, or a string from ["LinearPolicy", "MlpPolicy"]
- **env** (Union[Env, VecEnv, str]) – The environment to train on, may be a string if registered with gym
- **n\_delta** (int) – How many random perturbations of the policy to try at each update step.
- **n\_top** (Optional[int]) – How many of the top delta to use in each update step. Default is n\_delta
- **learning\_rate** (Union[float, Callable[[float], float]]) – Float or schedule for the step size

- **delta\_std** (Union[float, Callable[[float], float]]) – Float or schedule for the exploration noise
- **zero\_policy** (bool) – Boolean determining if the passed policy should have it's weights zeroed before training.
- **alive\_bonus\_offset** (float) – Constant added to the reward at each step, used to cancel out alive bonuses.
- **n\_eval\_episodes** (int) – Number of episodes to evaluate each candidate.
- **policy\_kwargs** (Optional[Dict[str, Any]]) – Keyword arguments to pass to the policy on creation
- **tensorboard\_log** (Optional[str]) – String with the directory to put tensorboard logs:
- **seed** (Optional[int]) – Random seed for the training
- **verbose** (int) – Verbosity level: 0 no output, 1 info, 2 debug
- **device** (Union[device, str]) – Torch device to use for training, defaults to “cpu”
- **\_init\_setup\_model** (bool) – Whether or not to build the network at the creation of the instance

**evaluate\_candidates**(*candidate\_weights, callback, async\_eval*)

Evaluate each candidate.

**Parameters**

- **candidate\_weights** (Tensor) – The candidate weights to be evaluated.
- **callback** (BaseCallback) – Callback that will be called at each step (or after evaluation in the multiprocessing version)
- **async\_eval** (Optional[AsyncEval]) – The object for asynchronous evaluation of candidates.

**Return type**

Tensor

**Returns**

The episodic return for each candidate.

**get\_env()**

Returns the current environment (can be None if not defined).

**Return type**

Optional[VecEnv]

**Returns**

The current environment

**get\_parameters()**

Return the parameters of the agent. This includes parameters from different networks, e.g. critics (value functions) and policies (pi functions).

**Return type**

Dict[str, Dict]

**Returns**

Mapping of from names of the objects to PyTorch state-dicts.



**get\_vec\_normalize\_env()**

Return the VecNormalize wrapper of the training env if it exists.

**Return type**

Optional[VecNormalize]

**Returns**

The VecNormalize env.

**learn**(*total\_timesteps*, *callback=None*, *log\_interval=1*, *tb\_log\_name='ARS'*, *eval\_env=None*, *eval\_freq=-1*, *n\_eval\_episodes=5*, *eval\_log\_path=None*, *reset\_num\_timesteps=True*, *async\_eval=None*, *progress\_bar=False*)

Return a trained model.

**Parameters**

- **total\_timesteps** (int) – The total number of samples (env steps) to train on
- **callback** (Union[None, Callable, List[BaseCallback], BaseCallback]) – call-back(s) called at every step with state of the algorithm.
- **log\_interval** (int) – The number of timesteps before logging.
- **tb\_log\_name** (str) – the name of the run for TensorBoard logging
- **eval\_env** (Union[Env, VecEnv, None]) – Environment to use for evaluation. Caution, this parameter is deprecated and will be removed in the future. Please use *EvalCallback* or a custom Callback instead.
- **eval\_freq** (int) – Evaluate the agent every *eval\_freq* timesteps (this may vary a little). Caution, this parameter is deprecated and will be removed in the future. Please use *EvalCallback* or a custom Callback instead.
- **n\_eval\_episodes** (int) – Number of episode to evaluate the agent
- **eval\_log\_path** (Optional[str]) – Path to a folder where the evaluations will be saved
- **reset\_num\_timesteps** (bool) – whether or not to reset the current timestep number (used in logging)
- **async\_eval** (Optional[AsyncEval]) – The object for asynchronous evaluation of candidates.
- **progress\_bar** (bool) – Display a progress bar using tqdm and rich.

**Return type**

TypeVar(ARSSelf, bound= ARS)

**Returns**

the trained model

**classmethod load**(*path*, *env=None*, *device='auto'*, *custom\_objects=None*, *print\_system\_info=False*, *force\_reset=True*, *\*\*kwargs*)

Load the model from a zip-file. Warning: load re-creates the model from scratch, it does not update it in-place! For an in-place load use *set\_parameters* instead.

**Parameters**

- **path** (Union[str, Path, BufferedIOBase]) – path to the file (or a file-like) where to load the agent from
- **env** (Union[Env, VecEnv, None]) – the new environment to run the loaded model on (can be None if you only need prediction from a trained model) has priority over any saved environment

- **device** (Union[device, str]) – Device on which the code should run.
- **custom\_objects** (Optional[Dict[str, Any]]) – Dictionary of objects to replace upon loading. If a variable is present in this dictionary as a key, it will not be deserialized and the corresponding item will be used instead. Similar to custom\_objects in `keras.models.load_model`. Useful when you have an object in file that can not be deserialized.
- **print\_system\_info** (bool) – Whether to print system info from the saved model and the current system info (useful to debug loading issues)
- **force\_reset** (bool) – Force call to `reset()` before training to avoid unexpected behavior. See <https://github.com/DLR-RM/stable-baselines3/issues/597>
- **kwargs** – extra arguments to change the model when loading

**Return type**

TypeVar(SelfBaseAlgorithm, bound= BaseAlgorithm)

**Returns**

new model instance with loaded parameters

**property logger: Logger**

Getter for the logger object.

**predict**(*observation, state=None, episode\_start=None, deterministic=False*)

Get the policy action from an observation (and optional hidden state). Includes sugar-coating to handle different observations (e.g. normalizing images).

**Parameters**

- **observation** (Union[ndarray, Dict[str, ndarray]]) – the input observation
- **state** (Optional[Tuple[ndarray, ...]]) – The last hidden states (can be None, used in recurrent policies)
- **episode\_start** (Optional[ndarray]) – The last masks (can be None, used in recurrent policies) this correspond to beginning of episodes, where the hidden states of the RNN must be reset.
- **deterministic** (bool) – Whether or not to return deterministic actions.

**Return type**

Tuple[ndarray, Optional[Tuple[ndarray, ...]]]

**Returns**

the model's action and the next hidden state (used in recurrent policies)

**save**(*path, exclude=None, include=None*)

Save all the attributes of the object and the model parameters in a zip-file.

**Parameters**

- **path** (Union[str, Path, BufferedIOBase]) – path to the file where the rl agent should be saved
- **exclude** (Optional[Iterable[str]]) – name of parameters that should be excluded in addition to the default ones
- **include** (Optional[Iterable[str]]) – name of parameters that might be excluded but should be included anyway

**Return type**

None

**set\_env**(*env*, *force\_reset=True*)

Checks the validity of the environment, and if it is coherent, set it as the current environment. Furthermore wrap any non vectorized env into a vectorized checked parameters: - observation\_space - action\_space

**Parameters**

- **env** (Union[Env, VecEnv]) – The environment for learning a policy
- **force\_reset** (bool) – Force call to `reset()` before training to avoid unexpected behavior. See issue <https://github.com/DLR-RM/stable-baselines3/issues/597>

**Return type**

None

**set\_logger**(*logger*)

Setter for for logger object. :rtype: None

**Warning:** When passing a custom logger object, this will overwrite `tensorboard_log` and `verbose` settings passed to the constructor.

**set\_parameters**(*load\_path\_or\_dict*, *exact\_match=True*, *device='auto'*)

Load parameters from a given zip-file or a nested dictionary containing parameters for different modules (see `get_parameters`).

**Parameters**

- **load\_path\_or\_iter** – Location of the saved data (path or file-like, see `save`), or a nested dictionary containing `nn.Module` parameters used by the policy. The dictionary maps object names to a state-dictionary returned by `torch.nn.Module.state_dict()`.
- **exact\_match** (bool) – If True, the given parameters should include parameters for each module and each of their parameters, otherwise raises an Exception. If set to False, this can be used to update only specific parameters.
- **device** (Union[device, str]) – Device on which the code should run.

**Return type**

None

**set\_random\_seed**(*seed=None*)

Set the seed of the pseudo-random generators (python, numpy, pytorch, gym, action\_space)

**Parameters**

**seed** (Optional[int]) –

**Return type**

None

## 4.6 ARS Policies

```
class sb3_contrib.ars.policies.ARSPolicy(observation_space, action_space, net_arch=None,  
                                         activation_fn=<class 'torch.nn.modules.activation.ReLU'>,  
                                         squash_output=True)
```

Policy network for ARS.

### Parameters

- **observation\_space** (Space) – The observation space of the environment
- **action\_space** (Space) – The action space of the environment
- **net\_arch** (Optional[List[int]]) – Network architecture, defaults to a 2 layers MLP with 64 hidden nodes.
- **activation\_fn** (Type[Module]) – Activation function
- **squash\_output** (bool) – For continuous actions, whether the output is squashed or not using a `tanh()` function. If not squashed with `tanh` the output will instead be clipped.

### **forward**(obs)

Defines the computation performed at every call.

Should be overridden by all subclasses. :rtype: Tensor

---

**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.

---

**sb3\_contrib.ars.LinearPolicy**

alias of `ARSLinearPolicy`

**sb3\_contrib.ars.MlpPolicy**

alias of `ARSPolicy`

## MASKABLE PPO

Implementation of **invalid action masking** for the Proximal Policy Optimization (PPO) algorithm. Other than adding support for action masking, the behavior is the same as in SB3's core PPO algorithm.

### Available Policies

<i>MlpPolicy</i>	alias of MaskableActorCriticPolicy
<i>CnnPolicy</i>	alias of MaskableActorCriticCnnPolicy
<i>MultiInputPolicy</i>	alias of MaskableMultiInputActorCriticPolicy

## 5.1 Notes

- Paper: <https://arxiv.org/abs/2006.14171>
- Blog post: <https://costa.sh/blog-a-closer-look-at-invalid-action-masking-in-policy-gradient-algorithms.html>
- Additional Blog post: <https://boring-guy.sh/posts/masking-rl/>

## 5.2 Can I use?

- Recurrent policies:
- Multi processing: ✓
- Gym spaces:

Space	Action	Observation
Discrete	✓	✓
Box		✓
MultiDiscrete	✓	✓
MultiBinary	✓	✓
Dict		✓

## 5.3 Example

Train a PPO agent on `InvalidActionEnvDiscrete`. `InvalidActionEnvDiscrete` has a `action_masks` method that returns the invalid action mask (True if the action is valid, False otherwise).

```
from sb3_contrib import MaskablePPO
from sb3_contrib.common.envs import InvalidActionEnvDiscrete
from sb3_contrib.common.maskable.evaluation import evaluate_policy
from sb3_contrib.common.maskable.utils import get_action_masks

env = InvalidActionEnvDiscrete(dim=80, n_invalid_actions=60)
model = MaskablePPO("MlpPolicy", env, gamma=0.4, seed=32, verbose=1)
model.learn(5000)

evaluate_policy(model, env, n_eval_episodes=20, reward_threshold=90, warn=False)

model.save("ppo_mask")
del model # remove to demonstrate saving and loading

model = MaskablePPO.load("ppo_mask")

obs = env.reset()
while True:
    # Retrieve current action mask
    action_masks = get_action_masks(env)
    action, _states = model.predict(obs, action_masks=action_masks)
    obs, rewards, dones, info = env.step(action)
    env.render()
```

If the environment implements the invalid action mask but using a different name, you can use the `ActionMasker` to specify the name (see [PR #25](#)):

```
import gym
import numpy as np

from sb3_contrib.common.maskable.policies import MaskableActorCriticPolicy
from sb3_contrib.common.wrappers import ActionMasker
from sb3_contrib.ppo_mask import MaskablePPO

def mask_fn(env: gym.Env) -> np.ndarray:
    # Do whatever you'd like in this function to return the action mask
    # for the current env. In this example, we assume the env has a
    # helpful method we can rely on.
    return env.valid_action_mask()

env = ... # Initialize env
env = ActionMasker(env, mask_fn) # Wrap to enable masking

# MaskablePPO behaves the same as SB3's PPO unless the env is wrapped
# with ActionMasker. If the wrapper is detected, the masks are automatically
```

(continues on next page)

(continued from previous page)

```
# retrieved and used when learning. Note that MaskablePPO does not accept
# a new action_mask_fn kwarg, as it did in an earlier draft.
model = MaskablePPO(MaskableActorCriticPolicy, env, verbose=1)
model.learn()


# Note that use of masks is manual and optional outside of learning,
# so masking can be "removed" at testing time
model.predict(observation, action_masks=valid_action_array)
```

## 5.4 Results


Results are shown for two MicroRTS benchmarks: MicrortsMining4x4F9-v0 (600K steps) and MicrortsMining10x10F9-v0 (1.5M steps). For each, models were trained with and without masking, using 3 seeds.

### 5.4.1 4x4


#### No masking



## With masking




## Combined






## 5.4.2 10x10


### No masking



### With masking



## Combined



More information may be found in the [associated PR](#).

### 5.4.3 How to replicate the results?

Clone the repo for the experiment:

```
git clone git@github.com:kronion/microrsts-ppo-comparison.git
cd microrsts-ppo-comparison
```

Install dependencies:

```
# Install MicroRTS:
rm -fR ~/microrsts && mkdir ~/microrsts && \
  wget -O ~/microrsts/microrsts.zip http://microrsts.s3.amazonaws.com/microrsts/artifacts/
  ↪ 202004222224.microrsts.zip && \
  unzip ~/microrsts/microrsts.zip -d ~/microrsts/

# You may want to make a venv before installing packages
pip install -r requirements.txt
```

Train several times with various seeds, with and without masking:

```
# python sb/train_ppo.py [output dir] [MicroRTS map size] [--mask] [--seed int]

# 4x4 unmasked
python sb3/train_ppo.py zoo 4 --seed 42
python sb3/train_ppo.py zoo 4 --seed 43
python sb3/train_ppo.py zoo 4 --seed 44

# 4x4 masked
python sb3/train_ppo.py zoo 4 --mask --seed 42
python sb3/train_ppo.py zoo 4 --mask --seed 43
python sb3/train_ppo.py zoo 4 --mask --seed 44

# 10x10 unmasked
python sb3/train_ppo.py zoo 10 --seed 42
python sb3/train_ppo.py zoo 10 --seed 43
python sb3/train_ppo.py zoo 10 --seed 44

# 10x10 masked
python sb3/train_ppo.py zoo 10 --mask --seed 42
python sb3/train_ppo.py zoo 10 --mask --seed 43
python sb3/train_ppo.py zoo 10 --mask --seed 44
```

View the tensorboard log output:

```
# For 4x4 environment
tensorboard --logdir zoo/4x4/runs

# For 10x10 environment
tensorboard --logdir zoo/10x10/runs
```

## 5.5 Parameters

```
class sb3_contrib.ppo_mask.MaskablePPO(policy, env, learning_rate=0.0003, n_steps=2048, batch_size=64,
                                       n_epochs=10, gamma=0.99, gae_lambda=0.95, clip_range=0.2,
                                       clip_range_vf=None, normalize_advantage=True, ent_coef=0.0,
                                       vf_coef=0.5, max_grad_norm=0.5, target_kl=None,
                                       tensorboard_log=None, create_eval_env=False,
                                       policy_kwargs=None, verbose=0, seed=None, device='auto',
                                       _init_setup_model=True)
```

Proximal Policy Optimization algorithm (PPO) (clip version) with Invalid Action Masking.

Based on the original Stable Baselines 3 implementation.

Introduction to PPO: <https://spinningup.openai.com/en/latest/algorithms/ppo.html> Background on Invalid Action Masking: <https://arxiv.org/abs/2006.14171>

### Parameters

- **policy** (Union[str, Type[MaskableActorCriticPolicy]]) – The policy model to use (MlpPolicy, CnnPolicy, ...)
- **env** (Union[Env, VecEnv, str]) – The environment to learn from (if registered in Gym, can be str)
- **learning\_rate** (Union[float, Callable[[float], float]]) – The learning rate, it can be a function of the current progress remaining (from 1 to 0)
- **n\_steps** (int) – The number of steps to run for each environment per update (i.e. batch size is `n_steps * n_env` where `n_env` is number of environment copies running in parallel)
- **batch\_size** (Optional[int]) – Minibatch size
- **n\_epochs** (int) – Number of epoch when optimizing the surrogate loss
- **gamma** (float) – Discount factor
- **gae\_lambda** (float) – Factor for trade-off of bias vs variance for Generalized Advantage Estimator
- **clip\_range** (Union[float, Callable[[float], float]]) – Clipping parameter, it can be a function of the current progress remaining (from 1 to 0).
- **clip\_range\_vf** (Union[None, float, Callable[[float], float]]) – Clipping parameter for the value function, it can be a function of the current progress remaining (from 1 to 0). This is a parameter specific to the OpenAI implementation. If None is passed (default), no clipping will be done on the value function. **IMPORTANT:** this clipping depends on the reward scaling.
- **normalize\_advantage** (bool) – Whether to normalize or not the advantage
- **ent\_coef** (float) – Entropy coefficient for the loss calculation
- **vf\_coef** (float) – Value function coefficient for the loss calculation
- **max\_grad\_norm** (float) – The maximum value for the gradient clipping
- **target\_kl** (Optional[float]) – Limit the KL divergence between updates, because the clipping is not enough to prevent large update see issue #213 (cf <https://github.com/hill-a/stable-baselines/issues/213>) By default, there is no limit on the kl div.
- **tensorboard\_log** (Optional[str]) – the log location for tensorboard (if None, no logging)

- **create\_eval\_env** (bool) – Whether to create a second environment that will be used for evaluating the agent periodically (Only available when passing string for the environment). Caution, this parameter is deprecated and will be removed in the future.
- **policy\_kwargs** (Optional[Dict[str, Any]]) – additional arguments to be passed to the policy on creation
- **verbose** (int) – the verbosity level: 0 no output, 1 info, 2 debug
- **seed** (Optional[int]) – Seed for the pseudo random generators
- **device** (Union[device, str]) – Device (cpu, cuda, ...) on which the code should be run. Setting it to auto, the code will be run on the GPU if possible.
- **\_init\_setup\_model** (bool) – Whether or not to build the network at the creation of the instance

**collect\_rollouts**(*env, callback, rollout\_buffer, n\_rollout\_steps, use\_masking=True*)

Collect experiences using the current policy and fill a `RolloutBuffer`. The term rollout here refers to the model-free notion and should not be used with the concept of rollout used in model-based RL or planning.

This method is largely identical to the implementation found in the parent class.

#### Parameters

- **env** (VecEnv) – The training environment
- **callback** (BaseCallback) – Callback that will be called at each step (and at the beginning and end of the rollout)
- **rollout\_buffer** (RolloutBuffer) – Buffer to fill with rollouts
- **n\_steps** – Number of experiences to collect per environment
- **use\_masking** (bool) – Whether or not to use invalid action masks during training

#### Return type

bool

#### Returns

True if function returned with at least *n\_rollout\_steps* collected, False if callback terminated rollout prematurely.

**get\_env()**

Returns the current environment (can be None if not defined).

#### Return type

Optional[VecEnv]

#### Returns

The current environment

**get\_parameters()**

Return the parameters of the agent. This includes parameters from different networks, e.g. critics (value functions) and policies (pi functions).

#### Return type

Dict[str, Dict]

#### Returns

Mapping of from names of the objects to PyTorch state-dicts.

**get\_vec\_normalize\_env()**

Return the VecNormalize wrapper of the training env if it exists.

**Return type**

Optional[VecNormalize]

**Returns**

The VecNormalize env.

**learn**(*total\_timesteps*, *callback=None*, *log\_interval=1*, *eval\_env=None*, *eval\_freq=-1*, *n\_eval\_episodes=5*, *tb\_log\_name='PPO'*, *eval\_log\_path=None*, *reset\_num\_timesteps=True*, *use\_masking=True*, *progress\_bar=False*)

Return a trained model.

**Parameters**

- **total\_timesteps** (int) – The total number of samples (env steps) to train on
- **callback** (Union[None, Callable, List[BaseCallback], BaseCallback]) – call-back(s) called at every step with state of the algorithm.
- **log\_interval** (int) – The number of timesteps before logging.
- **tb\_log\_name** (str) – the name of the run for TensorBoard logging
- **reset\_num\_timesteps** (bool) – whether or not to reset the current timestep number (used in logging)
- **progress\_bar** (bool) – Display a progress bar using tqdm and rich.

**Return type**

TypeVar(MaskablePPOSelf, bound= MaskablePPO)

**Returns**

the trained model

**classmethod load**(*path*, *env=None*, *device='auto'*, *custom\_objects=None*, *print\_system\_info=False*, *force\_reset=True*, *\*\*kwargs*)

Load the model from a zip-file. Warning: load re-creates the model from scratch, it does not update it in-place! For an in-place load use `set_parameters` instead.

**Parameters**

- **path** (Union[str, Path, BufferedIOBase]) – path to the file (or a file-like) where to load the agent from
- **env** (Union[Env, VecEnv, None]) – the new environment to run the loaded model on (can be None if you only need prediction from a trained model) has priority over any saved environment
- **device** (Union[device, str]) – Device on which the code should run.
- **custom\_objects** (Optional[Dict[str, Any]]) – Dictionary of objects to replace upon loading. If a variable is present in this dictionary as a key, it will not be deserialized and the corresponding item will be used instead. Similar to `custom_objects` in `keras.models.load_model`. Useful when you have an object in file that can not be deserialized.
- **print\_system\_info** (bool) – Whether to print system info from the saved model and the current system info (useful to debug loading issues)
- **force\_reset** (bool) – Force call to `reset()` before training to avoid unexpected behavior. See <https://github.com/DLR-RM/stable-baselines3/issues/597>
- **kwargs** – extra arguments to change the model when loading

**Return type**

TypeVar(SelfBaseAlgorithm, bound= BaseAlgorithm)

**Returns**

new model instance with loaded parameters

**property logger: Logger**

Getter for the logger object.

**predict**(*observation*, *state=None*, *episode\_start=None*, *deterministic=False*, *action\_masks=None*)

Get the policy action from an observation (and optional hidden state). Includes sugar-coating to handle different observations (e.g. normalizing images).

**Parameters**

- **observation** (ndarray) – the input observation
- **state** (Optional[Tuple[ndarray, ...]]) – The last hidden states (can be None, used in recurrent policies)
- **episode\_start** (Optional[ndarray]) – The last masks (can be None, used in recurrent policies) this correspond to beginning of episodes, where the hidden states of the RNN must be reset.
- **deterministic** (bool) – Whether or not to return deterministic actions.

**Return type**

Tuple[ndarray, Optional[Tuple[ndarray, ...]]]

**Returns**

the model's action and the next hidden state (used in recurrent policies)

**save**(*path*, *exclude=None*, *include=None*)

Save all the attributes of the object and the model parameters in a zip-file.

**Parameters**

- **path** (Union[str, Path, BufferedIOBase]) – path to the file where the rl agent should be saved
- **exclude** (Optional[Iterable[str]]) – name of parameters that should be excluded in addition to the default ones
- **include** (Optional[Iterable[str]]) – name of parameters that might be excluded but should be included anyway

**Return type**

None

**set\_env**(*env*, *force\_reset=True*)

Checks the validity of the environment, and if it is coherent, set it as the current environment. Furthermore wrap any non vectorized env into a vectorized checked parameters: - observation\_space - action\_space

**Parameters**

- **env** (Union[Env, VecEnv]) – The environment for learning a policy
- **force\_reset** (bool) – Force call to `reset()` before training to avoid unexpected behavior. See issue <https://github.com/DLR-RM/stable-baselines3/issues/597>

**Return type**

None

**set\_logger**(*logger*)

Setter for for logger object. :rtype: None

**Warning:** When passing a custom logger object, this will overwrite `tensorboard_log` and `verbose` settings passed to the constructor.

**set\_parameters**(*load\_path\_or\_dict*, *exact\_match=True*, *device='auto'*)

Load parameters from a given zip-file or a nested dictionary containing parameters for different modules (see `get_parameters`).

**Parameters**

- **load\_path\_or\_iter** – Location of the saved data (path or file-like, see `save`), or a nested dictionary containing `nn.Module` parameters used by the policy. The dictionary maps object names to a state-dictionary returned by `torch.nn.Module.state_dict()`.
- **exact\_match** (bool) – If True, the given parameters should include parameters for each module and each of their parameters, otherwise raises an Exception. If set to False, this can be used to update only specific parameters.
- **device** (Union[device, str]) – Device on which the code should run.

**Return type**

None

**set\_random\_seed**(*seed=None*)

Set the seed of the pseudo-random generators (python, numpy, pytorch, gym, action\_space)

**Parameters**

**seed** (Optional[int]) –

**Return type**

None

**train**()

Update policy using the currently gathered rollout buffer.

**Return type**

None

## 5.6 MaskablePPO Policies

`sb3_contrib.ppo_mask.MlpPolicy`

alias of `MaskableActorCriticPolicy`



```
class sb3_contrib.common.maskable.policies.MaskableActorCriticPolicy(observation_space,
                                                                    action_space, lr_schedule,
                                                                    net_arch=None,
                                                                    activation_fn=<class
                                                                    'torch.nn.modules.activation.Tanh'>,
                                                                    ortho_init=True, fea-
                                                                    tures_extractor_class=<class
                                                                    'sta-
                                                                    ble_baselines3.common.torch_layers.Flatten
                                                                    fea-
                                                                    tures_extractor'>,
                                                                    normalize_images=True,
                                                                    optimizer_class=<class
                                                                    'torch.optim.adam.Adam'>,
                                                                    optimizer_kwargs=None)
```

Policy class for actor-critic algorithms (has both policy and value prediction). Used by A2C, PPO and the likes.

#### Parameters

- **observation\_space** (Space) – Observation space
- **action\_space** (Space) – Action space
- **lr\_schedule** (Callable[[float], float]) – Learning rate schedule (could be constant)
- **net\_arch** (Optional[List[Union[int, Dict[str, List[int]]]]) – The specification of the policy and value networks.
- **activation\_fn** (Type[Module]) – Activation function
- **ortho\_init** (bool) – Whether to use or not orthogonal initialization
- **features\_extractor\_class** (Type[BaseFeaturesExtractor]) – Features extractor to use.
- **features\_extractor\_kwargs** (Optional[Dict[str, Any]]) – Keyword arguments to pass to the features extractor.
- **normalize\_images** (bool) – Whether to normalize images or not, dividing by 255.0 (True by default)
- **optimizer\_class** (Type[Optimizer]) – The optimizer to use, `th.optim.Adam` by default
- **optimizer\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments, excluding the learning rate, to pass to the optimizer

**evaluate\_actions**(*obs*, *actions*, *action\_masks=None*)

Evaluate actions according to the current policy, given the observations.

#### Parameters

- **obs** (Tensor) –
- **actions** (Tensor) –

#### Return type

Tuple[Tensor, Tensor, Tensor]

#### Returns

estimated value, log likelihood of taking those actions and entropy of the action distribution.

**forward**(*obs*, *deterministic=False*, *action\_masks=None*)

Forward pass in all the networks (actor and critic)

**Parameters**

- **obs** (Tensor) – Observation
- **deterministic** (bool) – Whether to sample or use deterministic actions
- **action\_masks** (Optional[ndarray]) – Action masks to apply to the action distribution

**Return type**

Tuple[Tensor, Tensor, Tensor]

**Returns**

action, value and log probability of the action

**get\_distribution**(*obs*, *action\_masks=None*)

Get the current policy distribution given the observations.

**Parameters**

- **obs** (Tensor) –
- **action\_masks** (Optional[ndarray]) –

**Return type**

MaskableDistribution

**Returns**

the action distribution.

**predict**(*observation*, *state=None*, *episode\_start=None*, *deterministic=False*, *action\_masks=None*)

Get the policy action from an observation (and optional hidden state). Includes sugar-coating to handle different observations (e.g. normalizing images).

**Parameters**

- **observation** (Union[ndarray, Dict[str, ndarray]]) – the input observation
- **state** (Optional[Tuple[ndarray, ...]]) – The last states (can be None, used in recurrent policies)
- **episode\_start** (Optional[ndarray]) – The last masks (can be None, used in recurrent policies)
- **deterministic** (bool) – Whether or not to return deterministic actions.
- **action\_masks** (Optional[ndarray]) – Action masks to apply to the action distribution

**Return type**

Tuple[ndarray, Optional[Tuple[ndarray, ...]]]

**Returns**

the model's action and the next state (used in recurrent policies)

**predict\_values**(*obs*)

Get the estimated values according to the current policy given the observations.

**Parameters**

- **obs** (Tensor) –

**Return type**

Tensor

**Returns**

the estimated values.

`sb3_contrib.ppo_mask.CnnPolicy`

alias of `MaskableActorCriticCnnPolicy`

```
class sb3_contrib.common.maskable.policies.MaskableActorCriticCnnPolicy(observation_space,
                                                                    action_space,
                                                                    lr_schedule,
                                                                    net_arch=None,
                                                                    activation_fn=<class
                                                                    'torch.nn.modules.activation.Tanh'>,
                                                                    ortho_init=True, fea-
                                                                    tures_extractor_class=<class
                                                                    'sta-
                                                                    ble_baselines3.common.torch_layers.Nat
                                                                    fea-
                                                                    tures_extractor_kwargs=None,
                                                                    normal-
                                                                    ize_images=True,
                                                                    opti-
                                                                    mizer_class=<class
                                                                    'torch.optim.adam.Adam'>,
                                                                    opti-
                                                                    mizer_kwargs=None)
```

CNN policy class for actor-critic algorithms (has both policy and value prediction). Used by A2C, PPO and the likes.

**Parameters**

- **observation\_space** (Space) – Observation space
- **action\_space** (Space) – Action space
- **lr\_schedule** (Callable[[float], float]) – Learning rate schedule (could be constant)
- **net\_arch** (Optional[List[Union[int, Dict[str, List[int]]]]) – The specification of the policy and value networks.
- **activation\_fn** (Type[Module]) – Activation function
- **ortho\_init** (bool) – Whether to use or not orthogonal initialization
- **features\_extractor\_class** (Type[BaseFeaturesExtractor]) – Features extractor to use.
- **features\_extractor\_kwargs** (Optional[Dict[str, Any]]) – Keyword arguments to pass to the features extractor.
- **normalize\_images** (bool) – Whether to normalize images or not, dividing by 255.0 (True by default)
- **optimizer\_class** (Type[Optimizer]) – The optimizer to use, `th.optim.Adam` by default
- **optimizer\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments, excluding the learning rate, to pass to the optimizer

`sb3_contrib.ppo_mask.MultiInputPolicy`

alias of `MaskableMultiInputActorCriticPolicy`

```
class sb3_contrib.common.maskable.policies.MaskableMultiInputActorCriticPolicy(observation_space,
                                                                                action_space,
                                                                                lr_schedule,
                                                                                net_arch=None,
                                                                                activation_fn=<class
'torch.nn.modules.activation.Tanh'>,
                                                                                ortho_init=True,
                                                                                features_extractor_class=<class
'stable_baselines3.common.torch_layers.CombinedExtractor'>,
                                                                                features_extractor_kwargs=None,
                                                                                normalize_images=True,
                                                                                optimizer_class=<class
'torch.optim.adam.Adam'>,
                                                                                optimizer_kwargs=None)
```

MultiInputActorClass policy class for actor-critic algorithms (has both policy and value prediction). Used by A2C, PPO and the likes.

#### Parameters

- **observation\_space** (Dict) – Observation space (Tuple)
- **action\_space** (Space) – Action space
- **lr\_schedule** (Callable[[float], float]) – Learning rate schedule (could be constant)
- **net\_arch** (Optional[List[Union[int, Dict[str, List[int]]]]) – The specification of the policy and value networks.
- **activation\_fn** (Type[Module]) – Activation function
- **ortho\_init** (bool) – Whether to use or not orthogonal initialization
- **features\_extractor\_class** (Type[BaseFeaturesExtractor]) – Uses the CombinedExtractor
- **features\_extractor\_kwargs** (Optional[Dict[str, Any]]) – Keyword arguments to pass to the feature extractor.
- **normalize\_images** (bool) – Whether to normalize images or not, dividing by 255.0 (True by default)
- **optimizer\_class** (Type[Optimizer]) – The optimizer to use, th.optim.Adam by default
- **optimizer\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments, excluding the learning rate, to pass to the optimizer

## RECURRENT PPO

Implementation of recurrent policies for the Proximal Policy Optimization (PPO) algorithm. Other than adding support for recurrent policies (LSTM here), the behavior is the same as in SB3's core PPO algorithm.

### Available Policies

<i>MlpLstmPolicy</i>	alias of <code>RecurrentActorCriticPolicy</code>
<i>CnnLstmPolicy</i>	alias of <code>RecurrentActorCriticCnnPolicy</code>
<i>MultiInputLstmPolicy</i>	alias of <code>RecurrentMultiInputActorCriticPolicy</code>

## 6.1 Notes

- Blog post: <https://iclr-blog-track.github.io/2022/03/25/ppo-implementation-details/>

## 6.2 Can I use?

- Recurrent policies: ✓
- Multi processing: ✓
- Gym spaces:

Space	Action	Observation
Discrete	✓	✓
Box	✓	✓
MultiDiscrete	✓	✓
MultiBinary	✓	✓
Dict		✓

## 6.3 Example

**Note:** It is particularly important to pass the `lstm_states` and `episode_start` argument to the `predict()` method, so the cell and hidden states of the LSTM are correctly updated.

---

```
import numpy as np

from sb3_contrib import RecurrentPPO
from stable_baselines3.common.evaluation import evaluate_policy

model = RecurrentPPO("MlpLstmPolicy", "CartPole-v1", verbose=1)
model.learn(5000)

env = model.get_env()
mean_reward, std_reward = evaluate_policy(model, env, n_eval_episodes=20, warn=False)
print(mean_reward)

model.save("ppo_recurrent")
del model # remove to demonstrate saving and loading

model = RecurrentPPO.load("ppo_recurrent")

obs = env.reset()
# cell and hidden state of the LSTM
lstm_states = None
num_envs = 1
# Episode start signals are used to reset the lstm states
episode_starts = np.ones((num_envs,), dtype=bool)
while True:
    action, lstm_states = model.predict(obs, state=lstm_states, episode_start=episode_
    ↪starts, deterministic=True)
    obs, rewards, dones, info = env.step(action)
    episode_starts = dones
    env.render()
```

## 6.4 Results

Report on environments with masked velocity (with and without framestack) can be found here: <https://wandb.ai/sb3/no-vel-envs/reports/PPO-vs-RecurrentPPO-aka-PPO-LSTM-on-environments-with-masked-velocity-VmldzoxOTI4NjE4>

RecurrentPPO was evaluated against PPO on:

- PendulumNoVel-v1
- LunarLanderNoVel-v2
- CartPoleNoVel-v1
- MountainCarContinuousNoVel-v0
- CarRacing-v0

### 6.4.1 How to replicate the results?

Clone the repo for the experiment:

```
git clone https://github.com/DLR-RM/rl-baselines3-zoo
cd rl-baselines3-zoo
git checkout feat/recurrent-ppo
```

Run the benchmark (replace \$ENV\_ID by the envs mentioned above):

```
python train.py --algo ppo_lstm --env $ENV_ID --eval-episodes 10 --eval-freq 10000
```

## 6.5 Parameters

```
class sb3_contrib.ppo_recurrent.RecurrentPPO(policy, env, learning_rate=0.0003, n_steps=128,
                                             batch_size=128, n_epochs=10, gamma=0.99,
                                             gae_lambda=0.95, clip_range=0.2, clip_range_vf=None,
                                             normalize_advantage=True, ent_coef=0.0, vf_coef=0.5,
                                             max_grad_norm=0.5, use_sde=False,
                                             sde_sample_freq=-1, target_kl=None,
                                             tensorboard_log=None, create_eval_env=False,
                                             policy_kwargs=None, verbose=0, seed=None,
                                             device='auto', _init_setup_model=True)
```

Proximal Policy Optimization algorithm (PPO) (clip version) with support for recurrent policies (LSTM).

Based on the original Stable Baselines 3 implementation.

Introduction to PPO: <https://spinningup.openai.com/en/latest/algorithms/ppo.html>

#### Parameters

- **policy** (Union[str, Type[RecurrentActorCriticPolicy]]) – The policy model to use (MlpPolicy, CnnPolicy, ...)
- **env** (Union[Env, VecEnv, str]) – The environment to learn from (if registered in Gym, can be str)
- **learning\_rate** (Union[float, Callable[[float], float]]) – The learning rate, it can be a function of the current progress remaining (from 1 to 0)
- **n\_steps** (int) – The number of steps to run for each environment per update (i.e. batch size is  $n\_steps * n\_env$  where  $n\_env$  is number of environment copies running in parallel)
- **batch\_size** (Optional[int]) – Minibatch size
- **n\_epochs** (int) – Number of epoch when optimizing the surrogate loss
- **gamma** (float) – Discount factor
- **gae\_lambda** (float) – Factor for trade-off of bias vs variance for Generalized Advantage Estimator
- **clip\_range** (Union[float, Callable[[float], float]]) – Clipping parameter, it can be a function of the current progress remaining (from 1 to 0).
- **clip\_range\_vf** (Union[None, float, Callable[[float], float]]) – Clipping parameter for the value function, it can be a function of the current progress remaining (from 1 to 0). This is a parameter specific to the OpenAI implementation. If None is passed (default), no

clipping will be done on the value function. IMPORTANT: this clipping depends on the reward scaling.

- **normalize\_advantage** (bool) – Whether to normalize or not the advantage
- **ent\_coef** (float) – Entropy coefficient for the loss calculation
- **vf\_coef** (float) – Value function coefficient for the loss calculation
- **max\_grad\_norm** (float) – The maximum value for the gradient clipping
- **target\_kl** (Optional[float]) – Limit the KL divergence between updates, because the clipping is not enough to prevent large update see issue #213 (cf <https://github.com/hill-a/stable-baselines/issues/213>) By default, there is no limit on the kl div.
- **tensorboard\_log** (Optional[str]) – the log location for tensorboard (if None, no logging)
- **create\_eval\_env** (bool) – Whether to create a second environment that will be used for evaluating the agent periodically (Only available when passing string for the environment). Caution, this parameter is deprecated and will be removed in the future.
- **policy\_kwargs** (Optional[Dict[str, Any]]) – additional arguments to be passed to the policy on creation
- **verbose** (int) – the verbosity level: 0 no output, 1 info, 2 debug
- **seed** (Optional[int]) – Seed for the pseudo random generators
- **device** (Union[device, str]) – Device (cpu, cuda, ...) on which the code should be run. Setting it to auto, the code will be run on the GPU if possible.
- **\_init\_setup\_model** (bool) – Whether or not to build the network at the creation of the instance

**collect\_rollouts**(*env, callback, rollout\_buffer, n\_rollout\_steps*)

Collect experiences using the current policy and fill a `RolloutBuffer`. The term rollout here refers to the model-free notion and should not be used with the concept of rollout used in model-based RL or planning.

#### Parameters

- **env** (VecEnv) – The training environment
- **callback** (BaseCallback) – Callback that will be called at each step (and at the beginning and end of the rollout)
- **rollout\_buffer** (RolloutBuffer) – Buffer to fill with rollouts
- **n\_steps** – Number of experiences to collect per environment

#### Return type

bool

#### Returns

True if function returned with at least *n\_rollout\_steps* collected, False if callback terminated rollout prematurely.

**get\_env**()

Returns the current environment (can be None if not defined).

#### Return type

Optional[VecEnv]

#### Returns

The current environment



**get\_parameters()**

Return the parameters of the agent. This includes parameters from different networks, e.g. critics (value functions) and policies (pi functions).

**Return type**

Dict[str, Dict]

**Returns**

Mapping of from names of the objects to PyTorch state-dicts.

**get\_vec\_normalize\_env()**

Return the VecNormalize wrapper of the training env if it exists.

**Return type**

Optional[VecNormalize]

**Returns**

The VecNormalize env.

**learn**(*total\_timesteps*, *callback=None*, *log\_interval=1*, *eval\_env=None*, *eval\_freq=-1*, *n\_eval\_episodes=5*, *tb\_log\_name='RecurrentPPO'*, *eval\_log\_path=None*, *reset\_num\_timesteps=True*, *progress\_bar=False*)

Return a trained model.

**Parameters**

- **total\_timesteps** (int) – The total number of samples (env steps) to train on
- **callback** (Union[None, Callable, List[BaseCallback], BaseCallback]) – call-back(s) called at every step with state of the algorithm.
- **log\_interval** (int) – The number of timesteps before logging.
- **tb\_log\_name** (str) – the name of the run for TensorBoard logging
- **reset\_num\_timesteps** (bool) – whether or not to reset the current timestep number (used in logging)
- **progress\_bar** (bool) – Display a progress bar using tqdm and rich.

**Return type**

TypeVar(RecurrentPPOSelf, bound= RecurrentPPO)

**Returns**

the trained model

**classmethod load**(*path*, *env=None*, *device='auto'*, *custom\_objects=None*, *print\_system\_info=False*, *force\_reset=True*, *\*\*kwargs*)

Load the model from a zip-file. Warning: load re-creates the model from scratch, it does not update it in-place! For an in-place load use `set_parameters` instead.

**Parameters**

- **path** (Union[str, Path, BufferedIOBase]) – path to the file (or a file-like) where to load the agent from
- **env** (Union[Env, VecEnv, None]) – the new environment to run the loaded model on (can be None if you only need prediction from a trained model) has priority over any saved environment
- **device** (Union[device, str]) – Device on which the code should run.

- **custom\_objects** (Optional[Dict[str, Any]]) – Dictionary of objects to replace upon loading. If a variable is present in this dictionary as a key, it will not be deserialized and the corresponding item will be used instead. Similar to custom\_objects in `keras.models.load_model`. Useful when you have an object in file that can not be deserialized.
- **print\_system\_info** (bool) – Whether to print system info from the saved model and the current system info (useful to debug loading issues)
- **force\_reset** (bool) – Force call to `reset()` before training to avoid unexpected behavior. See <https://github.com/DLR-RM/stable-baselines3/issues/597>
- **kwargs** – extra arguments to change the model when loading

**Return type**

TypeVar(SelfBaseAlgorithm, bound= BaseAlgorithm)

**Returns**

new model instance with loaded parameters

**property logger: Logger**

Getter for the logger object.

**predict**(*observation, state=None, episode\_start=None, deterministic=False*)

Get the policy action from an observation (and optional hidden state). Includes sugar-coating to handle different observations (e.g. normalizing images).

**Parameters**

- **observation** (Union[ndarray, Dict[str, ndarray]]) – the input observation
- **state** (Optional[Tuple[ndarray, ...]]) – The last hidden states (can be None, used in recurrent policies)
- **episode\_start** (Optional[ndarray]) – The last masks (can be None, used in recurrent policies) this correspond to beginning of episodes, where the hidden states of the RNN must be reset.
- **deterministic** (bool) – Whether or not to return deterministic actions.

**Return type**

Tuple[ndarray, Optional[Tuple[ndarray, ...]]]

**Returns**

the model's action and the next hidden state (used in recurrent policies)

**save**(*path, exclude=None, include=None*)

Save all the attributes of the object and the model parameters in a zip-file.

**Parameters**

- **path** (Union[str, Path, BufferedIOBase]) – path to the file where the rl agent should be saved
- **exclude** (Optional[Iterable[str]]) – name of parameters that should be excluded in addition to the default ones
- **include** (Optional[Iterable[str]]) – name of parameters that might be excluded but should be included anyway

**Return type**

None

**set\_env**(*env*, *force\_reset=True*)

Checks the validity of the environment, and if it is coherent, set it as the current environment. Furthermore wrap any non vectorized env into a vectorized checked parameters: - observation\_space - action\_space

**Parameters**

- **env** (Union[Env, VecEnv]) – The environment for learning a policy
- **force\_reset** (bool) – Force call to `reset()` before training to avoid unexpected behavior. See issue <https://github.com/DLR-RM/stable-baselines3/issues/597>

**Return type**

None

**set\_logger**(*logger*)

Setter for for logger object. :rtype: None

**Warning:** When passing a custom logger object, this will overwrite `tensorboard_log` and `verbose` settings passed to the constructor.

**set\_parameters**(*load\_path\_or\_dict*, *exact\_match=True*, *device='auto'*)

Load parameters from a given zip-file or a nested dictionary containing parameters for different modules (see `get_parameters`).

**Parameters**

- **load\_path\_or\_iter** – Location of the saved data (path or file-like, see `save`), or a nested dictionary containing `nn.Module` parameters used by the policy. The dictionary maps object names to a state-dictionary returned by `torch.nn.Module.state_dict()`.
- **exact\_match** (bool) – If True, the given parameters should include parameters for each module and each of their parameters, otherwise raises an Exception. If set to False, this can be used to update only specific parameters.
- **device** (Union[device, str]) – Device on which the code should run.

**Return type**

None

**set\_random\_seed**(*seed=None*)

Set the seed of the pseudo-random generators (python, numpy, pytorch, gym, action\_space)

**Parameters**

**seed** (Optional[int]) –

**Return type**

None

**train**()

Update policy using the currently gathered rollout buffer.

**Return type**

None

## 6.6 RecurrentPPO Policies

`sb3_contrib.ppo_recurrent.MlpLstmPolicy`

alias of `RecurrentActorCriticPolicy`

```
class sb3_contrib.common.recurrent.policies.RecurrentActorCriticPolicy(observation_space,
    action_space,
    lr_schedule,
    net_arch=None,
    activation_fn=<class
        'torch.nn.modules.activation.Tanh'>,
    ortho_init=True,
    use_sde=False,
    log_std_init=0.0,
    full_std=True,
    sde_net_arch=None,
    use_expln=False,
    squash_output=False,
    features_extractor_class=<class
        'stable_baselines3.common.torch_layers.FlattenExtractor'>,
    features_extractor_kwargs=None,
    normalize_images=True,
    optimizer_class=<class
        'torch.optim.adam.Adam'>,
    optimizer_kwargs=None,
    lstm_hidden_size=256,
    n_lstm_layers=1,
    shared_lstm=False,
    enable_critic_lstm=True,
    lstm_kwargs=None)
```

Recurrent policy class for actor-critic algorithms (has both policy and value prediction). To be used with A2C, PPO and the likes. It assumes that both the actor and the critic LSTM have the same architecture.

### Parameters

- **observation\_space** (Space) – Observation space
- **action\_space** (Space) – Action space
- **lr\_schedule** (Callable[[float], float]) – Learning rate schedule (could be constant)
- **net\_arch** (Optional[List[Union[int, Dict[str, List[int]]]]) – The specification of the policy and value networks.
- **activation\_fn** (Type[Module]) – Activation function
- **ortho\_init** (bool) – Whether to use or not orthogonal initialization
- **use\_sde** (bool) – Whether to use State Dependent Exploration or not
- **log\_std\_init** (float) – Initial value for the log standard deviation

- **full\_std** (bool) – Whether to use (n\_features x n\_actions) parameters for the std instead of only (n\_features,) when using gSDE
- **sde\_net\_arch** (Optional[List[int]]) – Network architecture for extracting features when using gSDE. If None, the latent features from the policy will be used. Pass an empty list to use the states as features.
- **use\_expln** (bool) – Use `expln()` function instead of `exp()` to ensure a positive standard deviation (cf paper). It allows to keep variance above zero and prevent it from growing too fast. In practice, `exp()` is usually enough.
- **squash\_output** (bool) – Whether to squash the output using a tanh function, this allows to ensure boundaries when using gSDE.
- **features\_extractor\_class** (Type[BaseFeaturesExtractor]) – Features extractor to use.
- **features\_extractor\_kwargs** (Optional[Dict[str, Any]]) – Keyword arguments to pass to the features extractor.
- **normalize\_images** (bool) – Whether to normalize images or not, dividing by 255.0 (True by default)
- **optimizer\_class** (Type[Optimizer]) – The optimizer to use, `th.optim.Adam` by default
- **optimizer\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments, excluding the learning rate, to pass to the optimizer
- **lstm\_hidden\_size** (int) – Number of hidden units for each LSTM layer.
- **n\_lstm\_layers** (int) – Number of LSTM layers.
- **shared\_lstm** (bool) – Whether the LSTM is shared between the actor and the critic (in that case, only the actor gradient is used) By default, the actor and the critic have two separate LSTM.
- **enable\_critic\_lstm** (bool) – Use a separate LSTM for the critic.
- **lstm\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments to pass the the LSTM constructor.

**evaluate\_actions**(*obs, actions, lstm\_states, episode\_starts*)

Evaluate actions according to the current policy, given the observations.

#### Parameters

- **obs** (Tensor) – Observation.
- **actions** (Tensor) –
- **lstm\_states** (RNNStates) – The last hidden and memory states for the LSTM.
- **episode\_starts** (Tensor) – Whether the observations correspond to new episodes or not (we reset the lstm states in that case).

#### Return type

Tuple[Tensor, Tensor, Tensor]

#### Returns

estimated value, log likelihood of taking those actions and entropy of the action distribution.

**forward**(*obs, lstm\_states, episode\_starts, deterministic=False*)

Forward pass in all the networks (actor and critic)

#### Parameters

- **obs** (Tensor) – Observation. Observation
- **lstm\_states** (RNNStates) – The last hidden and memory states for the LSTM.
- **episode\_starts** (Tensor) – Whether the observations correspond to new episodes or not (we reset the lstm states in that case).
- **deterministic** (bool) – Whether to sample or use deterministic actions

**Return type**

Tuple[Tensor, Tensor, Tensor, RNNStates]

**Returns**

action, value and log probability of the action

**get\_distribution**(*obs*, *lstm\_states*, *episode\_starts*)

Get the current policy distribution given the observations.

**Parameters**

- **obs** (Tensor) – Observation.
- **lstm\_states** (Tuple[Tensor, Tensor]) – The last hidden and memory states for the LSTM.
- **episode\_starts** (Tensor) – Whether the observations correspond to new episodes or not (we reset the lstm states in that case).

**Return type**

Tuple[Distribution, Tuple[Tensor, ...]]

**Returns**

the action distribution and new hidden states.

**predict**(*observation*, *state=None*, *episode\_start=None*, *deterministic=False*)

Get the policy action from an observation (and optional hidden state). Includes sugar-coating to handle different observations (e.g. normalizing images).

**Parameters**

- **observation** (Union[ndarray, Dict[str, ndarray]]) – the input observation
- **lstm\_states** – The last hidden and memory states for the LSTM.
- **episode\_starts** – Whether the observations correspond to new episodes or not (we reset the lstm states in that case).
- **deterministic** (bool) – Whether or not to return deterministic actions.

**Return type**

Tuple[ndarray, Optional[Tuple[ndarray, ...]]]

**Returns**

the model's action and the next hidden state (used in recurrent policies)

**predict\_values**(*obs*, *lstm\_states*, *episode\_starts*)

Get the estimated values according to the current policy given the observations.

**Parameters**

- **obs** (Tensor) – Observation.
- **lstm\_states** (Tuple[Tensor, Tensor]) – The last hidden and memory states for the LSTM.

- **episode\_starts** (Tensor) – Whether the observations correspond to new episodes or not (we reset the lstm states in that case).

**Return type**

Tensor

**Returns**

the estimated values.

`sb3_contrib.ppo_recurrent.CnnLstmPolicy`alias of `RecurrentActorCriticCnnPolicy`

```
class sb3_contrib.common.recurrent.policies.RecurrentActorCriticCnnPolicy(observation_space,
    action_space,
    lr_schedule,
    net_arch=None, activation_fn=<class
    'torch.nn.modules.activation.Tanh'>,
    ortho_init=True,
    use_sde=False,
    log_std_init=0.0,
    full_std=True,
    sde_net_arch=None,
    use_expln=False,
    squash_output=False,
    features_extractor_class=<class
    'stable_baselines3.common.torch_layers.LstmWrapper'>,
    features_extractor_kwargs=None,
    normalize_images=True,
    optimizer_class=<class
    'torch.optim.adam.Adam'>,
    optimizer_kwargs=None,
    lstm_hidden_size=256,
    n_lstm_layers=1,
    shared_lstm=False,
    enable_critic_lstm=True,
    lstm_kwargs=None)
```

CNN recurrent policy class for actor-critic algorithms (has both policy and value prediction). Used by A2C, PPO and the likes.

**Parameters**

- **observation\_space** (Space) – Observation space
- **action\_space** (Space) – Action space
- **lr\_schedule** (Callable[[float], float]) – Learning rate schedule (could be constant)
- **net\_arch** (Optional[List[Union[int, Dict[str, List[int]]]]) – The specification of the policy and value networks.

- **activation\_fn** (Type[Module]) – Activation function
- **ortho\_init** (bool) – Whether to use or not orthogonal initialization
- **use\_sde** (bool) – Whether to use State Dependent Exploration or not
- **log\_std\_init** (float) – Initial value for the log standard deviation
- **full\_std** (bool) – Whether to use (n\_features x n\_actions) parameters for the std instead of only (n\_features,) when using gSDE
- **sde\_net\_arch** (Optional[List[int]]) – Network architecture for extracting features when using gSDE. If None, the latent features from the policy will be used. Pass an empty list to use the states as features.
- **use\_expln** (bool) – Use `expln()` function instead of `exp()` to ensure a positive standard deviation (cf paper). It allows to keep variance above zero and prevent it from growing too fast. In practice, `exp()` is usually enough.
- **squash\_output** (bool) – Whether to squash the output using a tanh function, this allows to ensure boundaries when using gSDE.
- **features\_extractor\_class** (Type[BaseFeaturesExtractor]) – Features extractor to use.
- **features\_extractor\_kwargs** (Optional[Dict[str, Any]]) – Keyword arguments to pass to the features extractor.
- **normalize\_images** (bool) – Whether to normalize images or not, dividing by 255.0 (True by default)
- **optimizer\_class** (Type[Optimizer]) – The optimizer to use, `th.optim.Adam` by default
- **optimizer\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments, excluding the learning rate, to pass to the optimizer
- **lstm\_hidden\_size** (int) – Number of hidden units for each LSTM layer.
- **n\_lstm\_layers** (int) – Number of LSTM layers.
- **shared\_lstm** (bool) – Whether the LSTM is shared between the actor and the critic. By default, only the actor has a recurrent network.
- **enable\_critic\_lstm** (bool) – Use a separate LSTM for the critic.
- **lstm\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments to pass the the LSTM constructor.

`sb3_contrib.ppo_recurrent.MultiInputLstmPolicy`

alias of `RecurrentMultiInputActorCriticPolicy`



```

class sb3_contrib.common.recurrent.policies.RecurrentMultiInputActorCriticPolicy(observation_space,
                                                                                  ac-
                                                                                  tion_space,
                                                                                  lr_schedule,
                                                                                  net_arch=None,
                                                                                  activa-
                                                                                  tion_fn=<class
                                                                                  'torch.nn.modules.activation.
                                                                                  or-
                                                                                  tho_init=True,
                                                                                  use_sde=False,
                                                                                  log_std_init=0.0,
                                                                                  full_std=True,
                                                                                  sde_net_arch=None,
                                                                                  use_expln=False,
                                                                                  squash_output=False,
                                                                                  fea-
                                                                                  tures_extractor_class=<class
                                                                                  'sta-
                                                                                  ble_baselines3.common.torch
                                                                                  fea-
                                                                                  tures_extractor_kwargs=None,
                                                                                  normal-
                                                                                  ize_images=True,
                                                                                  opti-
                                                                                  mizer_class=<class
                                                                                  'torch.optim.adam.Adam'>,
                                                                                  opti-
                                                                                  mizer_kwargs=None,
                                                                                  lstm_hidden_size=256,
                                                                                  n_lstm_layers=1,
                                                                                  shared_lstm=False,
                                                                                  en-
                                                                                  able_critic_lstm=True,
                                                                                  lstm_kwargs=None)

```

MultiInputActorClass policy class for actor-critic algorithms (has both policy and value prediction). Used by A2C, PPO and the likes.

#### Parameters

- **observation\_space** (Space) – Observation space
- **action\_space** (Space) – Action space
- **lr\_schedule** (Callable[[float], float]) – Learning rate schedule (could be constant)
- **net\_arch** (Optional[List[Union[int, Dict[str, List[int]]]]) – The specification of the policy and value networks.
- **activation\_fn** (Type[Module]) – Activation function
- **ortho\_init** (bool) – Whether to use or not orthogonal initialization
- **use\_sde** (bool) – Whether to use State Dependent Exploration or not
- **log\_std\_init** (float) – Initial value for the log standard deviation
- **full\_std** (bool) – Whether to use (n\_features x n\_actions) parameters for the std instead of only (n\_features,) when using gSDE

- **sde\_net\_arch** (Optional[List[int]]) – Network architecture for extracting features when using gSDE. If None, the latent features from the policy will be used. Pass an empty list to use the states as features.
- **use\_expln** (bool) – Use `expln()` function instead of `exp()` to ensure a positive standard deviation (cf paper). It allows to keep variance above zero and prevent it from growing too fast. In practice, `exp()` is usually enough.
- **squash\_output** (bool) – Whether to squash the output using a tanh function, this allows to ensure boundaries when using gSDE.
- **features\_extractor\_class** (Type[BaseFeaturesExtractor]) – Features extractor to use.
- **features\_extractor\_kwargs** (Optional[Dict[str, Any]]) – Keyword arguments to pass to the features extractor.
- **normalize\_images** (bool) – Whether to normalize images or not, dividing by 255.0 (True by default)
- **optimizer\_class** (Type[Optimizer]) – The optimizer to use, `th.optim.Adam` by default
- **optimizer\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments, excluding the learning rate, to pass to the optimizer
- **lstm\_hidden\_size** (int) – Number of hidden units for each LSTM layer.
- **n\_lstm\_layers** (int) – Number of LSTM layers.
- **shared\_lstm** (bool) – Whether the LSTM is shared between the actor and the critic. By default, only the actor has a recurrent network.
- **enable\_critic\_lstm** (bool) – Use a separate LSTM for the critic.
- **lstm\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments to pass the the LSTM constructor.

## QR-DQN

Quantile Regression DQN (QR-DQN) builds on Deep Q-Network (DQN) and make use of quantile regression to explicitly model the distribution over returns, instead of predicting the mean return (DQN).

### Available Policies

<i>MlpPolicy</i>	alias of QRDQNPolicy
<i>CnnPolicy</i>	Policy class for QR-DQN when using images as input.
<i>MultiInputPolicy</i>	Policy class for QR-DQN when using dict observations as input.

## 7.1 Notes

- Original paper: <https://arxiv.org/abs/1710.100442>
- Distributional RL (C51): <https://arxiv.org/abs/1707.06887>
- Further reference: [https://github.com/amy12xx/ml\\_notes\\_and\\_reports/blob/master/distributional\\_rl/QRDQN.pdf](https://github.com/amy12xx/ml_notes_and_reports/blob/master/distributional_rl/QRDQN.pdf)

## 7.2 Can I use?

- Recurrent policies:
- Multi processing: ✓
- Gym spaces:

Space	Action	Observation
Discrete	✓	✓
Box		✓
MultiDiscrete		✓
MultiBinary		✓
Dict		✓

## 7.3 Example

```
import gym

from sb3_contrib import QRDQN

env = gym.make("CartPole-v1")

policy_kwargs = dict(n_quantiles=50)
model = QRDQN("MlpPolicy", env, policy_kwargs=policy_kwargs, verbose=1)
model.learn(total_timesteps=10000, log_interval=4)
model.save("qrdqn_cartpole")

del model # remove to demonstrate saving and loading

model = QRDQN.load("qrdqn_cartpole")

obs = env.reset()
while True:
    action, _states = model.predict(obs, deterministic=True)
    obs, reward, done, info = env.step(action)
    env.render()
    if done:
        obs = env.reset()
```

## 7.4 Results

Result on Atari environments (10M steps, Pong and Breakout) and classic control tasks using 3 and 5 seeds.

The complete learning curves are available in the [associated PR](#).

---

**Note:** QR-DQN implementation was validated against [Intel Coach](#) one which roughly compare to the original paper results (we trained the agent with a smaller budget).

---

Environments	QR-DQN	DQN
Breakout	413 +/- 21	~300
Pong	20 +/- 0	~20
CartPole	386 +/- 64	500 +/- 0
MountainCar	-111 +/- 4	-107 +/- 4
LunarLander	168 +/- 39	195 +/- 28
Acrobot	-73 +/- 2	-74 +/- 2

### 7.4.1 How to replicate the results?

Clone RL-Zoo fork and checkout the branch `feat/qrdqn`:

```
git clone https://github.com/ku2482/rl-baselines3-zoo/
cd rl-baselines3-zoo/
git checkout feat/qrdqn
```

Run the benchmark (replace `$ENV_ID` by the envs mentioned above):

```
python train.py --algo qrdqn --env $ENV_ID --eval-episodes 10 --eval-freq 10000
```

Plot the results:

```
python scripts/all_plots.py -a qrdqn -e Breakout Pong -f logs/ -o logs/qrdqn_results
python scripts/plot_from_file.py -i logs/qrdqn_results.pkl -latex -l QR-DQN
```

## 7.5 Parameters

```
class sb3_contrib.qrdqn.QRDQN(policy, env, learning_rate=5e-05, buffer_size=1000000,
                               learning_starts=50000, batch_size=32, tau=1.0, gamma=0.99, train_freq=4,
                               gradient_steps=1, replay_buffer_class=None, replay_buffer_kwargs=None,
                               optimize_memory_usage=False, target_update_interval=10000,
                               exploration_fraction=0.005, exploration_initial_eps=1.0,
                               exploration_final_eps=0.01, max_grad_norm=None, tensorboard_log=None,
                               create_eval_env=False, policy_kwargs=None, verbose=0, seed=None,
                               device='auto', _init_setup_model=True)
```

Quantile Regression Deep Q-Network (QR-DQN) Paper: <https://arxiv.org/abs/1710.10044> Default hyperparameters are taken from the paper and are tuned for Atari games.

#### Parameters

- **policy** (Union[str, Type[QRDQNPolicy]]) – The policy model to use (MlpPolicy, CnnPolicy, ...)
- **env** (Union[Env, VecEnv, str]) – The environment to learn from (if registered in Gym, can be str)
- **learning\_rate** (Union[float, Callable[[float], float]]) – The learning rate, it can be a function of the current progress remaining (from 1 to 0)
- **buffer\_size** (int) – size of the replay buffer
- **learning\_starts** (int) – how many steps of the model to collect transitions for before learning starts
- **batch\_size** (Optional[int]) – Minibatch size for each gradient update
- **tau** (float) – the soft update coefficient (“Polyak update”, between 0 and 1) default 1 for hard update
- **gamma** (float) – the discount factor
- **train\_freq** (int) – Update the model every `train_freq` steps. Alternatively pass a tuple of frequency and unit like (5, "step") or (2, "episode").

- **gradient\_steps** (int) – How many gradient steps to do after each rollout (see `train_freq` and `n_episodes_rollout`) Set to -1 means to do as many gradient steps as steps done in the environment during the rollout.
- **replay\_buffer\_class** (Optional[ReplayBuffer]) – Replay buffer class to use (for instance `HerReplayBuffer`). If None, it will be automatically selected.
- **replay\_buffer\_kwargs** (Optional[Dict[str, Any]]) – Keyword arguments to pass to the replay buffer on creation.
- **optimize\_memory\_usage** (bool) – Enable a memory efficient variant of the replay buffer at a cost of more complexity. See <https://github.com/DLR-RM/stable-baselines3/issues/37#issuecomment-637501195>
- **target\_update\_interval** (int) – update the target network every `target_update_interval` environment steps.
- **exploration\_fraction** (float) – fraction of entire training period over which the exploration rate is reduced
- **exploration\_initial\_eps** (float) – initial value of random action probability
- **exploration\_final\_eps** (float) – final value of random action probability
- **max\_grad\_norm** (Optional[float]) – The maximum value for the gradient clipping (if None, no clipping)
- **tensorboard\_log** (Optional[str]) – the log location for tensorboard (if None, no logging)
- **create\_eval\_env** (bool) – Whether to create a second environment that will be used for evaluating the agent periodically (Only available when passing string for the environment). Caution, this parameter is deprecated and will be removed in the future.
- **policy\_kwargs** (Optional[Dict[str, Any]]) – additional arguments to be passed to the policy on creation
- **verbose** (int) – the verbosity level: 0 no output, 1 info, 2 debug
- **seed** (Optional[int]) – Seed for the pseudo random generators
- **device** (Union[device, str]) – Device (cpu, cuda, ...) on which the code should be run. Setting it to auto, the code will be run on the GPU if possible.
- **\_init\_setup\_model** (bool) – Whether or not to build the network at the creation of the instance

**collect\_rollouts**(*env, callback, train\_freq, replay\_buffer, action\_noise=None, learning\_starts=0, log\_interval=None*)

Collect experiences and store them into a `ReplayBuffer`.

#### Parameters

- **env** (VecEnv) – The training environment
- **callback** (BaseCallback) – Callback that will be called at each step (and at the beginning and end of the rollout)
- **train\_freq** (TrainFreq) – How much experience to collect by doing rollouts of current policy. Either `TrainFreq(<n>, TrainFrequencyUnit.STEP)` or `TrainFreq(<n>, TrainFrequencyUnit.EPISODE)` with `<n>` being an integer greater than 0.

- **action\_noise** (Optional[ActionNoise]) – Action noise that will be used for exploration Required for deterministic policy (e.g. TD3). This can also be used in addition to the stochastic policy for SAC.
- **learning\_starts** (int) – Number of steps before learning for the warm-up phase.
- **replay\_buffer** (ReplayBuffer) –
- **log\_interval** (Optional[int]) – Log data every log\_interval episodes

**Return type**

RolloutReturn

**Returns****get\_env()**

Returns the current environment (can be None if not defined).

**Return type**

Optional[VecEnv]

**Returns**

The current environment

**get\_parameters()**

Return the parameters of the agent. This includes parameters from different networks, e.g. critics (value functions) and policies (pi functions).

**Return type**

Dict[str, Dict]

**Returns**

Mapping of from names of the objects to PyTorch state-dicts.

**get\_vec\_normalize\_env()**

Return the VecNormalize wrapper of the training env if it exists.

**Return type**

Optional[VecNormalize]

**Returns**

The VecNormalize env.

**learn**(total\_timesteps, callback=None, log\_interval=4, eval\_env=None, eval\_freq=-1, n\_eval\_episodes=5, tb\_log\_name='QRDQN', eval\_log\_path=None, reset\_num\_timesteps=True, progress\_bar=False)

Return a trained model.

**Parameters**

- **total\_timesteps** (int) – The total number of samples (env steps) to train on
- **callback** (Union[None, Callable, List[BaseCallback], BaseCallback]) – call-back(s) called at every step with state of the algorithm.
- **log\_interval** (int) – The number of timesteps before logging.
- **tb\_log\_name** (str) – the name of the run for TensorBoard logging
- **reset\_num\_timesteps** (bool) – whether or not to reset the current timestep number (used in logging)
- **progress\_bar** (bool) – Display a progress bar using tqdm and rich.

**Return type**

TypeVar(QRDQNSelf, bound= QRDQN)

**Returns**

the trained model

**classmethod** `load(path, env=None, device='auto', custom_objects=None, print_system_info=False, force_reset=True, **kwargs)`

Load the model from a zip-file. Warning: `load` re-creates the model from scratch, it does not update it in-place! For an in-place load use `set_parameters` instead.

**Parameters**

- **path** (Union[str, Path, BufferedIOBase]) – path to the file (or a file-like) where to load the agent from
- **env** (Union[Env, VecEnv, None]) – the new environment to run the loaded model on (can be None if you only need prediction from a trained model) has priority over any saved environment
- **device** (Union[device, str]) – Device on which the code should run.
- **custom\_objects** (Optional[Dict[str, Any]]) – Dictionary of objects to replace upon loading. If a variable is present in this dictionary as a key, it will not be deserialized and the corresponding item will be used instead. Similar to `custom_objects` in `keras.models.load_model`. Useful when you have an object in file that can not be deserialized.
- **print\_system\_info** (bool) – Whether to print system info from the saved model and the current system info (useful to debug loading issues)
- **force\_reset** (bool) – Force call to `reset()` before training to avoid unexpected behavior. See <https://github.com/DLR-RM/stable-baselines3/issues/597>
- **kwargs** – extra arguments to change the model when loading

**Return type**

TypeVar(SelfBaseAlgorithm, bound= BaseAlgorithm)

**Returns**

new model instance with loaded parameters

**load\_replay\_buffer**(*path, truncate\_last\_traj=True*)

Load a replay buffer from a pickle file.

**Parameters**

- **path** (Union[str, Path, BufferedIOBase]) – Path to the pickled replay buffer.
- **truncate\_last\_traj** (bool) – When using `HerReplayBuffer` with online sampling: If set to `True`, we assume that the last trajectory in the replay buffer was finished (and truncate it). If set to `False`, we assume that we continue the same trajectory (same episode).

**Return type**

None

**property logger:** `Logger`

Getter for the logger object.

**predict**(*observation, state=None, episode\_start=None, deterministic=False*)

Get the policy action from an observation (and optional hidden state). Includes sugar-coating to handle different observations (e.g. normalizing images).

**Parameters**



- **observation** (ndarray) – the input observation
- **state** (Optional[Tuple[ndarray, ...]]) – The last hidden states (can be None, used in recurrent policies)
- **episode\_start** (Optional[ndarray]) – The last masks (can be None, used in recurrent policies) this correspond to beginning of episodes, where the hidden states of the RNN must be reset.
- **deterministic** (bool) – Whether or not to return deterministic actions.

**Return type**

Tuple[ndarray, Optional[Tuple[ndarray, ...]]]

**Returns**

the model's action and the next hidden state (used in recurrent policies)

**save**(*path*, *exclude=None*, *include=None*)

Save all the attributes of the object and the model parameters in a zip-file.

**Parameters**

- **path** (Union[str, Path, BufferedIOBase]) – path to the file where the rl agent should be saved
- **exclude** (Optional[Iterable[str]]) – name of parameters that should be excluded in addition to the default ones
- **include** (Optional[Iterable[str]]) – name of parameters that might be excluded but should be included anyway

**Return type**

None

**save\_replay\_buffer**(*path*)

Save the replay buffer as a pickle file.

**Parameters**

**path** (Union[str, Path, BufferedIOBase]) – Path to the file where the replay buffer should be saved. if path is a str or pathlib.Path, the path is automatically created if necessary.

**Return type**

None

**set\_env**(*env*, *force\_reset=True*)

Checks the validity of the environment, and if it is coherent, set it as the current environment. Furthermore wrap any non vectorized env into a vectorized checked parameters: - observation\_space - action\_space

**Parameters**

- **env** (Union[Env, VecEnv]) – The environment for learning a policy
- **force\_reset** (bool) – Force call to `reset()` before training to avoid unexpected behavior. See issue <https://github.com/DLR-RM/stable-baselines3/issues/597>

**Return type**

None

**set\_logger**(*logger*)

Setter for for logger object. :rtype: None

**Warning:** When passing a custom logger object, this will overwrite `tensorboard_log` and `verbose` settings passed to the constructor.

**set\_parameters**(*load\_path\_or\_dict*, *exact\_match=True*, *device='auto'*)

Load parameters from a given zip-file or a nested dictionary containing parameters for different modules (see `get_parameters`).

**Parameters**

- **load\_path\_or\_iter** – Location of the saved data (path or file-like, see `save`), or a nested dictionary containing `nn.Module` parameters used by the policy. The dictionary maps object names to a state-dictionary returned by `torch.nn.Module.state_dict()`.
- **exact\_match** (bool) – If `True`, the given parameters should include parameters for each module and each of their parameters, otherwise raises an `Exception`. If set to `False`, this can be used to update only specific parameters.
- **device** (Union[device, str]) – Device on which the code should run.

**Return type**

None

**set\_random\_seed**(*seed=None*)

Set the seed of the pseudo-random generators (python, numpy, pytorch, gym, action\_space)

**Parameters**

**seed** (Optional[int]) –

**Return type**

None

**train**(*gradient\_steps*, *batch\_size=100*)

Sample the replay buffer and do the updates (gradient descent and update target networks)

**Return type**

None

## 7.6 QR-DQN Policies

`sb3_contrib.qrdqn.MlpPolicy`

alias of `QRDQNPoly`

```
class sb3_contrib.qrdqn.policies.QRDQNPoly(observation_space, action_space, lr_schedule,
                                           n_quantiles=200, net_arch=None, activation_fn=<class
                                           'torch.nn.modules.activation.ReLU'>,
                                           features_extractor_class=<class 'stable_baselines3.common.torch_layers.FlattenExtractor'>,
                                           features_extractor_kwargs=None,
                                           normalize_images=True, optimizer_class=<class
                                           'torch.optim.adam.Adam'>, optimizer_kwargs=None)
```

Policy class with quantile and target networks for QR-DQN.

**Parameters**

- **observation\_space** (Space) – Observation space
- **action\_space** (Space) – Action space

- **lr\_schedule** (Callable[[float], float]) – Learning rate schedule (could be constant)
- **n\_quantiles** (int) – Number of quantiles
- **net\_arch** (Optional[List[int]]) – The specification of the network architecture.
- **activation\_fn** (Type[Module]) – Activation function
- **features\_extractor\_class** (Type[BaseFeaturesExtractor]) – Features extractor to use.
- **features\_extractor\_kwargs** (Optional[Dict[str, Any]]) – Keyword arguments to pass to the features extractor.
- **normalize\_images** (bool) – Whether to normalize images or not, dividing by 255.0 (True by default)
- **optimizer\_class** (Type[Optimizer]) – The optimizer to use, `th.optim.Adam` by default
- **optimizer\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments, excluding the learning rate, to pass to the optimizer

**forward**(*obs*, *deterministic=True*)

Defines the computation performed at every call.

Should be overridden by all subclasses. :rtype: Tensor

---

**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.

---

**set\_training\_mode**(*mode*)

Put the policy in either training or evaluation mode. This affects certain modules, such as batch normalisation and dropout. :type mode: bool :param mode: if true, set to training mode, else set to evaluation mode

**Return type**

None

```
class sb3_contrib.qrdqn.CnnPolicy(observation_space, action_space, lr_schedule, n_quantiles=200,
                                net_arch=None, activation_fn=<class
                                'torch.nn.modules.activation.ReLU'>, features_extractor_class=<class
                                'stable_baselines3.common.torch_layers.NatureCNN'>,
                                features_extractor_kwargs=None, normalize_images=True,
                                optimizer_class=<class 'torch.optim.adam.Adam'>,
                                optimizer_kwargs=None)
```

Policy class for QR-DQN when using images as input.

**Parameters**

- **observation\_space** (Space) – Observation space
- **action\_space** (Space) – Action space
- **lr\_schedule** (Callable[[float], float]) – Learning rate schedule (could be constant)
- **n\_quantiles** (int) – Number of quantiles
- **net\_arch** (Optional[List[int]]) – The specification of the network architecture.
- **activation\_fn** (Type[Module]) – Activation function

- **features\_extractor\_class** (Type[BaseFeaturesExtractor]) – Features extractor to use.
- **normalize\_images** (bool) – Whether to normalize images or not, dividing by 255.0 (True by default)
- **optimizer\_class** (Type[Optimizer]) – The optimizer to use, `th.optim.Adam` by default
- **optimizer\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments, excluding the learning rate, to pass to the optimizer

```
class sb3_contrib.qrdqn.MultiInputPolicy(observation_space, action_space, lr_schedule,
                                         n_quantiles=200, net_arch=None, activation_fn=<class
                                         'torch.nn.modules.activation.ReLU'>,
                                         features_extractor_class=<class 'stable_baselines3.common.torch_layers.CombinedExtractor'>,
                                         features_extractor_kwargs=None, normalize_images=True,
                                         optimizer_class=<class 'torch.optim.adam.Adam'>,
                                         optimizer_kwargs=None)
```

Policy class for QR-DQN when using dict observations as input.

#### Parameters

- **observation\_space** (Space) – Observation space
- **action\_space** (Space) – Action space
- **lr\_schedule** (Callable[[float], float]) – Learning rate schedule (could be constant)
- **n\_quantiles** (int) – Number of quantiles
- **net\_arch** (Optional[List[int]]) – The specification of the network architecture.
- **activation\_fn** (Type[Module]) – Activation function
- **features\_extractor\_class** (Type[BaseFeaturesExtractor]) – Features extractor to use.
- **normalize\_images** (bool) – Whether to normalize images or not, dividing by 255.0 (True by default)
- **optimizer\_class** (Type[Optimizer]) – The optimizer to use, `th.optim.Adam` by default
- **optimizer\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments, excluding the learning rate, to pass to the optimizer

## TQC

Controlling Overestimation Bias with Truncated Mixture of Continuous Distributional Quantile Critics (TQC). Truncated Quantile Critics (TQC) builds on SAC, TD3 and QR-DQN, making use of quantile regression to predict a distribution for the value function (instead of a mean value). It truncates the quantiles predicted by different networks (a bit as it is done in TD3).

### Available Policies

<i>MlpPolicy</i>	alias of TQCPolicy
<i>CnnPolicy</i>	Policy class (with both actor and critic) for TQC.
MultiInputPolicy	Policy class (with both actor and critic) for TQC.

## 8.1 Notes

- Original paper: <https://arxiv.org/abs/2005.04269>
- Original Implementation: [https://github.com/bayesgroup/tqc\\_pytorch](https://github.com/bayesgroup/tqc_pytorch)

## 8.2 Can I use?

- Recurrent policies:
- Multi processing: ✓
- Gym spaces:

Space	Action	Observation
Discrete		✓
Box	✓	✓
MultiDiscrete		✓
MultiBinary		✓
Dict		✓

## 8.3 Example

```
import gym
import numpy as np

from sb3_contrib import TQC

env = gym.make("Pendulum-v1")

policy_kwargs = dict(n_critics=2, n_quantiles=25)
model = TQC("MlpPolicy", env, top_quantiles_to_drop_per_net=2, verbose=1, policy_
↳kwargs=policy_kwargs)
model.learn(total_timesteps=10000, log_interval=4)
model.save("tqc_pendulum")

del model # remove to demonstrate saving and loading

model = TQC.load("tqc_pendulum")

obs = env.reset()
while True:
    action, _states = model.predict(obs, deterministic=True)
    obs, reward, done, info = env.step(action)
    env.render()
    if done:
        obs = env.reset()
```

## 8.4 Results

Result on the PyBullet benchmark (1M steps) and on BipedalWalkerHardcore-v3 (2M steps) using 3 seeds. The complete learning curves are available in the [associated PR](#).

The main difference with SAC is on harder environments (BipedalWalkerHardcore, Walker2D).

---

**Note:** Hyperparameters from the [gSDE paper](#) were used (as they are tuned for SAC on PyBullet envs), including using gSDE for the exploration and not the unstructured Gaussian noise but this should not affect results in simulation.

---

---

**Note:** We are using the open source PyBullet environments and not the MuJoCo simulator (as done in the original paper). You can find a complete benchmark on PyBullet envs in the [gSDE paper](#) if you want to compare TQC results to those of A2C/PPO/SAC/TD3.

---

Environments	SAC	TQC
	gSDE	gSDE
HalfCheetah	2984 +/- 202	3041 +/- 157
Ant	3102 +/- 37	3700 +/- 37
Hopper	2262 +/- 1	2401 +/- 62*
Walker2D	2136 +/- 67	2535 +/- 94
BipedalWalkerHardcore	13 +/- 18	228 +/- 18

\* with tuned hyperparameter `top_quantiles_to_drop_per_net` taken from the original paper

### 8.4.1 How to replicate the results?

Clone RL-Zoo and checkout the branch `feat/tqc`:

```
git clone https://github.com/DLR-RM/rl-baselines3-zoo
cd rl-baselines3-zoo/
git checkout feat/tqc
```

Run the benchmark (replace `$ENV_ID` by the envs mentioned above):

```
python train.py --algo tqc --env $ENV_ID --eval-episodes 10 --eval-freq 100000
```

Plot the results:

```
python scripts/all_plots.py -a tqc -e HalfCheetah Ant Hopper Walker2D_
↳BipedalWalkerHardcore -f logs/ -o logs/tqc_results
python scripts/plot_from_file.py -i logs/tqc_results.pkl -latex -l TQC
```

## 8.5 Comments

This implementation is based on SB3 SAC implementation and uses the code from the original TQC implementation for the quantile huber loss.

## 8.6 Parameters

```
class sb3_contrib.tqc.TQC(policy, env, learning_rate=0.0003, buffer_size=1000000, learning_starts=100,
    batch_size=256, tau=0.005, gamma=0.99, train_freq=1, gradient_steps=1,
    action_noise=None, replay_buffer_class=None, replay_buffer_kwargs=None,
    optimize_memory_usage=False, ent_coef='auto', target_update_interval=1,
    target_entropy='auto', top_quantiles_to_drop_per_net=2, use_sde=False,
    sde_sample_freq=-1, use_sde_at_warmup=False, tensorboard_log=None,
    create_eval_env=False, policy_kwargs=None, verbose=0, seed=None,
    device='auto', _init_setup_model=True)
```

Controlling Overestimation Bias with Truncated Mixture of Continuous Distributional Quantile Critics. Paper: <https://arxiv.org/abs/2005.04269> This implementation uses SB3 SAC implementation as base.

**Parameters**

- **policy** (Union[str, Type[TQCPolicy]]) – The policy model to use (MlpPolicy, CnnPolicy, ...)
- **env** (Union[Env, VecEnv, str]) – The environment to learn from (if registered in Gym, can be str)
- **learning\_rate** (Union[float, Callable]) – learning rate for adam optimizer, the same learning rate will be used for all networks (Q-Values, Actor and Value function) it can be a function of the current progress remaining (from 1 to 0)
- **buffer\_size** (int) – size of the replay buffer
- **learning\_starts** (int) – how many steps of the model to collect transitions for before learning starts
- **batch\_size** (int) – Minibatch size for each gradient update
- **tau** (float) – the soft update coefficient (“Polyak update”, between 0 and 1)
- **gamma** (float) – the discount factor
- **train\_freq** (int) – Update the model every `train_freq` steps. Alternatively pass a tuple of frequency and unit like (5, "step") or (2, "episode").
- **gradient\_steps** (int) – How many gradient update after each step
- **action\_noise** (Optional[ActionNoise]) – the action noise type (None by default), this can help for hard exploration problem. Cf `common.noise` for the different action noise type.
- **replay\_buffer\_class** (Optional[ReplayBuffer]) – Replay buffer class to use (for instance `HerReplayBuffer`). If None, it will be automatically selected.
- **replay\_buffer\_kwargs** (Optional[Dict[str, Any]]) – Keyword arguments to pass to the replay buffer on creation.
- **optimize\_memory\_usage** (bool) – Enable a memory efficient variant of the replay buffer at a cost of more complexity. See <https://github.com/DLR-RM/stable-baselines3/issues/37#issuecomment-637501195>
- **ent\_coef** (Union[str, float]) – Entropy regularization coefficient. (Equivalent to inverse of reward scale in the original SAC paper.) Controlling exploration/exploitation trade-off. Set it to ‘auto’ to learn it automatically (and ‘auto\_0.1’ for using 0.1 as initial value)
- **target\_update\_interval** (int) – update the target network every `target_network_update_freq` gradient steps.
- **target\_entropy** (Union[str, float]) – target entropy when learning `ent_coef` (`ent_coef = 'auto'`)
- **top\_quantiles\_to\_drop\_per\_net** (int) – Number of quantiles to drop per network
- **use\_sde** (bool) – Whether to use generalized State Dependent Exploration (gSDE) instead of action noise exploration (default: False)
- **sde\_sample\_freq** (int) – Sample a new noise matrix every `n` steps when using gSDE. Default: -1 (only sample at the beginning of the rollout)
- **use\_sde\_at\_warmup** (bool) – Whether to use gSDE instead of uniform sampling during the warm up phase (before learning starts)
- **create\_eval\_env** (bool) – Whether to create a second environment that will be used for evaluating the agent periodically (Only available when passing string for the environment). Caution, this parameter is deprecated and will be removed in the future.



- **policy\_kwargs** (Optional[Dict[str, Any]]) – additional arguments to be passed to the policy on creation
- **verbose** (int) – the verbosity level: 0 no output, 1 info, 2 debug
- **seed** (Optional[int]) – Seed for the pseudo random generators
- **device** (Union[device, str]) – Device (cpu, cuda, ...) on which the code should be run. Setting it to auto, the code will be run on the GPU if possible.
- **\_init\_setup\_model** (bool) – Whether or not to build the network at the creation of the instance

**collect\_rollouts**(*env, callback, train\_freq, replay\_buffer, action\_noise=None, learning\_starts=0, log\_interval=None*)

Collect experiences and store them into a ReplayBuffer.

#### Parameters

- **env** (VecEnv) – The training environment
- **callback** (BaseCallback) – Callback that will be called at each step (and at the beginning and end of the rollout)
- **train\_freq** (TrainFreq) – How much experience to collect by doing rollouts of current policy. Either TrainFreq(<n>, TrainFrequencyUnit.STEP) or TrainFreq(<n>, TrainFrequencyUnit.EPISODE) with <n> being an integer greater than 0.
- **action\_noise** (Optional[ActionNoise]) – Action noise that will be used for exploration Required for deterministic policy (e.g. TD3). This can also be used in addition to the stochastic policy for SAC.
- **learning\_starts** (int) – Number of steps before learning for the warm-up phase.
- **replay\_buffer** (ReplayBuffer) –
- **log\_interval** (Optional[int]) – Log data every log\_interval episodes

#### Return type

RolloutReturn

#### Returns

#### **get\_env()**

Returns the current environment (can be None if not defined).

#### Return type

Optional[VecEnv]

#### Returns

The current environment

#### **get\_parameters()**

Return the parameters of the agent. This includes parameters from different networks, e.g. critics (value functions) and policies (pi functions).

#### Return type

Dict[str, Dict]

#### Returns

Mapping of from names of the objects to PyTorch state-dicts.

**get\_vec\_normalize\_env()**

Return the VecNormalize wrapper of the training env if it exists.

**Return type**

Optional[VecNormalize]

**Returns**

The VecNormalize env.

**learn**(*total\_timesteps*, *callback=None*, *log\_interval=4*, *eval\_env=None*, *eval\_freq=-1*, *n\_eval\_episodes=5*, *tb\_log\_name='TQC'*, *eval\_log\_path=None*, *reset\_num\_timesteps=True*, *progress\_bar=False*)

Return a trained model.

**Parameters**

- **total\_timesteps** (int) – The total number of samples (env steps) to train on
- **callback** (Union[None, Callable, List[BaseCallback], BaseCallback]) – call-back(s) called at every step with state of the algorithm.
- **log\_interval** (int) – The number of timesteps before logging.
- **tb\_log\_name** (str) – the name of the run for TensorBoard logging
- **reset\_num\_timesteps** (bool) – whether or not to reset the current timestep number (used in logging)
- **progress\_bar** (bool) – Display a progress bar using tqdm and rich.

**Return type**

TypeVar(TQCSelf, bound= TQC)

**Returns**

the trained model

**classmethod load**(*path*, *env=None*, *device='auto'*, *custom\_objects=None*, *print\_system\_info=False*, *force\_reset=True*, *\*\*kwargs*)

Load the model from a zip-file. Warning: load re-creates the model from scratch, it does not update it in-place! For an in-place load use `set_parameters` instead.

**Parameters**

- **path** (Union[str, Path, BufferedIOBase]) – path to the file (or a file-like) where to load the agent from
- **env** (Union[Env, VecEnv, None]) – the new environment to run the loaded model on (can be None if you only need prediction from a trained model) has priority over any saved environment
- **device** (Union[device, str]) – Device on which the code should run.
- **custom\_objects** (Optional[Dict[str, Any]]) – Dictionary of objects to replace upon loading. If a variable is present in this dictionary as a key, it will not be deserialized and the corresponding item will be used instead. Similar to `custom_objects` in `keras.models.load_model`. Useful when you have an object in file that can not be deserialized.
- **print\_system\_info** (bool) – Whether to print system info from the saved model and the current system info (useful to debug loading issues)
- **force\_reset** (bool) – Force call to `reset()` before training to avoid unexpected behavior. See <https://github.com/DLR-RM/stable-baselines3/issues/597>
- **kwargs** – extra arguments to change the model when loading

**Return type**

TypeVar(SelfBaseAlgorithm, bound= BaseAlgorithm)

**Returns**

new model instance with loaded parameters

**load\_replay\_buffer**(*path, truncate\_last\_traj=True*)

Load a replay buffer from a pickle file.

**Parameters**

- **path** (Union[str, Path, BufferedIOBase]) – Path to the pickled replay buffer.
- **truncate\_last\_traj** (bool) – When using HerReplayBuffer with online sampling: If set to True, we assume that the last trajectory in the replay buffer was finished (and truncate it). If set to False, we assume that we continue the same trajectory (same episode).

**Return type**

None

**property logger: Logger**

Getter for the logger object.

**predict**(*observation, state=None, episode\_start=None, deterministic=False*)

Get the policy action from an observation (and optional hidden state). Includes sugar-coating to handle different observations (e.g. normalizing images).

**Parameters**

- **observation** (Union[ndarray, Dict[str, ndarray]]) – the input observation
- **state** (Optional[Tuple[ndarray, ...]]) – The last hidden states (can be None, used in recurrent policies)
- **episode\_start** (Optional[ndarray]) – The last masks (can be None, used in recurrent policies) this correspond to beginning of episodes, where the hidden states of the RNN must be reset.
- **deterministic** (bool) – Whether or not to return deterministic actions.

**Return type**

Tuple[ndarray, Optional[Tuple[ndarray, ...]]]

**Returns**

the model's action and the next hidden state (used in recurrent policies)

**save**(*path, exclude=None, include=None*)

Save all the attributes of the object and the model parameters in a zip-file.

**Parameters**

- **path** (Union[str, Path, BufferedIOBase]) – path to the file where the rl agent should be saved
- **exclude** (Optional[Iterable[str]]) – name of parameters that should be excluded in addition to the default ones
- **include** (Optional[Iterable[str]]) – name of parameters that might be excluded but should be included anyway

**Return type**

None

**save\_replay\_buffer**(*path*)

Save the replay buffer as a pickle file.

**Parameters**

**path** (Union[str, Path, BufferedIOBase]) – Path to the file where the replay buffer should be saved. if path is a str or pathlib.Path, the path is automatically created if necessary.

**Return type**

None

**set\_env**(*env*, *force\_reset=True*)

Checks the validity of the environment, and if it is coherent, set it as the current environment. Furthermore wrap any non vectorized env into a vectorized checked parameters: - observation\_space - action\_space

**Parameters**

- **env** (Union[Env, VecEnv]) – The environment for learning a policy
- **force\_reset** (bool) – Force call to `reset()` before training to avoid unexpected behavior. See issue <https://github.com/DLR-RM/stable-baselines3/issues/597>

**Return type**

None

**set\_logger**(*logger*)

Setter for for logger object. :rtype: None

**Warning:** When passing a custom logger object, this will overwrite `tensorboard_log` and `verbose` settings passed to the constructor.

**set\_parameters**(*load\_path\_or\_dict*, *exact\_match=True*, *device='auto'*)

Load parameters from a given zip-file or a nested dictionary containing parameters for different modules (see `get_parameters`).

**Parameters**

- **load\_path\_or\_iter** – Location of the saved data (path or file-like, see `save`), or a nested dictionary containing `nn.Module` parameters used by the policy. The dictionary maps object names to a state-dictionary returned by `torch.nn.Module.state_dict()`.
- **exact\_match** (bool) – If True, the given parameters should include parameters for each module and each of their parameters, otherwise raises an Exception. If set to False, this can be used to update only specific parameters.
- **device** (Union[device, str]) – Device on which the code should run.

**Return type**

None

**set\_random\_seed**(*seed=None*)

Set the seed of the pseudo-random generators (python, numpy, pytorch, gym, action\_space)

**Parameters**

**seed** (Optional[int]) –

**Return type**

None

**train**(*gradient\_steps*, *batch\_size*=64)

Sample the replay buffer and do the updates (gradient descent and update target networks)

**Return type**

None

## 8.7 TQC Policies

`sb3_contrib.tqc.MlpPolicy`

alias of TQCPolicy

```
class sb3_contrib.tqc.policies.TQCPolicy(observation_space, action_space, lr_schedule, net_arch=None,
                                         activation_fn=<class 'torch.nn.modules.activation.ReLU'>,
                                         use_sde=False, log_std_init=-3, sde_net_arch=None,
                                         use_expln=False, clip_mean=2.0,
                                         features_extractor_class=<class
                                         'stable_baselines3.common.torch_layers.FlattenExtractor'>,
                                         features_extractor_kwargs=None, normalize_images=True,
                                         optimizer_class=<class 'torch.optim.adam.Adam'>,
                                         optimizer_kwargs=None, n_quantiles=25, n_critics=2,
                                         share_features_extractor=False)
```

Policy class (with both actor and critic) for TQC.

### Parameters

- **observation\_space** (Space) – Observation space
- **action\_space** (Space) – Action space
- **lr\_schedule** (Callable[[float], float]) – Learning rate schedule (could be constant)
- **net\_arch** (Union[List[int], Dict[str, List[int]], None]) – The specification of the policy and value networks.
- **activation\_fn** (Type[Module]) – Activation function
- **use\_sde** (bool) – Whether to use State Dependent Exploration or not
- **log\_std\_init** (float) – Initial value for the log standard deviation
- **sde\_net\_arch** (Optional[List[int]]) – Network architecture for extracting features when using gSDE. If None, the latent features from the policy will be used. Pass an empty list to use the states as features.
- **use\_expln** (bool) – Use `expln()` function instead of `exp()` when using gSDE to ensure a positive standard deviation (cf paper). It allows to keep variance above zero and prevent it from growing too fast. In practice, `exp()` is usually enough.
- **clip\_mean** (float) – Clip the mean output when using gSDE to avoid numerical instability.
- **features\_extractor\_class** (Type[BaseFeaturesExtractor]) – Features extractor to use.
- **features\_extractor\_kwargs** (Optional[Dict[str, Any]]) – Keyword arguments to pass to the feature extractor.
- **normalize\_images** (bool) – Whether to normalize images or not, dividing by 255.0 (True by default)
- **optimizer\_class** (Type[Optimizer]) – The optimizer to use, `th.optim.Adam` by default

- **optimizer\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments, excluding the learning rate, to pass to the optimizer
- **n\_quantiles** (int) – Number of quantiles for the critic.
- **n\_critics** (int) – Number of critic networks to create.
- **share\_features\_extractor** (bool) – Whether to share or not the features extractor between the actor and the critic (this saves computation time)

**forward**(*obs*, *deterministic=False*)

Defines the computation performed at every call.

Should be overridden by all subclasses. :rtype: Tensor

---

**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.

---

**reset\_noise**(*batch\_size=1*)

Sample new weights for the exploration matrix, when using gSDE.

**Parameters**

**batch\_size** (int) –

**Return type**

None

**set\_training\_mode**(*mode*)

Put the policy in either training or evaluation mode. This affects certain modules, such as batch normalisation and dropout. :type mode: bool :param mode: if true, set to training mode, else set to evaluation mode

**Return type**

None

```
class sb3_contrib.tqc.CnnPolicy(observation_space, action_space, lr_schedule, net_arch=None,
                                activation_fn=<class 'torch.nn.modules.activation.ReLU'>,
                                use_sde=False, log_std_init=-3, sde_net_arch=None, use_expln=False,
                                clip_mean=2.0, features_extractor_class=<class
                                'stable_baselines3.common.torch_layers.NatureCNN'>,
                                features_extractor_kwargs=None, normalize_images=True,
                                optimizer_class=<class 'torch.optim.adam.Adam'>,
                                optimizer_kwargs=None, n_quantiles=25, n_critics=2,
                                share_features_extractor=False)
```

Policy class (with both actor and critic) for TQC.

**Parameters**

- **observation\_space** (Space) – Observation space
- **action\_space** (Space) – Action space
- **lr\_schedule** (Callable[[float], float]) – Learning rate schedule (could be constant)
- **net\_arch** (Union[List[int], Dict[str, List[int]], None]) – The specification of the policy and value networks.
- **activation\_fn** (Type[Module]) – Activation function

- **use\_sde** (bool) – Whether to use State Dependent Exploration or not
- **log\_std\_init** (float) – Initial value for the log standard deviation
- **sde\_net\_arch** (Optional[List[int]]) – Network architecture for extracting features when using gSDE. If None, the latent features from the policy will be used. Pass an empty list to use the states as features.
- **use\_expln** (bool) – Use `expln()` function instead of `exp()` when using gSDE to ensure a positive standard deviation (cf paper). It allows to keep variance above zero and prevent it from growing too fast. In practice, `exp()` is usually enough.
- **clip\_mean** (float) – Clip the mean output when using gSDE to avoid numerical instability.
- **features\_extractor\_class** (Type[BaseFeaturesExtractor]) – Features extractor to use.
- **normalize\_images** (bool) – Whether to normalize images or not, dividing by 255.0 (True by default)
- **optimizer\_class** (Type[Optimizer]) – The optimizer to use, `th.optim.Adam` by default
- **optimizer\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments, excluding the learning rate, to pass to the optimizer
- **n\_quantiles** (int) – Number of quantiles for the critic.
- **n\_critics** (int) – Number of critic networks to create.
- **share\_features\_extractor** (bool) – Whether to share or not the features extractor between the actor and the critic (this saves computation time)





## TRPO

Trust Region Policy Optimization (TRPO) is an iterative approach for optimizing policies with guaranteed monotonic improvement.

### Available Policies

<i>MlpPolicy</i>	alias of <code>ActorCriticPolicy</code>
<i>CnnPolicy</i>	alias of <code>ActorCriticCnnPolicy</code>
<i>MultiInputPolicy</i>	alias of <code>MultiInputActorCriticPolicy</code>

## 9.1 Notes

- Original paper: <https://arxiv.org/abs/1502.05477>
- OpenAI blog post: <https://blog.openai.com/openai-baselines-ppo/>

## 9.2 Can I use?

- Recurrent policies:
- Multi processing: ✓
- Gym spaces:

Space	Action	Observation
Discrete	✓	✓
Box	✓	✓
MultiDiscrete	✓	✓
MultiBinary	✓	✓
Dict		✓

## 9.3 Example

```
import gym
import numpy as np

from sb3_contrib import TRPO

env = gym.make("Pendulum-v1")

model = TRPO("MlpPolicy", env, verbose=1)
model.learn(total_timesteps=10000, log_interval=4)
model.save("trpo_pendulum")

del model # remove to demonstrate saving and loading

model = TRPO.load("trpo_pendulum")

obs = env.reset()
while True:
    action, _states = model.predict(obs, deterministic=True)
    obs, reward, done, info = env.step(action)
    env.render()
    if done:
        obs = env.reset()
```

## 9.4 Results

Result on the MuJoCo benchmark (1M steps on -v3 envs with MuJoCo v2.1.0) using 3 seeds. The complete learning curves are available in the [associated PR](#).

Environments	TRPO
HalfCheetah	1803 +/- 46
Ant	3554 +/- 591
Hopper	3372 +/- 215
Walker2d	4502 +/- 234
Swimmer	359 +/- 2

### 9.4.1 How to replicate the results?

Clone RL-Zoo and checkout the branch feat/trpo:

```
git clone https://github.com/cyprienc/rl-baselines3-zoo
cd rl-baselines3-zoo/
```

Run the benchmark (replace \$ENV\_ID by the envs mentioned above):

```
python train.py --algo tqc --env $ENV_ID --n-eval-envs 10 --eval-episodes 20 --eval-freq 50000
```

Plot the results:

```
python scripts/all_plots.py -a trpo -e HalfCheetah Ant Hopper Walker2d Swimmer -f logs/ -
  ↳o logs/trpo_results
python scripts/plot_from_file.py -i logs/trpo_results.pkl -latex -l TRPO
```

## 9.5 Parameters

```
class sb3_contrib.trpo.TRPO(policy, env, learning_rate=0.001, n_steps=2048, batch_size=128, gamma=0.99,
                             cg_max_steps=15, cg_damping=0.1, line_search_shrinking_factor=0.8,
                             line_search_max_iter=10, n_critic_updates=10, gae_lambda=0.95,
                             use_sde=False, sde_sample_freq=-1, normalize_advantage=True,
                             target_kl=0.01, sub_sampling_factor=1, tensorboard_log=None,
                             create_eval_env=False, policy_kwargs=None, verbose=0, seed=None,
                             device='auto', _init_setup_model=True)
```

Trust Region Policy Optimization (TRPO)

Paper: <https://arxiv.org/abs/1502.05477> Code: This implementation borrows code from OpenAI Spinning Up (<https://github.com/openai/spinningup/>) and Stable Baselines (TRPO from <https://github.com/hill-a/stable-baselines>)

Introduction to TRPO: <https://spinningup.openai.com/en/latest/algorithms/trpo.html>

### Parameters

- **policy** (Union[str, Type[ActorCriticPolicy]]) – The policy model to use (MlpPolicy, CnnPolicy, ...)
- **env** (Union[Env, VecEnv, str]) – The environment to learn from (if registered in Gym, can be str)
- **learning\_rate** (Union[float, Callable[[float], float]]) – The learning rate for the value function, it can be a function of the current progress remaining (from 1 to 0)
- **n\_steps** (int) – The number of steps to run for each environment per update (i.e. rollout buffer size is  $n\_steps * n\_envs$  where  $n\_envs$  is number of environment copies running in parallel) NOTE:  $n\_steps * n\_envs$  must be greater than 1 (because of the advantage normalization) See <https://github.com/pytorch/pytorch/issues/29372>
- **batch\_size** (int) – Minibatch size for the value function
- **gamma** (float) – Discount factor
- **cg\_max\_steps** (int) – maximum number of steps in the Conjugate Gradient algorithm for computing the Hessian vector product
- **cg\_damping** (float) – damping in the Hessian vector product computation
- **line\_search\_shrinking\_factor** (float) – step-size reduction factor for the line-search (i.e.,  $\theta_{new} = \theta + \alpha^i * step$ )
- **line\_search\_max\_iter** (int) – maximum number of iteration for the backtracking line-search
- **n\_critic\_updates** (int) – number of critic updates per policy update
- **gae\_lambda** (float) – Factor for trade-off of bias vs variance for Generalized Advantage Estimator
- **use\_sde** (bool) – Whether to use generalized State Dependent Exploration (gSDE) instead of action noise exploration (default: False)

- **sde\_sample\_freq** (int) – Sample a new noise matrix every *n* steps when using gSDE  
Default: -1 (only sample at the beginning of the rollout)
- **normalize\_advantage** (bool) – Whether to normalize or not the advantage
- **target\_kl** (float) – Target Kullback-Leibler divergence between updates. Should be small for stability. Values like 0.01, 0.05.
- **sub\_sampling\_factor** (int) – Sub-sample the batch to make computation faster see p40-42 of John Schulman thesis <http://joschu.net/docs/thesis.pdf>
- **tensorboard\_log** (Optional[str]) – the log location for tensorboard (if None, no logging)
- **create\_eval\_env** (bool) – Whether to create a second environment that will be used for evaluating the agent periodically (Only available when passing string for the environment). Caution, this parameter is deprecated and will be removed in the future.
- **policy\_kwargs** (Optional[Dict[str, Any]]) – additional arguments to be passed to the policy on creation
- **verbose** (int) – the verbosity level: 0 no output, 1 info, 2 debug
- **seed** (Optional[int]) – Seed for the pseudo random generators
- **device** (Union[device, str]) – Device (cpu, cuda, ...) on which the code should be run. Setting it to auto, the code will be run on the GPU if possible.
- **\_init\_setup\_model** (bool) – Whether or not to build the network at the creation of the instance

**collect\_rollouts**(*env, callback, rollout\_buffer, n\_rollout\_steps*)

Collect experiences using the current policy and fill a `RolloutBuffer`. The term rollout here refers to the model-free notion and should not be used with the concept of rollout used in model-based RL or planning.

#### Parameters

- **env** (VecEnv) – The training environment
- **callback** (BaseCallback) – Callback that will be called at each step (and at the beginning and end of the rollout)
- **rollout\_buffer** (RolloutBuffer) – Buffer to fill with rollouts
- **n\_rollout\_steps** (int) – Number of experiences to collect per environment

#### Return type

bool

#### Returns

True if function returned with at least *n\_rollout\_steps* collected, False if callback terminated rollout prematurely.

**get\_env**()

Returns the current environment (can be None if not defined).

#### Return type

Optional[VecEnv]

#### Returns

The current environment

**get\_parameters()**

Return the parameters of the agent. This includes parameters from different networks, e.g. critics (value functions) and policies (pi functions).

**Return type**

Dict[str, Dict]

**Returns**

Mapping of from names of the objects to PyTorch state-dicts.

**get\_vec\_normalize\_env()**

Return the VecNormalize wrapper of the training env if it exists.

**Return type**

Optional[VecNormalize]

**Returns**

The VecNormalize env.

**hessian\_vector\_product(params, grad\_kl, vector, retain\_graph=True)**

Computes the matrix-vector product with the Fisher information matrix.

**Parameters**

- **params** (List[Parameter]) – list of parameters used to compute the Hessian
- **grad\_kl** (Tensor) – flattened gradient of the KL divergence between the old and new policy
- **vector** (Tensor) – vector to compute the dot product the hessian-vector dot product with
- **retain\_graph** (bool) – if True, the graph will be kept after computing the Hessian

**Return type**

Tensor

**Returns**

Hessian-vector dot product (with damping)

**learn(total\_timesteps, callback=None, log\_interval=1, eval\_env=None, eval\_freq=-1, n\_eval\_episodes=5, tb\_log\_name='TRPO', eval\_log\_path=None, reset\_num\_timesteps=True, progress\_bar=False)**

Return a trained model.

**Parameters**

- **total\_timesteps** (int) – The total number of samples (env steps) to train on
- **callback** (Union[None, Callable, List[BaseCallback], BaseCallback]) – call-back(s) called at every step with state of the algorithm.
- **log\_interval** (int) – The number of timesteps before logging.
- **tb\_log\_name** (str) – the name of the run for TensorBoard logging
- **reset\_num\_timesteps** (bool) – whether or not to reset the current timestep number (used in logging)
- **progress\_bar** (bool) – Display a progress bar using tqdm and rich.

**Return type**

TypeVar(TRPOSelf, bound= TRPO)

**Returns**

the trained model

```
classmethod load(path, env=None, device='auto', custom_objects=None, print_system_info=False,
                  force_reset=True, **kwargs)
```

Load the model from a zip-file. Warning: load re-creates the model from scratch, it does not update it in-place! For an in-place load use `set_parameters` instead.

#### Parameters

- **path** (Union[str, Path, BufferedIOBase]) – path to the file (or a file-like) where to load the agent from
- **env** (Union[Env, VecEnv, None]) – the new environment to run the loaded model on (can be None if you only need prediction from a trained model) has priority over any saved environment
- **device** (Union[device, str]) – Device on which the code should run.
- **custom\_objects** (Optional[Dict[str, Any]]) – Dictionary of objects to replace upon loading. If a variable is present in this dictionary as a key, it will not be deserialized and the corresponding item will be used instead. Similar to `custom_objects` in `keras.models.load_model`. Useful when you have an object in file that can not be deserialized.
- **print\_system\_info** (bool) – Whether to print system info from the saved model and the current system info (useful to debug loading issues)
- **force\_reset** (bool) – Force call to `reset()` before training to avoid unexpected behavior. See <https://github.com/DLR-RM/stable-baselines3/issues/597>
- **kwargs** – extra arguments to change the model when loading

#### Return type

TypeVar(SelfBaseAlgorithm, bound= BaseAlgorithm)

#### Returns

new model instance with loaded parameters

#### property logger: **Logger**

Getter for the logger object.

```
predict(observation, state=None, episode_start=None, deterministic=False)
```

Get the policy action from an observation (and optional hidden state). Includes sugar-coating to handle different observations (e.g. normalizing images).

#### Parameters

- **observation** (Union[ndarray, Dict[str, ndarray]]) – the input observation
- **state** (Optional[Tuple[ndarray, ...]]) – The last hidden states (can be None, used in recurrent policies)
- **episode\_start** (Optional[ndarray]) – The last masks (can be None, used in recurrent policies) this correspond to beginning of episodes, where the hidden states of the RNN must be reset.
- **deterministic** (bool) – Whether or not to return deterministic actions.

#### Return type

Tuple[ndarray, Optional[Tuple[ndarray, ...]]]

#### Returns

the model's action and the next hidden state (used in recurrent policies)

**save**(*path*, *exclude=None*, *include=None*)

Save all the attributes of the object and the model parameters in a zip-file.

#### Parameters

- **path** (Union[str, Path, BufferedIOBase]) – path to the file where the rl agent should be saved
- **exclude** (Optional[Iterable[str]]) – name of parameters that should be excluded in addition to the default ones
- **include** (Optional[Iterable[str]]) – name of parameters that might be excluded but should be included anyway

#### Return type

None

**set\_env**(*env*, *force\_reset=True*)

Checks the validity of the environment, and if it is coherent, set it as the current environment. Furthermore wrap any non vectorized env into a vectorized checked parameters: - observation\_space - action\_space

#### Parameters

- **env** (Union[Env, VecEnv]) – The environment for learning a policy
- **force\_reset** (bool) – Force call to `reset()` before training to avoid unexpected behavior. See issue <https://github.com/DLR-RM/stable-baselines3/issues/597>

#### Return type

None

**set\_logger**(*logger*)

Setter for for logger object. :rtype: None

**Warning:** When passing a custom logger object, this will overwrite `tensorboard_log` and `verbose` settings passed to the constructor.

**set\_parameters**(*load\_path\_or\_dict*, *exact\_match=True*, *device='auto'*)

Load parameters from a given zip-file or a nested dictionary containing parameters for different modules (see `get_parameters`).

#### Parameters

- **load\_path\_or\_iter** – Location of the saved data (path or file-like, see `save`), or a nested dictionary containing `nn.Module` parameters used by the policy. The dictionary maps object names to a state-dictionary returned by `torch.nn.Module.state_dict()`.
- **exact\_match** (bool) – If True, the given parameters should include parameters for each module and each of their parameters, otherwise raises an Exception. If set to False, this can be used to update only specific parameters.
- **device** (Union[device, str]) – Device on which the code should run.

#### Return type

None

**set\_random\_seed**(*seed=None*)

Set the seed of the pseudo-random generators (python, numpy, pytorch, gym, action\_space)

#### Parameters

- **seed** (Optional[int]) –

**Return type**

None

**train()**

Update policy using the currently gathered rollout buffer.

**Return type**

None

## 9.6 TRPO Policies

`sb3_contrib.trpo.MlpPolicy`

alias of `ActorCriticPolicy`

```
class stable_baselines3.common.policies.ActorCriticPolicy(observation_space, action_space,  
                                                         lr_schedule, net_arch=None,  
                                                         activation_fn=<class  
                                                         'torch.nn.modules.activation.Tanh'>,  
                                                         ortho_init=True, use_sde=False,  
                                                         log_std_init=0.0, full_std=True,  
                                                         use_expln=False, squash_output=False,  
                                                         features_extractor_class=<class 'sta-  
                                                         ble_baselines3.common.torch_layers.FlattenExtractor'>,  
                                                         features_extractor_kwargs=None,  
                                                         share_features_extractor=True,  
                                                         normalize_images=True,  
                                                         optimizer_class=<class  
                                                         'torch.optim.adam.Adam'>,  
                                                         optimizer_kwargs=None)
```

Policy class for actor-critic algorithms (has both policy and value prediction). Used by A2C, PPO and the likes.

### Parameters

- **observation\_space** (Space) – Observation space
- **action\_space** (Space) – Action space
- **lr\_schedule** (Callable[[float], float]) – Learning rate schedule (could be constant)
- **net\_arch** (Union[List[int], Dict[str, List[int]], List[Dict[str, List[int]]], None]) – The specification of the policy and value networks.
- **activation\_fn** (Type[Module]) – Activation function
- **ortho\_init** (bool) – Whether to use or not orthogonal initialization
- **use\_sde** (bool) – Whether to use State Dependent Exploration or not
- **log\_std\_init** (float) – Initial value for the log standard deviation
- **full\_std** (bool) – Whether to use (n\_features x n\_actions) parameters for the std instead of only (n\_features,) when using gSDE
- **use\_expln** (bool) – Use `expln()` function instead of `exp()` to ensure a positive standard deviation (cf paper). It allows to keep variance above zero and prevent it from growing too fast. In practice, `exp()` is usually enough.
- **squash\_output** (bool) – Whether to squash the output using a tanh function, this allows to ensure boundaries when using gSDE.



- **features\_extractor\_class** (Type[BaseFeaturesExtractor]) – Features extractor to use.
- **features\_extractor\_kwargs** (Optional[Dict[str, Any]]) – Keyword arguments to pass to the features extractor.
- **share\_features\_extractor** (bool) – If True, the features extractor is shared between the policy and value networks.
- **normalize\_images** (bool) – Whether to normalize images or not, dividing by 255.0 (True by default)
- **optimizer\_class** (Type[Optimizer]) – The optimizer to use, `th.optim.Adam` by default
- **optimizer\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments, excluding the learning rate, to pass to the optimizer

**evaluate\_actions**(*obs*, *actions*)

Evaluate actions according to the current policy, given the observations.

**Parameters**

- **obs** (Tensor) – Observation
- **actions** (Tensor) – Actions

**Return type**

Tuple[Tensor, Tensor, Optional[Tensor]]

**Returns**

estimated value, log likelihood of taking those actions and entropy of the action distribution.

**extract\_features**(*obs*)

Preprocess the observation if needed and extract features.

**Parameters**

- **obs** (Tensor) – Observation

**Return type**

Union[Tensor, Tuple[Tensor, Tensor]]

**Returns**

the output of the features extractor(s)

**forward**(*obs*, *deterministic=False*)

Forward pass in all the networks (actor and critic)

**Parameters**

- **obs** (Tensor) – Observation
- **deterministic** (bool) – Whether to sample or use deterministic actions

**Return type**

Tuple[Tensor, Tensor, Tensor]

**Returns**

action, value and log probability of the action

**get\_distribution**(*obs*)

Get the current policy distribution given the observations.

**Parameters**

- **obs** (Tensor) –

**Return type**  
Distribution

**Returns**  
the action distribution.

**predict\_values**(*obs*)

Get the estimated values according to the current policy given the observations.

**Parameters**  
**obs** (Tensor) – Observation

**Return type**  
Tensor

**Returns**  
the estimated values.

**reset\_noise**(*n\_envs=1*)

Sample new weights for the exploration matrix.

**Parameters**  
**n\_envs** (int) –

**Return type**  
None

`sb3_contrib.trpo.CnnPolicy`  
alias of `ActorCriticCnnPolicy`

```
class stable_baselines3.common.policies.ActorCriticCnnPolicy(observation_space, action_space,  
                                                         lr_schedule, net_arch=None,  
                                                         activation_fn=<class  
                                                         'torch.nn.modules.activation.Tanh'>,  
                                                         ortho_init=True, use_sde=False,  
                                                         log_std_init=0.0, full_std=True,  
                                                         use_expln=False,  
                                                         squash_output=False,  
                                                         features_extractor_class=<class  
                                                         'stable_baselines3.common.torch_layers.NatureCNN'>,  
                                                         features_extractor_kwargs=None,  
                                                         share_features_extractor=True,  
                                                         normalize_images=True,  
                                                         optimizer_class=<class  
                                                         'torch.optim.adam.Adam'>,  
                                                         optimizer_kwargs=None)
```

CNN policy class for actor-critic algorithms (has both policy and value prediction). Used by A2C, PPO and the likes.

**Parameters**

- **observation\_space** (Space) – Observation space
- **action\_space** (Space) – Action space
- **lr\_schedule** (Callable[[float], float]) – Learning rate schedule (could be constant)
- **net\_arch** (Union[List[int], Dict[str, List[int]], List[Dict[str, List[int]]], None]) – The specification of the policy and value networks.

- **activation\_fn** (Type[Module]) – Activation function
- **ortho\_init** (bool) – Whether to use or not orthogonal initialization
- **use\_sde** (bool) – Whether to use State Dependent Exploration or not
- **log\_std\_init** (float) – Initial value for the log standard deviation
- **full\_std** (bool) – Whether to use (n\_features x n\_actions) parameters for the std instead of only (n\_features,) when using gSDE
- **use\_expln** (bool) – Use `expln()` function instead of `exp()` to ensure a positive standard deviation (cf paper). It allows to keep variance above zero and prevent it from growing too fast. In practice, `exp()` is usually enough.
- **squash\_output** (bool) – Whether to squash the output using a tanh function, this allows to ensure boundaries when using gSDE.
- **features\_extractor\_class** (Type[BaseFeaturesExtractor]) – Features extractor to use.
- **features\_extractor\_kwargs** (Optional[Dict[str, Any]]) – Keyword arguments to pass to the features extractor.
- **share\_features\_extractor** (bool) – If True, the features extractor is shared between the policy and value networks.
- **normalize\_images** (bool) – Whether to normalize images or not, dividing by 255.0 (True by default)
- **optimizer\_class** (Type[Optimizer]) – The optimizer to use, `th.optim.Adam` by default
- **optimizer\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments, excluding the learning rate, to pass to the optimizer

`sb3_contrib.trpo.MultiInputPolicy`

alias of `MultiInputActorCriticPolicy`

```
class stable_baselines3.common.policies.MultiInputActorCriticPolicy(observation_space,
                                                                    action_space, lr_schedule,
                                                                    net_arch=None,
                                                                    activation_fn=<class
                                                                    'torch.nn.modules.activation.Tanh'>,
                                                                    ortho_init=True,
                                                                    use_sde=False,
                                                                    log_std_init=0.0,
                                                                    full_std=True,
                                                                    use_expln=False,
                                                                    squash_output=False, fea-
                                                                    tures_extractor_class=<class
                                                                    'sta-
                                                                    ble_baselines3.common.torch_layers.Combine
                                                                    fea-
                                                                    tures_extractor_kwargs=None,
                                                                    share_features_extractor=True,
                                                                    normalize_images=True,
                                                                    optimizer_class=<class
                                                                    'torch.optim.adam.Adam'>,
                                                                    optimizer_kwargs=None)
```

MultiInputActorClass policy class for actor-critic algorithms (has both policy and value prediction). Used by A2C, PPO and the likes.

#### Parameters

- **observation\_space** (Dict) – Observation space (Tuple)
- **action\_space** (Space) – Action space
- **lr\_schedule** (Callable[[float], float]) – Learning rate schedule (could be constant)
- **net\_arch** (Union[List[int], Dict[str, List[int]], List[Dict[str, List[int]]], None]) – The specification of the policy and value networks.
- **activation\_fn** (Type[Module]) – Activation function
- **ortho\_init** (bool) – Whether to use or not orthogonal initialization
- **use\_sde** (bool) – Whether to use State Dependent Exploration or not
- **log\_std\_init** (float) – Initial value for the log standard deviation
- **full\_std** (bool) – Whether to use (n\_features x n\_actions) parameters for the std instead of only (n\_features,) when using gSDE
- **use\_expln** (bool) – Use `expln()` function instead of `exp()` to ensure a positive standard deviation (cf paper). It allows to keep variance above zero and prevent it from growing too fast. In practice, `exp()` is usually enough.
- **squash\_output** (bool) – Whether to squash the output using a tanh function, this allows to ensure boundaries when using gSDE.
- **features\_extractor\_class** (Type[BaseFeaturesExtractor]) – Uses the CombinedExtractor
- **features\_extractor\_kwargs** (Optional[Dict[str, Any]]) – Keyword arguments to pass to the features extractor.
- **share\_features\_extractor** (bool) – If True, the features extractor is shared between the policy and value networks.
- **normalize\_images** (bool) – Whether to normalize images or not, dividing by 255.0 (True by default)
- **optimizer\_class** (Type[Optimizer]) – The optimizer to use, `th.optim.Adam` by default
- **optimizer\_kwargs** (Optional[Dict[str, Any]]) – Additional keyword arguments, excluding the learning rate, to pass to the optimizer

`sb3_contrib.common.utils.conjugate_gradient_solver(matrix_vector_dot_fn, b, max_iter=10, residual_tol=1e-10)`

Finds an approximate solution to a set of linear equations  $Ax = b$

**Sources:**

- [https://github.com/ajlangley/trpo-pytorch/blob/master/conjugate\\_gradient.py](https://github.com/ajlangley/trpo-pytorch/blob/master/conjugate_gradient.py)
- [https://github.com/joschu/modular\\_rl/blob/master/modular\\_rl/trpo.py#L122](https://github.com/joschu/modular_rl/blob/master/modular_rl/trpo.py#L122)

**Reference:**

- <https://epubs.siam.org/doi/abs/10.1137/1.9781611971446.ch6>

**Parameters**

- **matrix\_vector\_dot\_fn** (Callable[[Tensor], Tensor]) – a function that right multiplies a matrix  $A$  by a vector  $v$
- **b** – the right hand term in the set of linear equations  $Ax = b$
- **max\_iter** – the maximum number of iterations (default is 10)
- **residual\_tol** – residual tolerance for early stopping of the solving (default is 1e-10)

**Return x**

the approximate solution to the system of equations defined by *matrix\_vector\_dot\_fn* and *b*

**Return type**

Tensor

`sb3_contrib.common.utils.flat_grad(output, parameters, create_graph=False, retain_graph=False)`

Returns the gradients of the passed sequence of parameters into a flat gradient. Order of parameters is preserved.

**Parameters**

- **output** – functional output to compute the gradient for
- **parameters** (Sequence[Parameter]) – sequence of Parameter
- **retain\_graph** (bool) – – If False, the graph used to compute the grad will be freed. Defaults to the value of *create\_graph*.
- **create\_graph** (bool) – – If True, graph of the derivative will be constructed, allowing to compute higher order derivative products. Default: False.

**Return type**

Tensor

**Returns**

Tensor containing the flattened gradients

`sb3_contrib.common.utils.quantile_huber_loss(current_quantiles, target_quantiles, cum_prob=None, sum_over_quantiles=True)`

The quantile-regression loss, as described in the QR-DQN and TQC papers. Partially taken from [https://github.com/bayesgroup/tqc\\_pytorch](https://github.com/bayesgroup/tqc_pytorch).

**Parameters**

- **current\_quantiles** (Tensor) – current estimate of quantiles, must be either (batch\_size, n\_quantiles) or (batch\_size, n\_critics, n\_quantiles)
- **target\_quantiles** (Tensor) – target of quantiles, must be either (batch\_size, n\_target\_quantiles), (batch\_size, 1, n\_target\_quantiles), or (batch\_size, n\_critics, n\_target\_quantiles)
- **cum\_prob** (Optional[Tensor]) – cumulative probabilities to calculate quantiles (also called midpoints in QR-DQN paper), must be either (batch\_size, n\_quantiles), (batch\_size, 1, n\_quantiles), or (batch\_size, n\_critics, n\_quantiles). (if None, calculating unit quantiles)
- **sum\_over\_quantiles** (bool) – if summing over the quantile dimension or not

**Return type**

Tensor

**Returns**

the loss

## GYM WRAPPERS

Additional [Gym Wrappers](#) to enhance Gym environments.

### 11.1 TimeFeatureWrapper

```
class sb3_contrib.common.wrappers.TimeFeatureWrapper(env, max_steps=1000, test_mode=False)
```

Add remaining, normalized time to observation space for fixed length episodes. See <https://arxiv.org/abs/1712.00378> and <https://github.com/aravindr93/mjrl/issues/13>.

---

**Note:** Only `gym.spaces.Box` and `gym.spaces.Dict` (`gym.GoalEnv`) 1D observation spaces are supported for now.

---

#### Parameters

- **env** (Env) – Gym env to wrap.
- **max\_steps** (int) – Max number of steps of an episode if it is not wrapped in a `TimeLimit` object.
- **test\_mode** (bool) – In test mode, the time feature is constant, equal to zero. This allow to check that the agent did not overfit this feature, learning a deterministic pre-defined sequence of actions.

#### reset()

Resets the environment to an initial state and returns an initial observation.

Note that this function should not reset the environment's random number generator(s); random variables in the environment's state should be sampled independently between multiple calls to `reset()`. In other words, each call of `reset()` should yield an environment suitable for a new episode, independent of previous episodes.

#### Return type

`Union[Tuple, Dict[str, Any], ndarray, int]`

#### Returns:

observation (object): the initial observation.

**step**(*action*)

Run one timestep of the environment's dynamics. When end of episode is reached, you are responsible for calling *reset()* to reset this environment's state.

Accepts an action and returns a tuple (observation, reward, done, info).

**Return type**

Tuple[Union[Tuple, Dict[str, Any], ndarray, int], float, bool, Dict]

**Args:**

action (object): an action provided by the agent

**Returns:**

observation (object): agent's observation of the current environment  
reward (float) : amount of reward returned after previous action  
done (bool): whether the episode has ended, in which case further step() calls will return undefined results  
info (dict): contains auxiliary diagnostic information (helpful for debugging, and sometimes learning)



## CHANGELOG

### 12.1 Release 1.6.2 (2022-10-10)

Progress bar and upgrade to latest SB3 version

#### 12.1.1 Breaking Changes:

- Upgraded to Stable-Baselines3  $\geq$  1.6.2

#### 12.1.2 New Features:

- Added `progress_bar` argument in the `learn()` method, displayed using TQDM and rich packages

#### 12.1.3 Bug Fixes:

#### 12.1.4 Deprecations:

- Deprecate parameters `eval_env`, `eval_freq` and `create_eval_env`

#### 12.1.5 Others:

- Fixed the return type of `.load()` methods so that they now use `TypeVar`

### 12.2 Release 1.6.1 (2022-09-29)

Bug fix release

### 12.2.1 Breaking Changes:

- Fixed the issue that `predict` does not always return action as `np.ndarray` (@qgallouedec)
- Upgraded to Stable-Baselines3  $\geq 1.6.1$

### 12.2.2 New Features:

### 12.2.3 Bug Fixes:

- Fixed the issue of wrongly passing policy arguments when using `CnnLstmPolicy` or `MultiInputLstmPolicy` with `RecurrentPPO` (@mlodel)
- Fixed division by zero error when computing FPS when a small number of time has elapsed in operating systems with low-precision timers.
- Fixed calling child callbacks in `MaskableEvalCallback` (@CppMaster)
- Fixed missing verbose parameter passing in the `MaskableEvalCallback` constructor (@burakdmb)
- Fixed the issue that when updating the target network in QRDQN, TQC, the `running_mean` and `running_var` properties of batch norm layers are not updated (@honglu2875)

### 12.2.4 Deprecations:

### 12.2.5 Others:

- Changed the default buffer device from "cpu" to "auto"

## 12.3 Release 1.6.0 (2022-07-11)

### Add RecurrentPPO (aka PPO LSTM)

### 12.3.1 Breaking Changes:

- Upgraded to Stable-Baselines3  $\geq 1.6.0$
- Changed the way policy “aliases” are handled (“MlpPolicy”, “CnnPolicy”, ...), removing the former `register_policy` helper, `policy_base` parameter and using `policy_aliases` static attributes instead (@Gregwar)
- Renamed `rollout/exploration rate` key to `rollout/exploration_rate` for QRDQN (to be consistent with SB3 DQN)
- Upgraded to python 3.7+ syntax using `pyupgrade`
- SB3 now requires PyTorch  $\geq 1.11$
- Changed the default network architecture when using `CnnPolicy` or `MultiInputPolicy` with TQC, `share_features_extractor` is now set to `False` by default and the `net_arch=[256, 256]` (instead of `net_arch=[]` that was before)

### 12.3.2 New Features:

- Added RecurrentPPO (aka PPO LSTM)

### 12.3.3 Bug Fixes:

- Fixed a bug in RecurrentPPO when calculating the masked loss functions (@rnderstigt)
- Fixed a bug in TRPO where kl divergence was not implemented for MultiDiscrete space

### 12.3.4 Deprecations:

## 12.4 Release 1.5.0 (2022-03-25)

### 12.4.1 Breaking Changes:

- Switched minimum Gym version to 0.21.0.
- Upgraded to Stable-Baselines3 >= 1.5.0

### 12.4.2 New Features:

- Allow PPO to turn of advantage normalization (see [PR #61](#)) (@vwxyzjn)

### 12.4.3 Bug Fixes:

- Removed explicit calls to `forward()` method as per pytorch guidelines

### 12.4.4 Deprecations:

### 12.4.5 Others:

### 12.4.6 Documentation:

## 12.5 Release 1.4.0 (2022-01-19)

**Add Trust Region Policy Optimization (TRPO) and Augmented Random Search (ARS) algorithms**

### 12.5.1 Breaking Changes:

- Dropped python 3.6 support
- Upgraded to Stable-Baselines3  $\geq 1.4.0$
- MaskablePPO was updated to match latest SB3 PPO version (timeout handling and new method for the policy object)

### 12.5.2 New Features:

- Added TRPO (@cyprienc)
- Added experimental support to train off-policy algorithms with multiple envs (note: HerReplayBuffer currently not supported)
- Added Augmented Random Search (ARS) (@sgillen)

### 12.5.3 Bug Fixes:

### 12.5.4 Deprecations:

### 12.5.5 Others:

- Improve test coverage for MaskablePPO

### 12.5.6 Documentation:

## 12.6 Release 1.3.0 (2021-10-23)

### Add Invalid action masking for PPO

<b>Warning:</b> This version will be the last one supporting Python 3.6 (end of life in Dec 2021). We highly recommended you to upgrade to Python $\geq 3.7$ .
--

### 12.6.1 Breaking Changes:

- Removed sde\_net\_arch
- Upgraded to Stable-Baselines3  $\geq 1.3.0$

### 12.6.2 New Features:

- Added MaskablePPO algorithm (@kronion)
- MaskablePPO Dictionary Observation support (@glmcdona)

### 12.6.3 Bug Fixes:

### 12.6.4 Deprecations:

### 12.6.5 Others:

### 12.6.6 Documentation:

## 12.7 Release 1.2.0 (2021-09-08)

### Train/Eval mode support

### 12.7.1 Breaking Changes:

- Upgraded to Stable-Baselines3 >= 1.2.0

### 12.7.2 Bug Fixes:

- QR-DQN and TQC updated so that their policies are switched between train and eval mode at the correct time (@ayeright)

### 12.7.3 Deprecations:

### 12.7.4 Others:

- Fixed type annotation
- Added python 3.9 to CI

### 12.7.5 Documentation:

## 12.8 Release 1.1.0 (2021-07-01)

### Dictionary observation support and timeout handling

### 12.8.1 Breaking Changes:

- Added support for Dictionary observation spaces (cf. SB3 doc)
- Upgraded to Stable-Baselines3  $\geq 1.1.0$
- Added proper handling of timeouts for off-policy algorithms (cf. SB3 doc)
- Updated usage of logger (cf. SB3 doc)

### 12.8.2 Bug Fixes:

- Removed unused code in TQC

### 12.8.3 Deprecations:

### 12.8.4 Others:

- SB3 docs and tests dependencies are no longer required for installing SB3 contrib

### 12.8.5 Documentation:

- updated QR-DQN docs checkmark typo (@minhlong94)

## 12.9 Release 1.0 (2021-03-17)

### 12.9.1 Breaking Changes:

- Upgraded to Stable-Baselines3  $\geq 1.0$

### 12.9.2 Bug Fixes:

- Fixed a bug with QR-DQN predict method when using `deterministic=False` with image space

## 12.10 Pre-Release 0.11.1 (2021-02-27)

### 12.10.1 Bug Fixes:

- Upgraded to Stable-Baselines3  $\geq 0.11.1$

## 12.11 Pre-Release 0.11.0 (2021-02-27)

### 12.11.1 Breaking Changes:

- Upgraded to Stable-Baselines3  $\geq$  0.11.0

### 12.11.2 New Features:

- Added TimeFeatureWrapper to the wrappers
- Added QR-DQN algorithm ([@ku2482](#))

### 12.11.3 Bug Fixes:

- Fixed bug in TQC when saving/loading the policy only with non-default number of quantiles
- Fixed bug in QR-DQN when calculating the target quantiles ([@ku2482](#), [@guyk1971](#))

### 12.11.4 Deprecations:

### 12.11.5 Others:

- Updated TQC to match new SB3 version
- Updated SB3 min version
- Moved quantile\_huber\_loss to common/utils.py ([@ku2482](#))

### 12.11.6 Documentation:

## 12.12 Pre-Release 0.10.0 (2020-10-28)

### Truncated Quantiles Critic (TQC)

### 12.12.1 Breaking Changes:

### 12.12.2 New Features:

- Added TQC algorithm ([@araffin](#))

### 12.12.3 Bug Fixes:

- Fixed features extractor issue (TQC with CnnPolicy)

### 12.12.4 Deprecations:

### 12.12.5 Others:

### 12.12.6 Documentation:

- Added initial documentation
- Added contribution guide and related PR templates

## 12.13 Maintainers

Stable-Baselines3 is currently maintained by [Antonin Raffin](#) (aka [@araffin](#)), [Ashley Hill](#) (aka [@hill-a](#)), [Maximilian Ernestus](#) (aka [@ernestum](#)), [Adam Gleave](#) ([@AdamGleave](#)) and [Anssi Kanervisto](#) (aka [@Miffyli](#)).

## 12.14 Contributors:

[@ku2482](#) [@guyk1971](#) [@minhlong94](#) [@ayeright](#) [@kronion](#) [@glmcdona](#) [@cyprienc](#) [@sgillen](#) [@Gregwar](#) [@rnderstigt](#)  
[@qgallouedec](#) [@mlodel](#) [@CppMaster](#) [@burakdmb](#) [@honglu2875](#)



## CITING STABLE BASELINES3

To cite this project in publications:

```
@misc{stable-baselines3,  
  author = {Raffin, Antonin and Hill, Ashley and Ernestus, Maximilian and Gleave, Adam,  
↵and Kanervisto, Anssi and Dormann, Noah},  
  title = {Stable Baselines3},  
  year = {2019},  
  publisher = {GitHub},  
  journal = {GitHub repository},  
  howpublished = {\url{https://github.com/DLR-RM/stable-baselines3}},  
}
```



## CONTRIBUTING

If you want to contribute, please read [CONTRIBUTING.md](#) first.



## INDICES AND TABLES

- `genindex`
- `search`
- `modindex`



## PYTHON MODULE INDEX

### S

- `sb3_contrib.ars`, [8](#)
- `sb3_contrib.common.utils`, [81](#)
- `sb3_contrib.common.wrappers`, [83](#)
- `sb3_contrib.ppo_mask`, [16](#)
- `sb3_contrib.ppo_recurrent`, [32](#)
- `sb3_contrib.qrdqn`, [46](#)
- `sb3_contrib.tqc`, [56](#)
- `sb3_contrib.trpo`, [67](#)





## A

ARS (class in *sb3\_contrib.ars*), 11

## C

CnnLstmPolicy (in module *sb3\_contrib.ppo\_recurrent*), 43

CnnPolicy (class in *sb3\_contrib.qrdqn*), 55

CnnPolicy (class in *sb3\_contrib.tqc*), 66

CnnPolicy (in module *sb3\_contrib.ppo\_mask*), 31

CnnPolicy (in module *sb3\_contrib.trpo*), 78

collect\_rollouts() (*sb3\_contrib.ppo\_mask.MaskablePPO* method), 25

collect\_rollouts() (*sb3\_contrib.ppo\_recurrent.RecurrentPPO* method), 36

collect\_rollouts() (*sb3\_contrib.qrdqn.QRDQN* method), 50

collect\_rollouts() (*sb3\_contrib.tqc.TQC* method), 61

collect\_rollouts() (*sb3\_contrib.trpo.TRPO* method), 72

conjugate\_gradient\_solver() (in module *sb3\_contrib.common.utils*), 81

## E

evaluate\_candidates() (*sb3\_contrib.ars.ARS* method), 12

## F

flat\_grad() (in module *sb3\_contrib.common.utils*), 81

## G

get\_env() (*sb3\_contrib.ars.ARS* method), 12

get\_env() (*sb3\_contrib.ppo\_mask.MaskablePPO* method), 25

get\_env() (*sb3\_contrib.ppo\_recurrent.RecurrentPPO* method), 36

get\_env() (*sb3\_contrib.qrdqn.QRDQN* method), 51

get\_env() (*sb3\_contrib.tqc.TQC* method), 61

get\_env() (*sb3\_contrib.trpo.TRPO* method), 72

get\_parameters() (*sb3\_contrib.ars.ARS* method), 12

get\_parameters() (*sb3\_contrib.ppo\_mask.MaskablePPO* method), 25

get\_parameters() (*sb3\_contrib.ppo\_recurrent.RecurrentPPO* method), 36

get\_parameters() (*sb3\_contrib.qrdqn.QRDQN* method), 51

get\_parameters() (*sb3\_contrib.tqc.TQC* method), 61

get\_parameters() (*sb3\_contrib.trpo.TRPO* method), 72

get\_vec\_normalize\_env() (*sb3\_contrib.ars.ARS* method), 12

get\_vec\_normalize\_env() (*sb3\_contrib.ppo\_mask.MaskablePPO* method), 25

get\_vec\_normalize\_env() (*sb3\_contrib.ppo\_recurrent.RecurrentPPO* method), 37

get\_vec\_normalize\_env() (*sb3\_contrib.qrdqn.QRDQN* method), 51

get\_vec\_normalize\_env() (*sb3\_contrib.tqc.TQC* method), 61

get\_vec\_normalize\_env() (*sb3\_contrib.trpo.TRPO* method), 73

## H

hessian\_vector\_product() (*sb3\_contrib.trpo.TRPO* method), 73

## L

learn() (*sb3\_contrib.ars.ARS* method), 13

learn() (*sb3\_contrib.ppo\_mask.MaskablePPO* method), 26

learn() (*sb3\_contrib.ppo\_recurrent.RecurrentPPO* method), 37

learn() (*sb3\_contrib.qrdqn.QRDQN* method), 51

learn() (*sb3\_contrib.tqc.TQC* method), 62

learn() (*sb3\_contrib.trpo.TRPO* method), 73

LinearPolicy (in module *sb3\_contrib.ars*), 16

load() (*sb3\_contrib.ars.ARS* class method), 13

load() (*sb3\_contrib.ppo\_mask.MaskablePPO* class method), 26

load() (*sb3\_contrib.ppo\_recurrent.RecurrentPPO* class method), 37

load() (*sb3\_contrib.qrdqn.QRDQN* class method), 52

load() (*sb3\_contrib.tqc.TQC class method*), 62  
 load() (*sb3\_contrib.trpo.TRPO class method*), 73  
 load\_replay\_buffer() (*sb3\_contrib.qrdqn.QRDQN method*), 52  
 load\_replay\_buffer() (*sb3\_contrib.tqc.TQC method*), 63  
 logger (*sb3\_contrib.ars.ARS property*), 14  
 logger (*sb3\_contrib.ppo\_mask.MaskablePPO property*), 27  
 logger (*sb3\_contrib.ppo\_recurrent.RecurrentPPO property*), 38  
 logger (*sb3\_contrib.qrdqn.QRDQN property*), 52  
 logger (*sb3\_contrib.tqc.TQC property*), 63  
 logger (*sb3\_contrib.trpo.TRPO property*), 74

## M

MaskablePPO (*class in sb3\_contrib.ppo\_mask*), 24  
 MlpLstmPolicy (*in module sb3\_contrib.ppo\_recurrent*), 40  
 MlpPolicy (*in module sb3\_contrib.ars*), 16  
 MlpPolicy (*in module sb3\_contrib.ppo\_mask*), 28  
 MlpPolicy (*in module sb3\_contrib.qrdqn*), 54  
 MlpPolicy (*in module sb3\_contrib.tqc*), 65  
 MlpPolicy (*in module sb3\_contrib.trpo*), 76  
 module  
   sb3\_contrib.ars, 8  
   sb3\_contrib.common.utils, 81  
   sb3\_contrib.common.wrappers, 83  
   sb3\_contrib.ppo\_mask, 16  
   sb3\_contrib.ppo\_recurrent, 32  
   sb3\_contrib.qrdqn, 46  
   sb3\_contrib.tqc, 56  
   sb3\_contrib.trpo, 67  
 MultiInputLstmPolicy (*in module sb3\_contrib.ppo\_recurrent*), 44  
 MultiInputPolicy (*class in sb3\_contrib.qrdqn*), 56  
 MultiInputPolicy (*in module sb3\_contrib.ppo\_mask*), 31  
 MultiInputPolicy (*in module sb3\_contrib.trpo*), 79

## P

predict() (*sb3\_contrib.ars.ARS method*), 14  
 predict() (*sb3\_contrib.ppo\_mask.MaskablePPO method*), 27  
 predict() (*sb3\_contrib.ppo\_recurrent.RecurrentPPO method*), 38  
 predict() (*sb3\_contrib.qrdqn.QRDQN method*), 52  
 predict() (*sb3\_contrib.tqc.TQC method*), 63  
 predict() (*sb3\_contrib.trpo.TRPO method*), 74

## Q

QRDQN (*class in sb3\_contrib.qrdqn*), 49  
 quantile\_huber\_loss() (*in module sb3\_contrib.common.utils*), 82

## R

RecurrentPPO (*class in sb3\_contrib.ppo\_recurrent*), 35  
 reset() (*sb3\_contrib.common.wrappers.TimeFeatureWrapper method*), 83

## S

save() (*sb3\_contrib.ars.ARS method*), 14  
 save() (*sb3\_contrib.ppo\_mask.MaskablePPO method*), 27  
 save() (*sb3\_contrib.ppo\_recurrent.RecurrentPPO method*), 38  
 save() (*sb3\_contrib.qrdqn.QRDQN method*), 53  
 save() (*sb3\_contrib.tqc.TQC method*), 63  
 save() (*sb3\_contrib.trpo.TRPO method*), 74  
 save\_replay\_buffer() (*sb3\_contrib.qrdqn.QRDQN method*), 53  
 save\_replay\_buffer() (*sb3\_contrib.tqc.TQC method*), 63  
 sb3\_contrib.ars  
   module, 8  
 sb3\_contrib.common.utils  
   module, 81  
 sb3\_contrib.common.wrappers  
   module, 83  
 sb3\_contrib.ppo\_mask  
   module, 16  
 sb3\_contrib.ppo\_recurrent  
   module, 32  
 sb3\_contrib.qrdqn  
   module, 46  
 sb3\_contrib.tqc  
   module, 56  
 sb3\_contrib.trpo  
   module, 67  
 set\_env() (*sb3\_contrib.ars.ARS method*), 14  
 set\_env() (*sb3\_contrib.ppo\_mask.MaskablePPO method*), 27  
 set\_env() (*sb3\_contrib.ppo\_recurrent.RecurrentPPO method*), 38  
 set\_env() (*sb3\_contrib.qrdqn.QRDQN method*), 53  
 set\_env() (*sb3\_contrib.tqc.TQC method*), 64  
 set\_env() (*sb3\_contrib.trpo.TRPO method*), 75  
 set\_logger() (*sb3\_contrib.ars.ARS method*), 15  
 set\_logger() (*sb3\_contrib.ppo\_mask.MaskablePPO method*), 27  
 set\_logger() (*sb3\_contrib.ppo\_recurrent.RecurrentPPO method*), 39  
 set\_logger() (*sb3\_contrib.qrdqn.QRDQN method*), 53  
 set\_logger() (*sb3\_contrib.tqc.TQC method*), 64  
 set\_logger() (*sb3\_contrib.trpo.TRPO method*), 75  
 set\_parameters() (*sb3\_contrib.ars.ARS method*), 15  
 set\_parameters() (*sb3\_contrib.ppo\_mask.MaskablePPO method*), 28

[set\\_parameters\(\) \(sb3\\_contrib.ppo\\_recurrent.RecurrentPPO method\), 39](#)  
[set\\_parameters\(\) \(sb3\\_contrib.qrdqn.QRDQN method\), 54](#)  
[set\\_parameters\(\) \(sb3\\_contrib.tqc.TQC method\), 64](#)  
[set\\_parameters\(\) \(sb3\\_contrib.trpo.TRPO method\), 75](#)  
[set\\_random\\_seed\(\) \(sb3\\_contrib.ars.ARS method\), 15](#)  
[set\\_random\\_seed\(\) \(sb3\\_contrib.ppo\\_mask.MaskablePPO method\), 28](#)  
[set\\_random\\_seed\(\) \(sb3\\_contrib.ppo\\_recurrent.RecurrentPPO method\), 39](#)  
[set\\_random\\_seed\(\) \(sb3\\_contrib.qrdqn.QRDQN method\), 54](#)  
[set\\_random\\_seed\(\) \(sb3\\_contrib.tqc.TQC method\), 64](#)  
[set\\_random\\_seed\(\) \(sb3\\_contrib.trpo.TRPO method\), 75](#)  
[step\(\) \(sb3\\_contrib.common.wrappers.TimeFeatureWrapper method\), 83](#)

## T

[TimeFeatureWrapper \(class in sb3\\_contrib.common.wrappers\), 83](#)  
[TQC \(class in sb3\\_contrib.tqc\), 59](#)  
[train\(\) \(sb3\\_contrib.ppo\\_mask.MaskablePPO method\), 28](#)  
[train\(\) \(sb3\\_contrib.ppo\\_recurrent.RecurrentPPO method\), 39](#)  
[train\(\) \(sb3\\_contrib.qrdqn.QRDQN method\), 54](#)  
[train\(\) \(sb3\\_contrib.tqc.TQC method\), 64](#)  
[train\(\) \(sb3\\_contrib.trpo.TRPO method\), 76](#)  
[TRPO \(class in sb3\\_contrib.trpo\), 71](#)