

---

# **MushroomRL Benchmark Documentation**

*Release 1.1.0*

**Carlo D'Eramo, Davide Tateo**

**Jun 10, 2021**



---

## Benchmarks Results:

---

<b>1</b>	<b>Reinforcement Learning python library</b>	<b>1</b>
<b>2</b>	<b>Download and installation</b>	<b>3</b>
<b>3</b>	<b>Benchmarks</b>	<b>5</b>
3.1	Policy Search Benchmarks . . . . .	5
3.2	Actor-Critic Benchmarks . . . . .	8
3.3	Value-Based Benchmarks . . . . .	40
3.4	Core functionality . . . . .	46
3.5	Builders . . . . .	55
3.6	Networks . . . . .	71
3.7	Experiment . . . . .	73
3.8	Utils . . . . .	75
	<b>Python Module Index</b>	<b>77</b>
	<b>Index</b>	<b>79</b>



---

## Reinforcement Learning python library

---

MushroomRL Benchmark is a benchmarking tool for the Mushroom RL library. The focus of this benchmarking tool is to benchmark the results of deep reinforcement learning algorithms, in particular Deep Actor-Critic. The idea behind MushroomRL Benchmarking is to have a complete platform to run batch comparisons of Deep RL algorithms implemented in MushroomRL under a set of standard benchmark tasks.

With MushroomRL Benchmarking you can:

- Run the benchmarks in a local machine, both sequentially and in parallel fashion
- Run experiments on a SLURM-based cluster.



## CHAPTER 2

---

### Download and installation

---

MushroomRL Benchmark can be downloaded from the [GitHub](#) repository. Installation can be done running

```
cd mushroom-rl-benchmark
pip install -e .[all]
```

To compile the documentation:

```
cd mushroom-rl-benchmark/docs
make html
```

or to compile the pdf version:

```
cd mushroom-rl-benchmark/docs
make latexpdf
```



### 3.1 Policy Search Benchmarks

We provide the benchmarks for the Policy Gradient algorithms:

- **REINFORCE**
- **GPOMDP**
- **eNAC**

We provide the benchmarks for the following Black-Box optimization algorithms:

- **RWR**
- **REPS**
- **PGPE**
- **ConstrainedREPS**

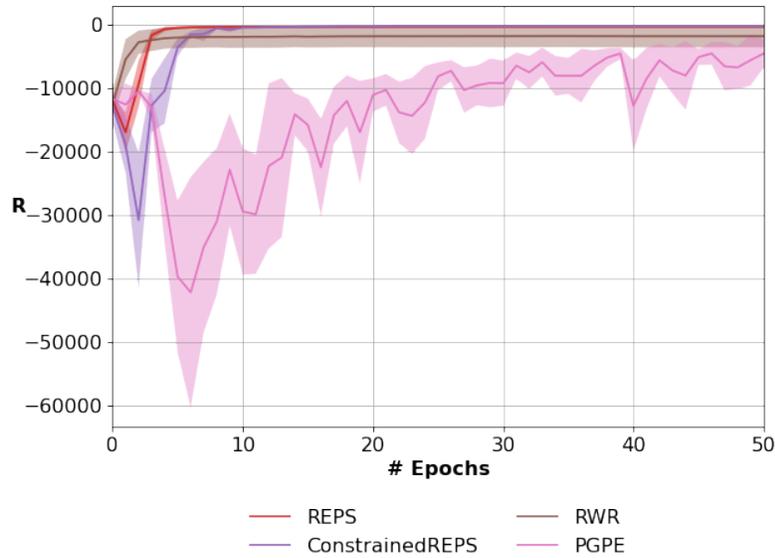
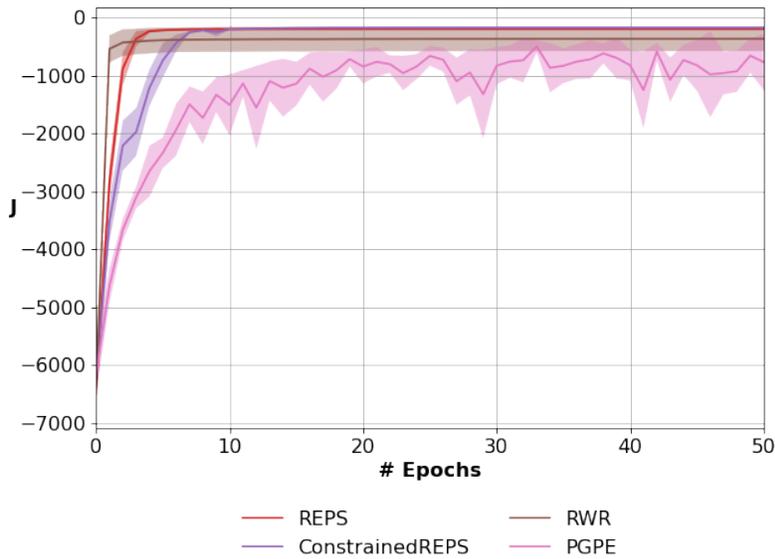
We consider the following environments in the benchmark

#### 3.1.1 Classic Control Environments Benchmarks

##### Segway

Run Parameters	
n_runs	25
n_epochs	50
n_episodes	100
n_episodes_test	10

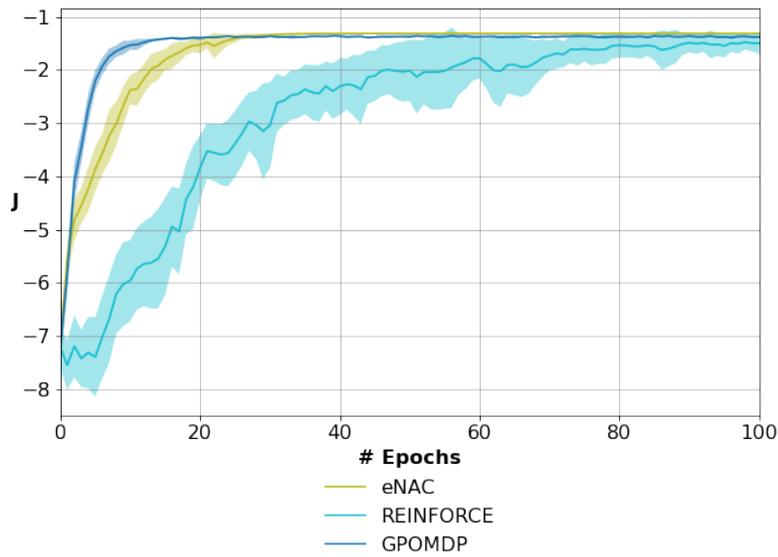
```
ConstrainedREPS:  
  eps: 0.5  
  kappa: 0.1  
  n_episodes_per_fit: 25  
PGPE:  
  alpha: 0.3  
  n_episodes_per_fit: 25  
REPS:  
  eps: 0.5  
  n_episodes_per_fit: 25  
RWR:  
  beta: 0.01  
  n_episodes_per_fit: 25
```

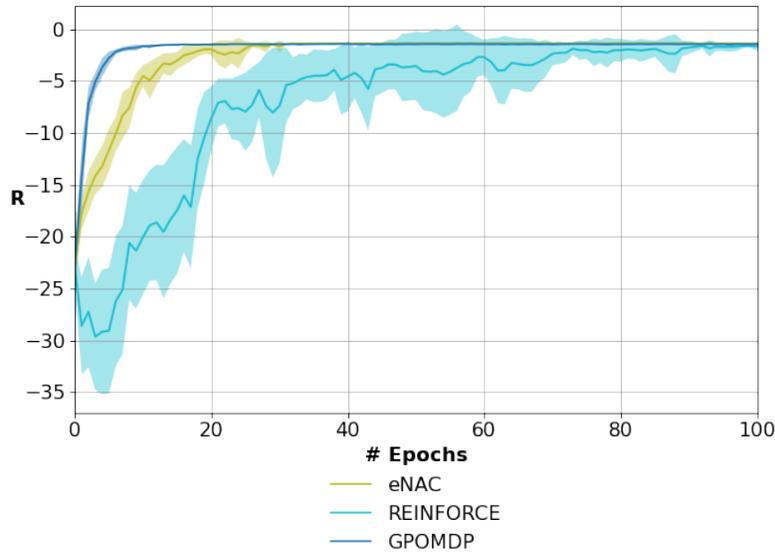


## LQR

Run Parameters	
n_runs	25
n_epochs	100
n_episodes	100
n_episodes_test	10

```
GPOMDP:  
  alpha: 0.01  
  n_episodes_per_fit: 25  
REINFORCE:  
  alpha: 0.01  
  n_episodes_per_fit: 25  
eNAC:  
  alpha: 0.01  
  n_episodes_per_fit: 25
```





## 3.2 Actor-Critic Benchmarks

We provide the benchmarks for the following Deep Actor-Critic algorithms:

- **StochasticAC**
- **COPDAC\_Q**

We provide the benchmarks for the following Deep Actor-Critic algorithms:

- **A2C**
- **PPO**
- **TRPO**
- **SAC**
- **DDPG**
- **TD3**

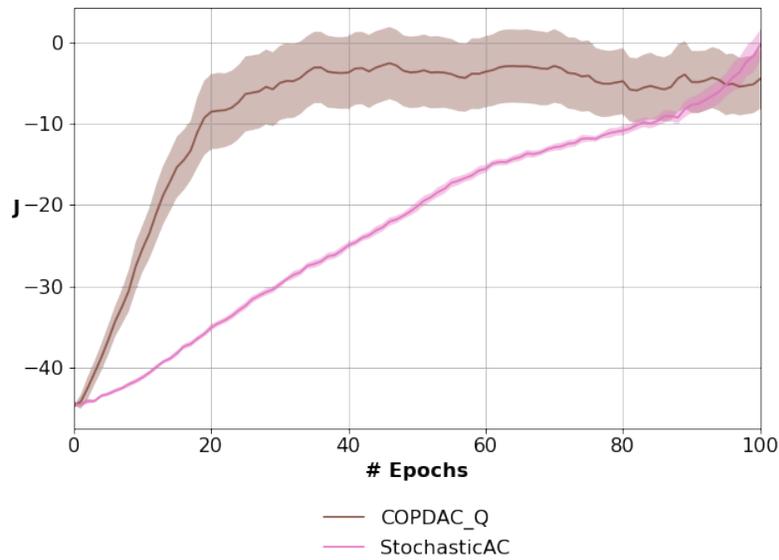
We consider the following environments in the benchmark

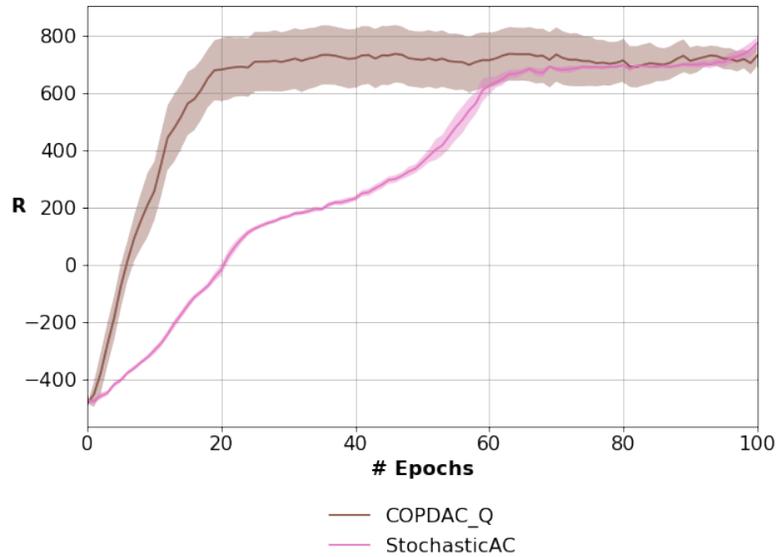
### 3.2.1 Classic Control Environments Benchmarks

Run Parameters	
n_runs	25
n_epochs	100
n_episodes	10
n_episodes_test	5

## InvertedPendulum

```
COPDAC_Q:
  alpha_omega: 0.5
  alpha_theta: 0.005
  alpha_v: 0.5
  n_tiles: 11
  n_tilings: 10
  std_eval: 0.001
  std_exp: 0.1
StochasticAC:
  alpha_theta: 0.001
  alpha_v: 0.1
  lambda_par: 0.9
  n_tiles: 11
  n_tilings: 10
  std_0: 1.0
```





### 3.2.2 Gym Control Environments Benchmarks

Run Parameters	
n_runs	25
n_epochs	10
n_steps	30000
n_episodes_test	10

#### Pendulum-v0

```

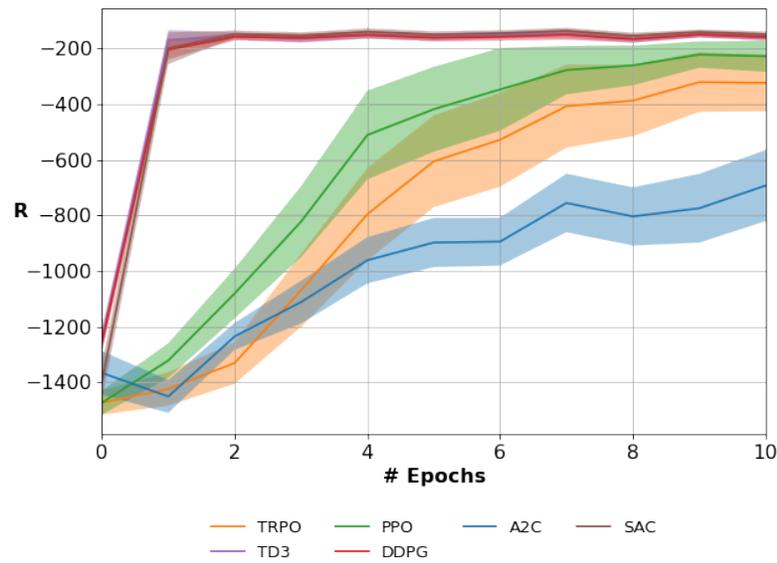
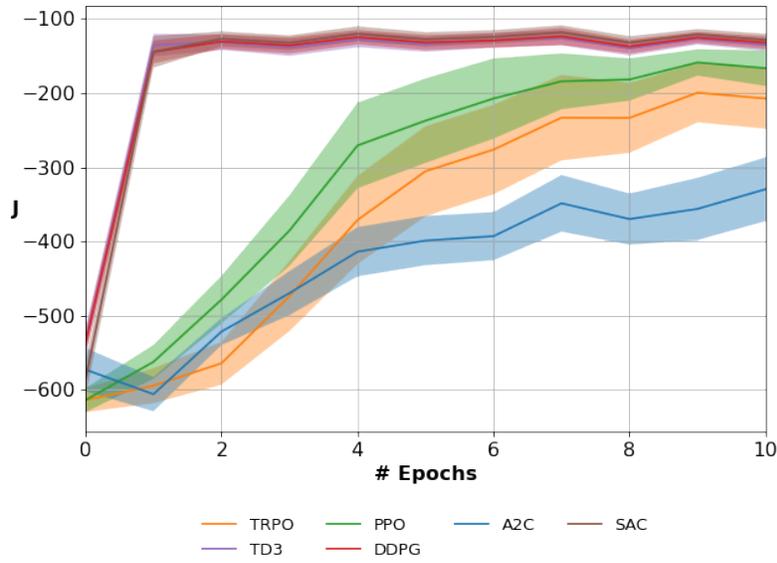
A2C:
  actor_lr: 0.0007
  batch_size: 64
  critic_lr: 0.0007
  critic_network: A2CNetwork
  ent_coeff: 0.01
  eps_actor: 0.003
  eps_critic: 1.0e-05
  max_grad_norm: 0.5
  n_features: 64
  preprocessors: null
DDPG:
  actor_lr: 0.0001
  actor_network: DDPGActorNetwork
  batch_size: 64
  critic_lr: 0.001
  critic_network: DDPGCriticNetwork
  initial_replay_size: 128
  max_replay_size: 1000000
  n_features:
    - 64
    - 64

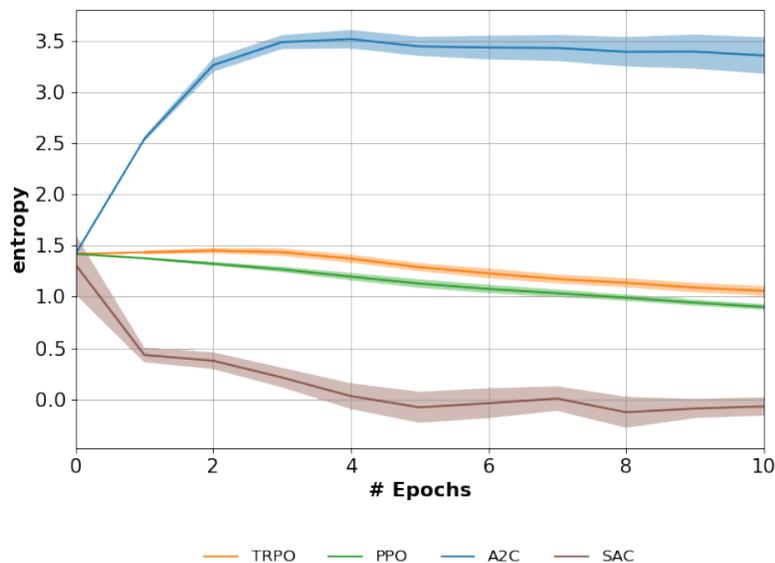
```

(continues on next page)

(continued from previous page)

```
tau: 0.001
PPO:
  actor_lr: 0.0003
  batch_size: 64
  critic_fit_params: null
  critic_lr: 0.0003
  critic_network: TRPONetwork
  eps: 0.2
  lam: 0.95
  n_epochs_policy: 4
  n_features: 32
  n_steps_per_fit: 3000
  preprocessors: null
SAC:
  actor_lr: 0.0001
  actor_network: SACActorNetwork
  batch_size: 64
  critic_lr: 0.0003
  critic_network: SACCriticNetwork
  initial_replay_size: 128
  lr_alpha: 0.0003
  max_replay_size: 500000
  n_features: 64
  preprocessors: null
  target_entropy: null
  tau: 0.001
  warmup_transitions: 128
TD3:
  actor_lr: 0.0001
  actor_network: TD3ActorNetwork
  batch_size: 64
  critic_lr: 0.001
  critic_network: TD3CriticNetwork
  initial_replay_size: 128
  max_replay_size: 1000000
  n_features:
    - 64
    - 64
  tau: 0.001
TRPO:
  batch_size: 64
  cg_damping: 0.01
  cg_residual_tol: 1.0e-10
  critic_fit_params: null
  critic_lr: 0.0003
  critic_network: TRPONetwork
  ent_coeff: 0.0
  lam: 0.95
  max_kl: 0.01
  n_epochs_cg: 100
  n_epochs_line_search: 10
  n_features: 32
  n_steps_per_fit: 3000
  preprocessors: null
```





### LunarLanderContinuous-v2

#### A2C:

```

actor_lr: 0.0007
batch_size: 64
critic_lr: 0.0007
critic_network: A2CNetwork
ent_coeff: 0.01
eps_actor: 0.003
eps_critic: 1.0e-05
max_grad_norm: 0.5
n_features: 64
preprocessors: StandardizationPreprocessor

```

#### DDPG:

```

actor_lr: 0.0001
actor_network: DDPGActorNetwork
batch_size: 64
critic_lr: 0.001
critic_network: DDPGCriticNetwork
initial_replay_size: 128
max_replay_size: 1000000
n_features:
- 64
- 64
tau: 0.001

```

#### PPO:

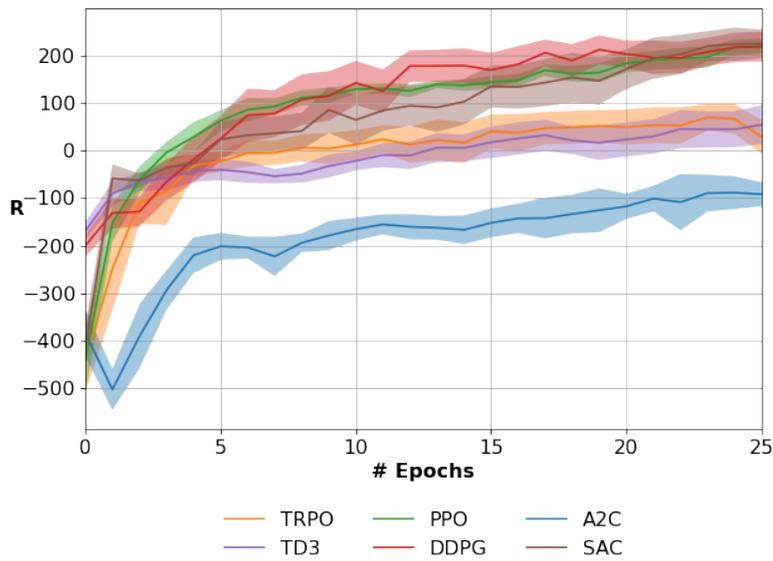
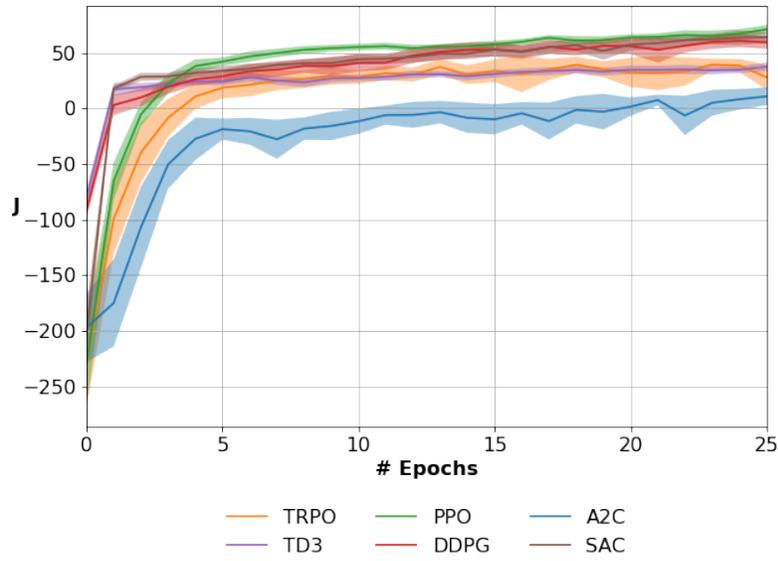
```

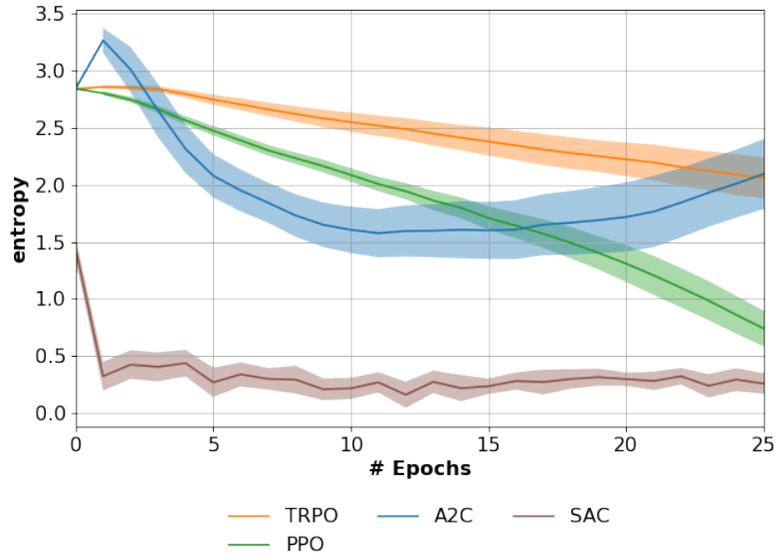
actor_lr: 0.0003
batch_size: 64
critic_fit_params: null
critic_lr: 0.003
critic_network: TRPONetwork
eps: 0.2
lam: 0.95
n_epochs_policy: 4
n_features: 32

```

(continues on next page)

```
n_steps_per_fit: 3000
preprocessors: StandardizationPreprocessor
SAC:
  actor_lr: 0.0001
  actor_network: SACActorNetwork
  batch_size: 64
  critic_lr: 0.0003
  critic_network: SACCriticNetwork
  initial_replay_size: 128
  lr_alpha: 0.0003
  max_replay_size: 500000
  n_features: 64
  preprocessors: null
  target_entropy: null
  tau: 0.005
  warmup_transitions: 128
TD3:
  actor_lr: 0.0001
  actor_network: TD3ActorNetwork
  batch_size: 64
  critic_lr: 0.001
  critic_network: TD3CriticNetwork
  initial_replay_size: 128
  max_replay_size: 1000000
  n_features:
  - 64
  - 64
  tau: 0.001
TRPO:
  batch_size: 64
  cg_damping: 0.01
  cg_residual_tol: 1.0e-10
  critic_fit_params: null
  critic_lr: 0.03
  critic_network: TRPONetwork
  ent_coeff: 0.0
  lam: 0.95
  max_kl: 0.01
  n_epochs_cg: 100
  n_epochs_line_search: 10
  n_features: 32
  n_steps_per_fit: 3000
  preprocessors: StandardizationPreprocessor
```





### 3.2.3 Mujoco Environments Benchmarks

Run Parameters	
n_runs	25
n_epochs	50
n_steps	30000
n_episodes_test	10

#### Hopper-v3

```

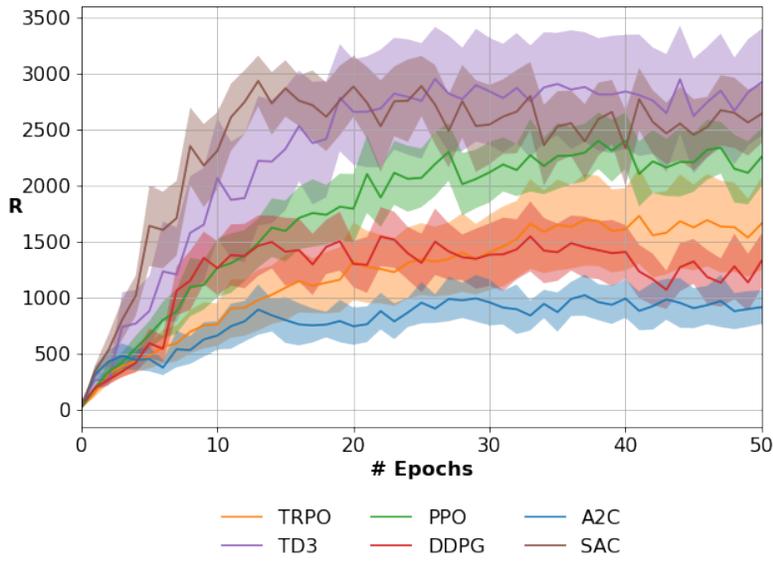
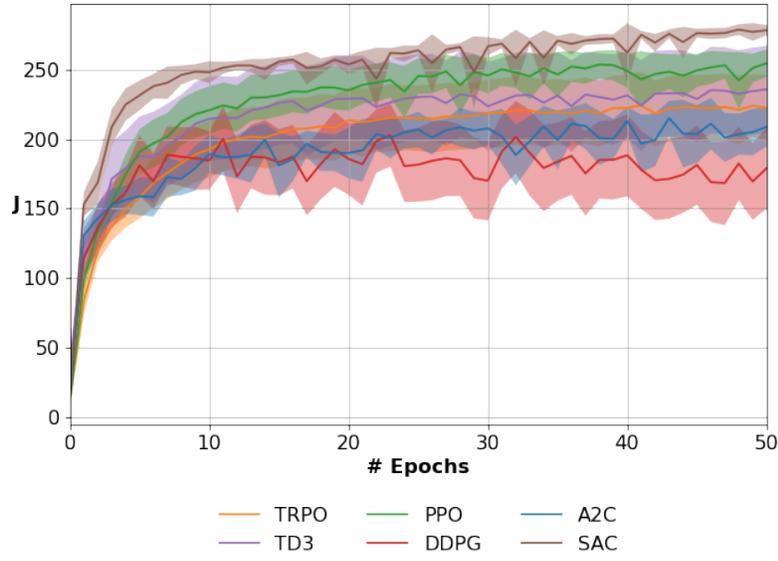
A2C:
  actor_lr: 0.0007
  batch_size: 64
  critic_lr: 0.0007
  critic_network: A2CNetwork
  ent_coeff: 0.01
  eps_actor: 0.003
  eps_critic: 1.0e-05
  max_grad_norm: 0.5
  n_features: 64
  preprocessors: StandardizationPreprocessor
DDPG:
  actor_lr: 0.0001
  actor_network: DDPGActorNetwork
  batch_size: 128
  critic_lr: 0.001
  critic_network: DDPGCriticNetwork
  initial_replay_size: 5000
  max_replay_size: 1000000
  n_features:
    - 400
    - 300

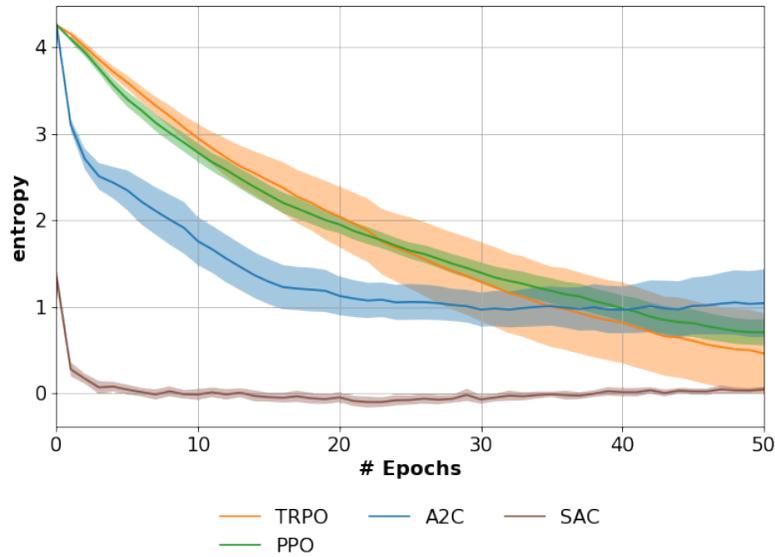
```

(continues on next page)

(continued from previous page)

```
tau: 0.001
PPO:
  actor_lr: 0.0003
  batch_size: 32
  critic_fit_params: null
  critic_lr: 0.0003
  critic_network: TRPONetwork
  eps: 0.2
  lam: 0.95
  n_epochs_policy: 10
  n_features: 32
  n_steps_per_fit: 2000
  preprocessors: StandardizationPreprocessor
SAC:
  actor_lr: 0.0001
  actor_network: SACActorNetwork
  batch_size: 256
  critic_lr: 0.0003
  critic_network: SACCriticNetwork
  initial_replay_size: 5000
  lr_alpha: 0.0003
  max_replay_size: 500000
  n_features: 256
  preprocessors: null
  target_entropy: null
  tau: 0.005
  warmup_transitions: 10000
TD3:
  actor_lr: 0.001
  actor_network: TD3ActorNetwork
  batch_size: 100
  critic_lr: 0.001
  critic_network: TD3CriticNetwork
  initial_replay_size: 1000
  max_replay_size: 1000000
  n_features:
  - 400
  - 300
  tau: 0.005
TRPO:
  batch_size: 64
  cg_damping: 0.01
  cg_residual_tol: 1.0e-10
  critic_fit_params: null
  critic_lr: 0.001
  critic_network: TRPONetwork
  ent_coeff: 0.0
  lam: 0.95
  max_kl: 0.01
  n_epochs_cg: 100
  n_epochs_line_search: 10
  n_features: 32
  n_steps_per_fit: 1000
  preprocessors: StandardizationPreprocessor
```





### Walker2d-v3

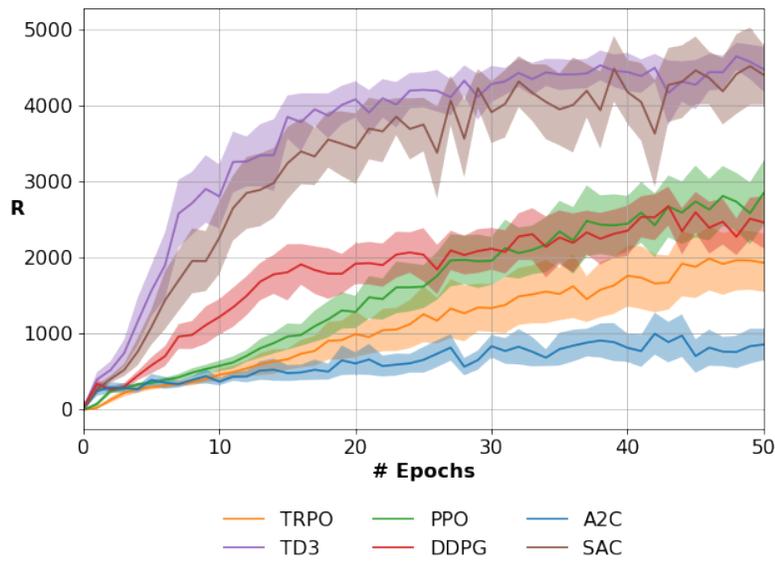
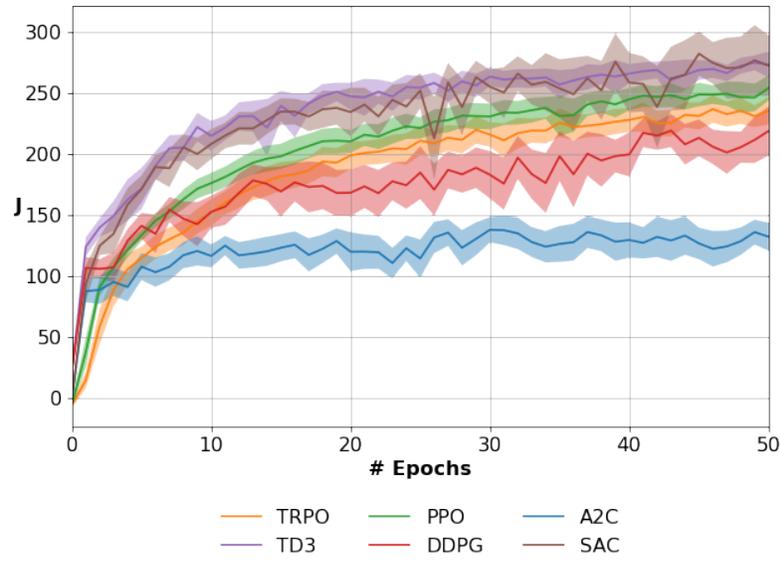
```

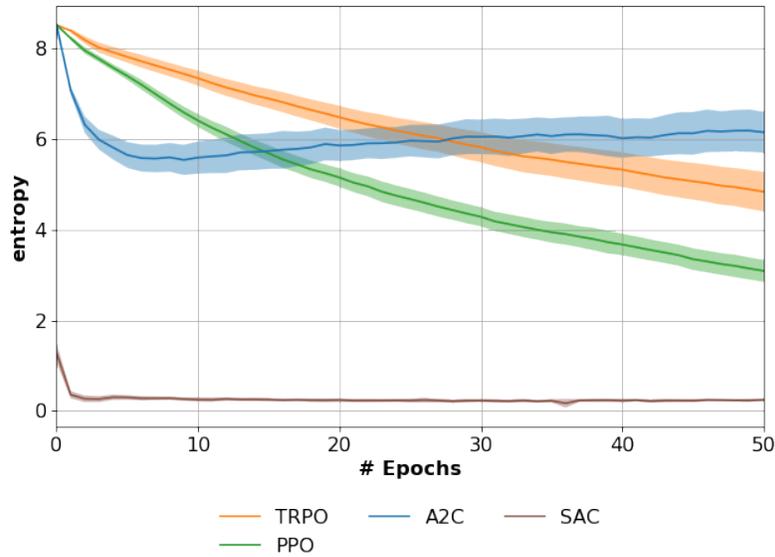
A2C:
  actor_lr: 0.0007
  batch_size: 64
  critic_lr: 0.0007
  critic_network: A2CNetwork
  ent_coeff: 0.01
  eps_actor: 0.003
  eps_critic: 1.0e-05
  max_grad_norm: 0.5
  n_features: 64
  preprocessors: StandardizationPreprocessor
DDPG:
  actor_lr: 0.0001
  actor_network: DDPGActorNetwork
  batch_size: 128
  critic_lr: 0.001
  critic_network: DDPGCriticNetwork
  initial_replay_size: 5000
  max_replay_size: 1000000
  n_features:
  - 400
  - 300
  tau: 0.001
PPO:
  actor_lr: 0.0003
  batch_size: 32
  critic_fit_params: null
  critic_lr: 0.0003
  critic_network: TRPONetwork
  eps: 0.2
  lam: 0.95
  n_epochs_policy: 10
  n_features: 32

```

(continues on next page)

```
n_steps_per_fit: 2000
preprocessors: StandardizationPreprocessor
SAC:
  actor_lr: 0.0001
  actor_network: SACActorNetwork
  batch_size: 256
  critic_lr: 0.0003
  critic_network: SACCriticNetwork
  initial_replay_size: 5000
  lr_alpha: 0.0003
  max_replay_size: 500000
  n_features: 256
  preprocessors: null
  target_entropy: null
  tau: 0.005
  warmup_transitions: 10000
TD3:
  actor_lr: 0.001
  actor_network: TD3ActorNetwork
  batch_size: 100
  critic_lr: 0.001
  critic_network: TD3CriticNetwork
  initial_replay_size: 1000
  max_replay_size: 1000000
  n_features:
    - 400
    - 300
  tau: 0.005
TRPO:
  batch_size: 64
  cg_damping: 0.01
  cg_residual_tol: 1.0e-10
  critic_fit_params: null
  critic_lr: 0.001
  critic_network: TRPONetwork
  ent_coeff: 0.0
  lam: 0.95
  max_kl: 0.01
  n_epochs_cg: 100
  n_epochs_line_search: 10
  n_features: 32
  n_steps_per_fit: 1000
  preprocessors: StandardizationPreprocessor
```





### HalfCheetah-v3

```

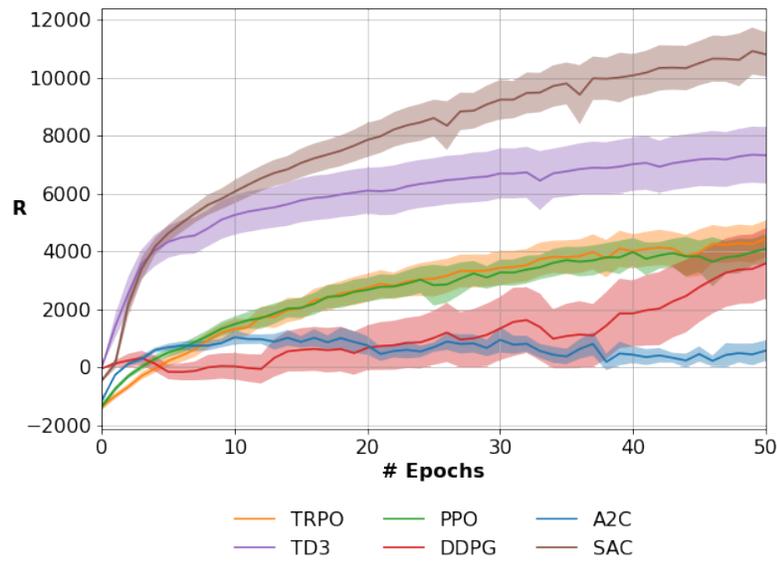
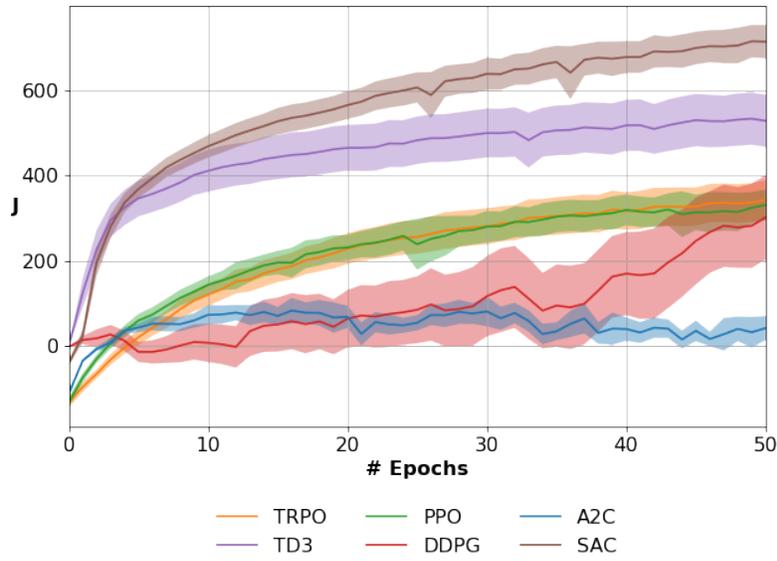
A2C:
  actor_lr: 0.0007
  batch_size: 64
  critic_lr: 0.0007
  critic_network: A2CNetwork
  ent_coeff: 0.01
  eps_actor: 0.003
  eps_critic: 1.0e-05
  max_grad_norm: 0.5
  n_features: 64
  preprocessors: StandardizationPreprocessor
DDPG:
  actor_lr: 0.0001
  actor_network: DDPGActorNetwork
  batch_size: 128
  critic_lr: 0.001
  critic_network: DDPGCriticNetwork
  initial_replay_size: 5000
  max_replay_size: 1000000
  n_features:
  - 400
  - 300
  tau: 0.001
PPO:
  actor_lr: 0.0003
  batch_size: 32
  critic_fit_params: null
  critic_lr: 0.0003
  critic_network: TRPONetwork
  eps: 0.2
  lam: 0.95
  n_epochs_policy: 10
  n_features: 32

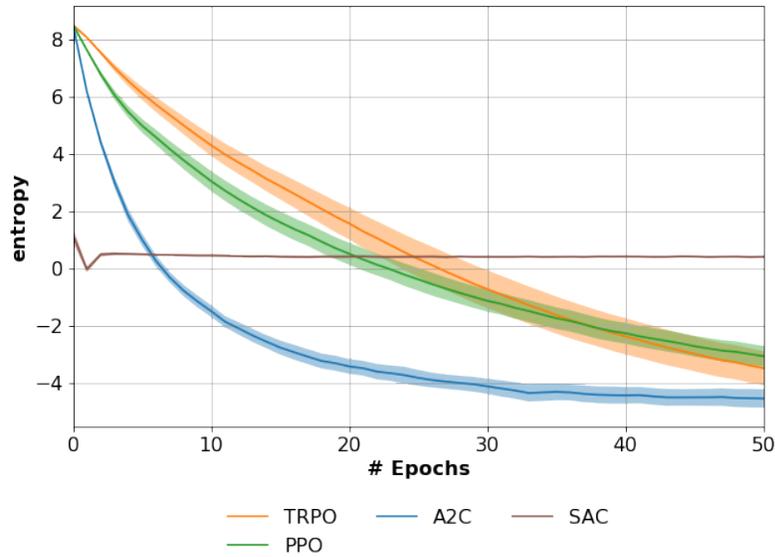
```

(continues on next page)

(continued from previous page)

```
n_steps_per_fit: 2000
preprocessors: StandardizationPreprocessor
SAC:
  actor_lr: 0.0001
  actor_network: SACActorNetwork
  batch_size: 256
  critic_lr: 0.0003
  critic_network: SACCriticNetwork
  initial_replay_size: 5000
  lr_alpha: 0.0003
  max_replay_size: 500000
  n_features: 256
  preprocessors: null
  target_entropy: null
  tau: 0.005
  warmup_transitions: 10000
TD3:
  actor_lr: 0.001
  actor_network: TD3ActorNetwork
  batch_size: 100
  critic_lr: 0.001
  critic_network: TD3CriticNetwork
  initial_replay_size: 10000
  max_replay_size: 1000000
  n_features:
    - 400
    - 300
  tau: 0.005
TRPO:
  batch_size: 64
  cg_damping: 0.01
  cg_residual_tol: 1.0e-10
  critic_fit_params: null
  critic_lr: 0.001
  critic_network: TRPONetwork
  ent_coeff: 0.0
  lam: 0.95
  max_kl: 0.01
  n_epochs_cg: 100
  n_epochs_line_search: 10
  n_features: 32
  n_steps_per_fit: 1000
  preprocessors: StandardizationPreprocessor
```





### Ant-v3

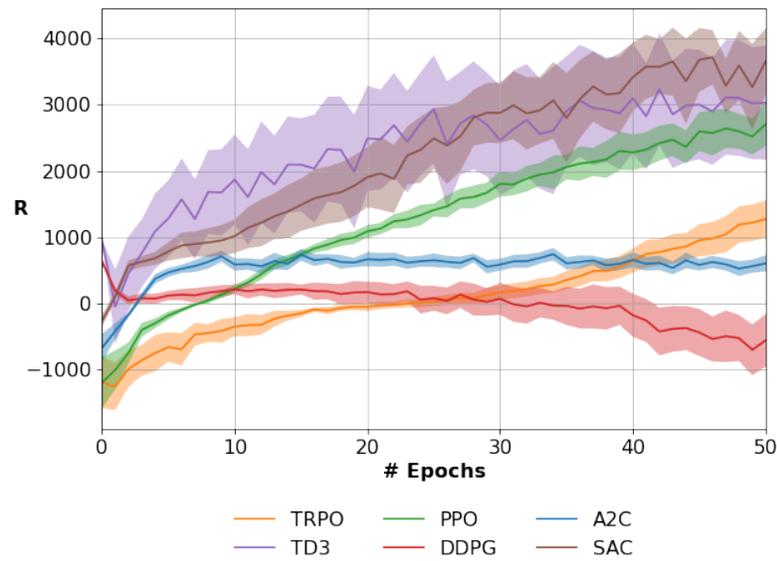
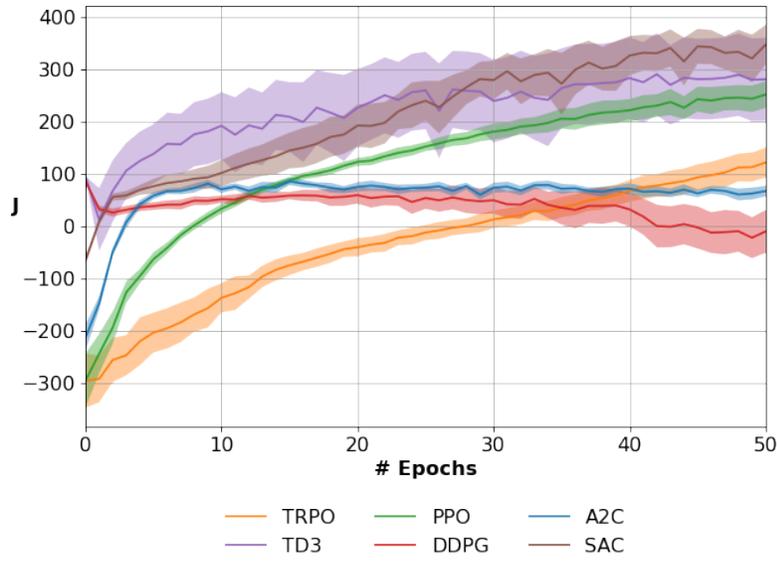
```

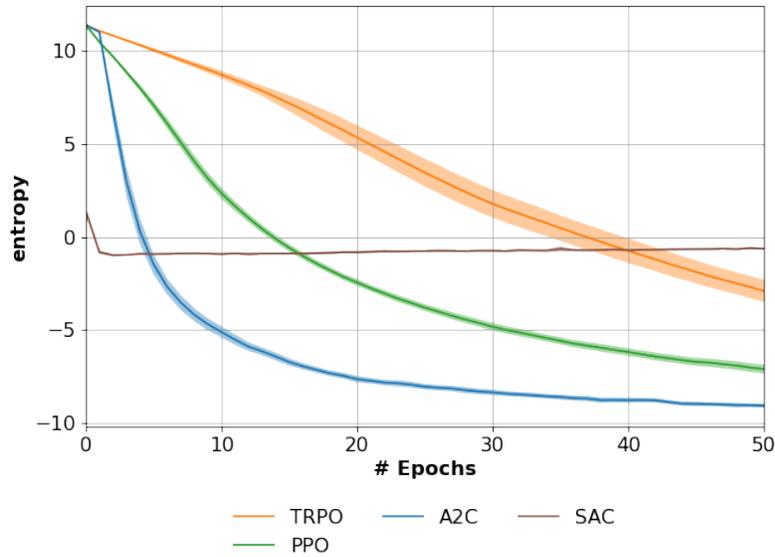
A2C:
  actor_lr: 0.0007
  batch_size: 64
  critic_lr: 0.0007
  critic_network: A2CNetwork
  ent_coeff: 0.01
  eps_actor: 0.003
  eps_critic: 1.0e-05
  max_grad_norm: 0.5
  n_features: 64
  preprocessors: StandardizationPreprocessor
DDPG:
  actor_lr: 0.0001
  actor_network: DDPGActorNetwork
  batch_size: 128
  critic_lr: 0.001
  critic_network: DDPGCriticNetwork
  initial_replay_size: 5000
  max_replay_size: 1000000
  n_features:
  - 400
  - 300
  tau: 0.001
PPO:
  actor_lr: 0.0003
  batch_size: 32
  critic_fit_params: null
  critic_lr: 0.0003
  critic_network: TRPONetwork
  eps: 0.2
  lam: 0.95
  n_epochs_policy: 10
  n_features: 32

```

(continues on next page)

```
n_steps_per_fit: 2000
preprocessors: StandardizationPreprocessor
SAC:
  actor_lr: 0.0001
  actor_network: SACActorNetwork
  batch_size: 256
  critic_lr: 0.0003
  critic_network: SACCriticNetwork
  initial_replay_size: 5000
  lr_alpha: 0.0003
  max_replay_size: 500000
  n_features: 256
  preprocessors: null
  target_entropy: null
  tau: 0.005
  warmup_transitions: 10000
TD3:
  actor_lr: 0.001
  actor_network: TD3ActorNetwork
  batch_size: 100
  critic_lr: 0.001
  critic_network: TD3CriticNetwork
  initial_replay_size: 10000
  max_replay_size: 1000000
  n_features:
    - 400
    - 300
  tau: 0.005
TRPO:
  batch_size: 64
  cg_damping: 0.01
  cg_residual_tol: 1.0e-10
  critic_fit_params: null
  critic_lr: 0.001
  critic_network: TRPONetwork
  ent_coeff: 0.0
  lam: 0.95
  max_kl: 0.01
  n_epochs_cg: 100
  n_epochs_line_search: 10
  n_features: 32
  n_steps_per_fit: 1000
  preprocessors: StandardizationPreprocessor
```





### 3.2.4 Bullet Environments Benchmarks

Run Parameters	
n_runs	25
n_epochs	50
n_steps	30000
n_episodes_test	10

#### HopperBulletEnv-v0

##### A2C:

```

actor_lr: 0.0007
batch_size: 64
critic_lr: 0.0007
critic_network: A2CNetwork
ent_coeff: 0.01
eps_actor: 0.003
eps_critic: 1.0e-05
max_grad_norm: 0.5
n_features: 64
preprocessors: null

```

##### DDPG:

```

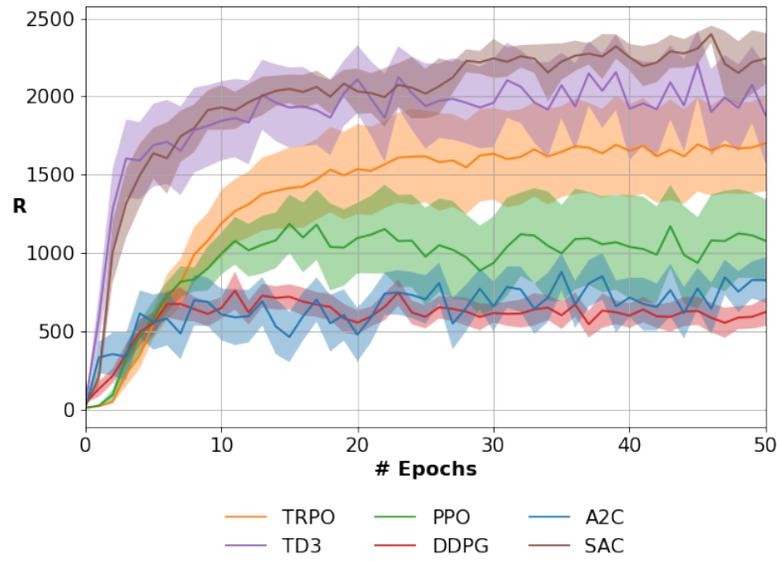
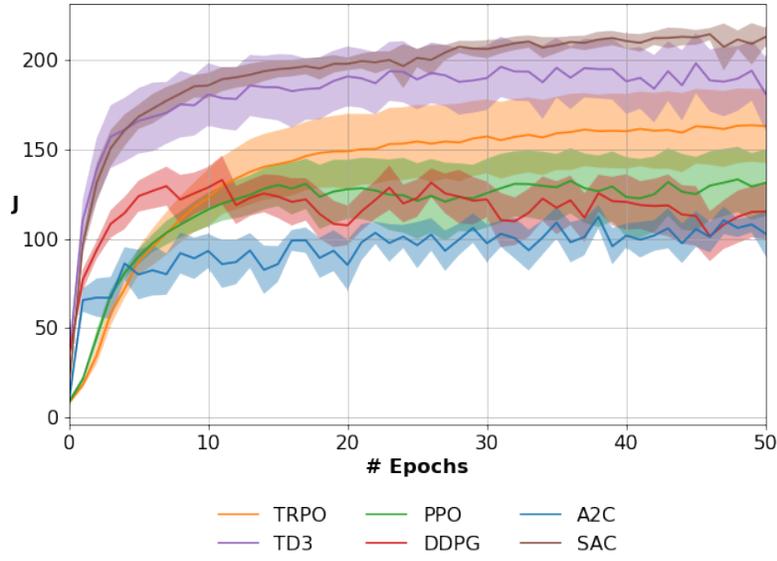
actor_lr: 0.0001
actor_network: DDPGActorNetwork
batch_size: 128
critic_lr: 0.001
critic_network: DDPGCriticNetwork
initial_replay_size: 5000
max_replay_size: 1000000
n_features:
- 400
- 300

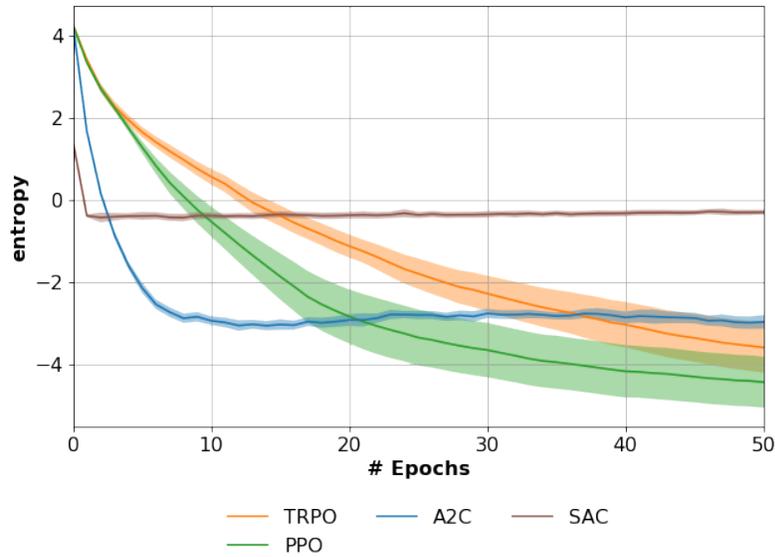
```

(continues on next page)

(continued from previous page)

```
tau: 0.001
PPO:
  actor_lr: 0.0003
  batch_size: 64
  critic_fit_params: null
  critic_lr: 0.0003
  critic_network: TRPONetwork
  eps: 0.2
  lam: 0.95
  n_epochs_policy: 4
  n_features: 32
  n_steps_per_fit: 3000
  preprocessors: null
SAC:
  actor_lr: 0.0001
  actor_network: SACActorNetwork
  batch_size: 256
  critic_lr: 0.0003
  critic_network: SACCriticNetwork
  initial_replay_size: 5000
  lr_alpha: 0.0003
  max_replay_size: 500000
  n_features: 256
  preprocessors: null
  target_entropy: null
  tau: 0.005
  warmup_transitions: 10000
TD3:
  actor_lr: 0.001
  actor_network: TD3ActorNetwork
  batch_size: 100
  critic_lr: 0.001
  critic_network: TD3CriticNetwork
  initial_replay_size: 1000
  max_replay_size: 1000000
  n_features:
  - 400
  - 300
  tau: 0.005
TRPO:
  batch_size: 64
  cg_damping: 0.01
  cg_residual_tol: 1.0e-10
  critic_fit_params: null
  critic_lr: 0.003
  critic_network: TRPONetwork
  ent_coeff: 0.0
  lam: 0.95
  max_kl: 0.01
  n_epochs_cg: 100
  n_epochs_line_search: 10
  n_features: 32
  n_steps_per_fit: 3000
  preprocessors: null
```





### Walker2DBulletEnv-v0

#### A2C:

```

actor_lr: 0.0007
batch_size: 64
critic_lr: 0.0007
critic_network: A2CNetwork
ent_coeff: 0.01
eps_actor: 0.003
eps_critic: 1.0e-05
max_grad_norm: 0.5
n_features: 64
preprocessors: null

```

#### DDPG:

```

actor_lr: 0.0001
actor_network: DDPGActorNetwork
batch_size: 128
critic_lr: 0.001
critic_network: DDPGCriticNetwork
initial_replay_size: 5000
max_replay_size: 1000000
n_features:
- 400
- 300
tau: 0.001

```

#### PPO:

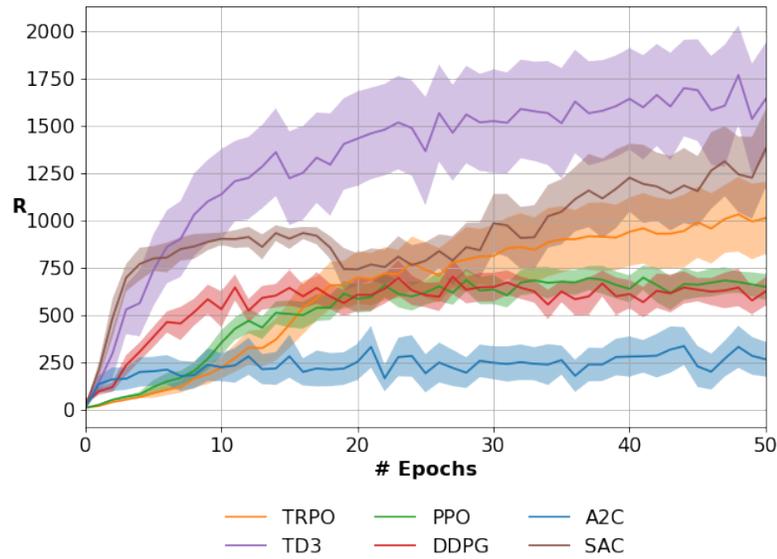
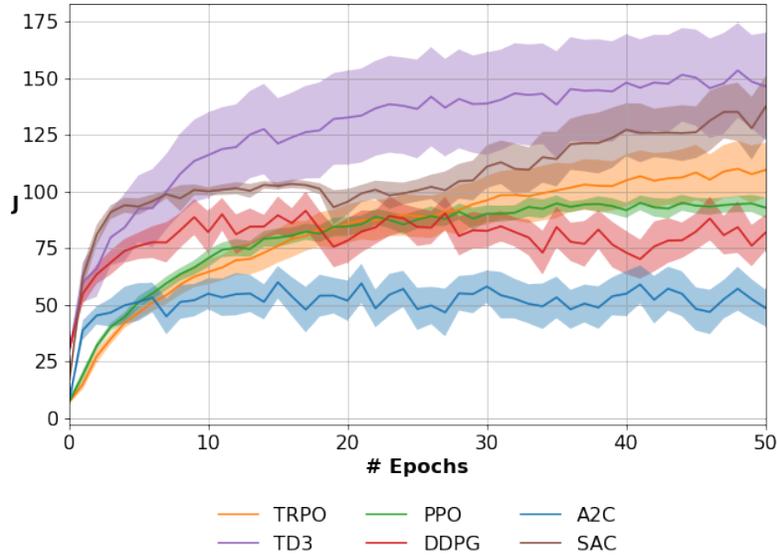
```

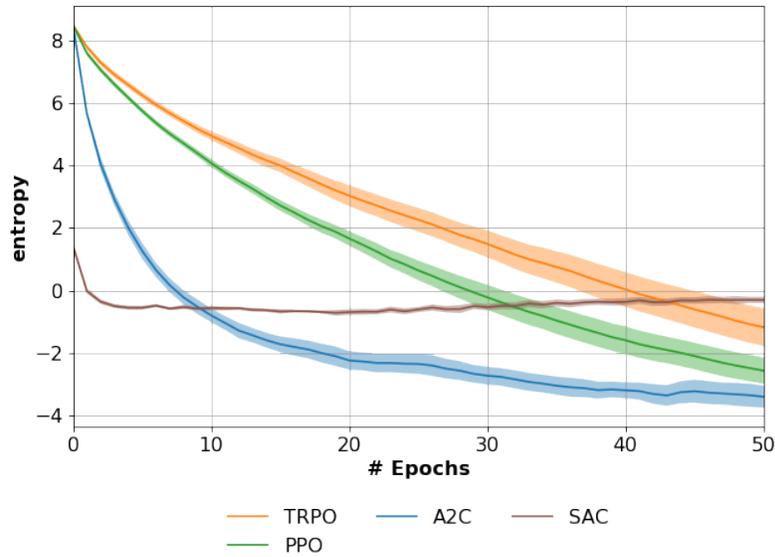
actor_lr: 0.0003
batch_size: 64
critic_fit_params: null
critic_lr: 0.0003
critic_network: TRPONetwork
eps: 0.2
lam: 0.95
n_epochs_policy: 4
n_features: 32

```

(continues on next page)

```
n_steps_per_fit: 3000
preprocessors: null
SAC:
  actor_lr: 0.0001
  actor_network: SACActorNetwork
  batch_size: 256
  critic_lr: 0.0003
  critic_network: SACCriticNetwork
  initial_replay_size: 5000
  lr_alpha: 0.0003
  max_replay_size: 500000
  n_features: 256
  preprocessors: null
  target_entropy: null
  tau: 0.005
  warmup_transitions: 10000
TD3:
  actor_lr: 0.001
  actor_network: TD3ActorNetwork
  batch_size: 100
  critic_lr: 0.001
  critic_network: TD3CriticNetwork
  initial_replay_size: 1000
  max_replay_size: 1000000
  n_features:
    - 400
    - 300
  tau: 0.005
TRPO:
  batch_size: 64
  cg_damping: 0.01
  cg_residual_tol: 1.0e-10
  critic_fit_params: null
  critic_lr: 0.003
  critic_network: TRPONetwork
  ent_coeff: 0.0
  lam: 0.95
  max_kl: 0.01
  n_epochs_cg: 100
  n_epochs_line_search: 10
  n_features: 32
  n_steps_per_fit: 3000
  preprocessors: null
```





### HalfCheetahBulletEnv-v0

#### A2C:

```

actor_lr: 0.0007
batch_size: 64
critic_lr: 0.0007
critic_network: A2CNetwork
ent_coeff: 0.01
eps_actor: 0.003
eps_critic: 1.0e-05
max_grad_norm: 0.5
n_features: 64
preprocessors: null

```

#### DDPG:

```

actor_lr: 0.0001
actor_network: DDPGActorNetwork
batch_size: 128
critic_lr: 0.001
critic_network: DDPGCriticNetwork
initial_replay_size: 5000
max_replay_size: 1000000
n_features:
- 400
- 300
tau: 0.001

```

#### PPO:

```

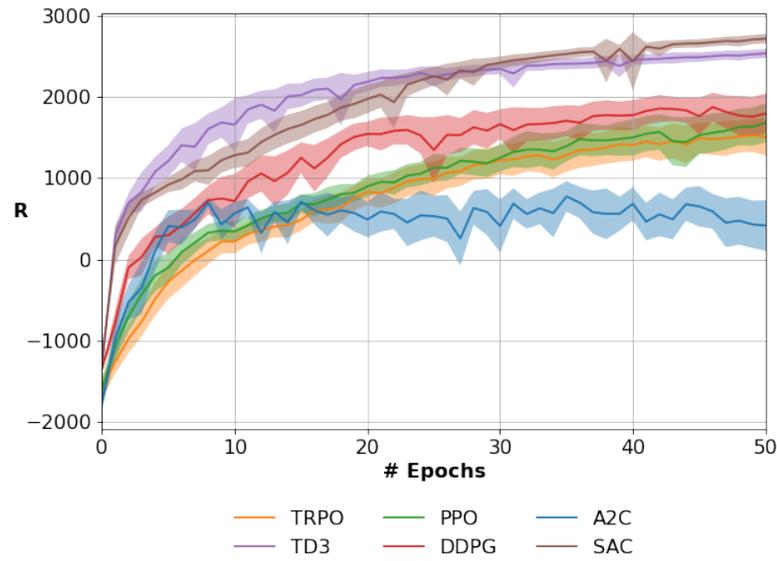
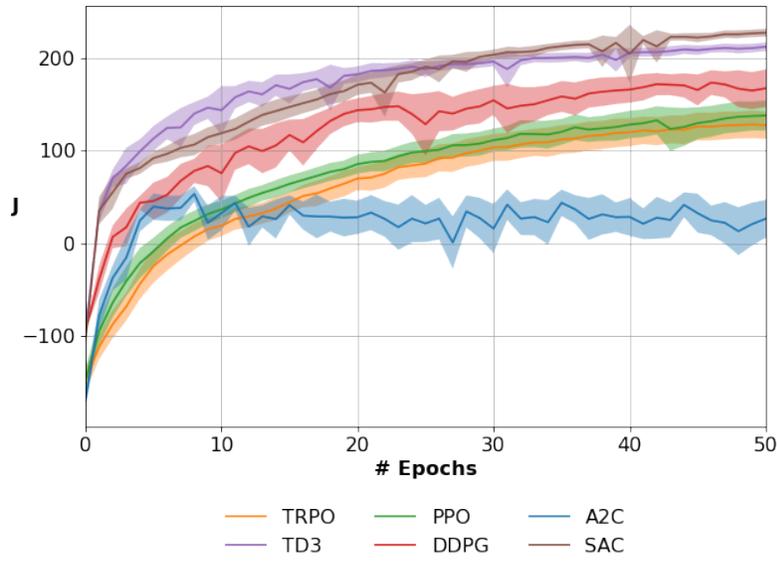
actor_lr: 0.0003
batch_size: 64
critic_fit_params: null
critic_lr: 0.0003
critic_network: TRPONetwork
eps: 0.2
lam: 0.95
n_epochs_policy: 4
n_features: 32

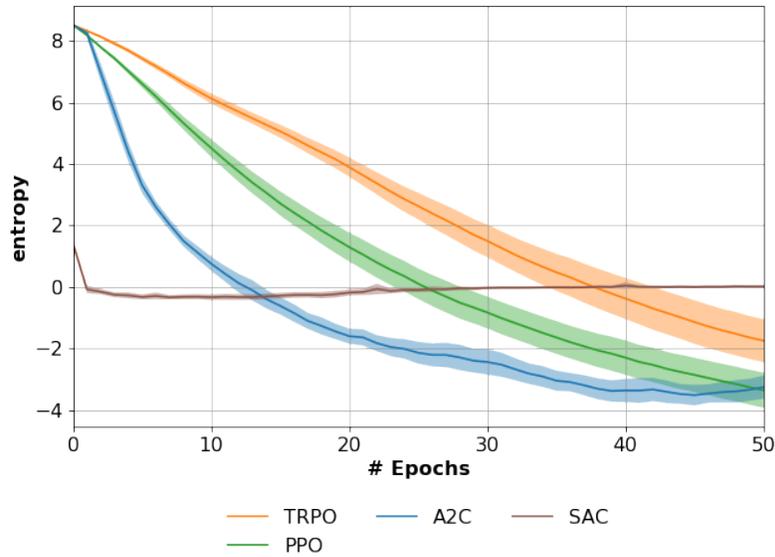
```

(continues on next page)

(continued from previous page)

```
n_steps_per_fit: 3000
preprocessors: null
SAC:
  actor_lr: 0.0001
  actor_network: SACActorNetwork
  batch_size: 256
  critic_lr: 0.0003
  critic_network: SACCriticNetwork
  initial_replay_size: 5000
  lr_alpha: 0.0003
  max_replay_size: 500000
  n_features: 256
  preprocessors: null
  target_entropy: null
  tau: 0.005
  warmup_transitions: 10000
TD3:
  actor_lr: 0.001
  actor_network: TD3ActorNetwork
  batch_size: 100
  critic_lr: 0.001
  critic_network: TD3CriticNetwork
  initial_replay_size: 10000
  max_replay_size: 1000000
  n_features:
    - 400
    - 300
  tau: 0.005
TRPO:
  batch_size: 64
  cg_damping: 0.01
  cg_residual_tol: 1.0e-10
  critic_fit_params: null
  critic_lr: 0.003
  critic_network: TRPONetwork
  ent_coeff: 0.0
  lam: 0.95
  max_kl: 0.01
  n_epochs_cg: 100
  n_epochs_line_search: 10
  n_features: 32
  n_steps_per_fit: 3000
  preprocessors: null
```





### AntBulletEnv-v0

#### A2C:

```

actor_lr: 0.0007
batch_size: 64
critic_lr: 0.0007
critic_network: A2CNetwork
ent_coeff: 0.01
eps_actor: 0.003
eps_critic: 1.0e-05
max_grad_norm: 0.5
n_features: 64
preprocessors: null

```

#### DDPG:

```

actor_lr: 0.0001
actor_network: DDPGActorNetwork
batch_size: 128
critic_lr: 0.001
critic_network: DDPGCriticNetwork
initial_replay_size: 5000
max_replay_size: 1000000
n_features:
- 400
- 300
tau: 0.001

```

#### PPO:

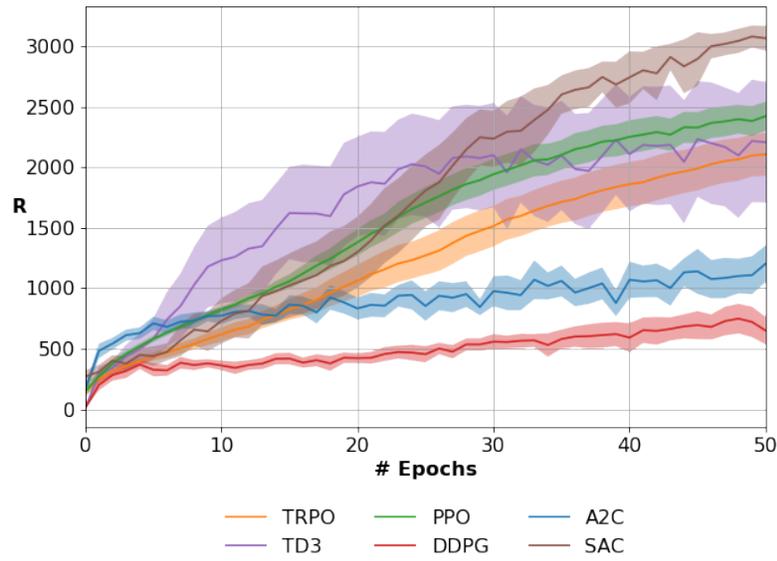
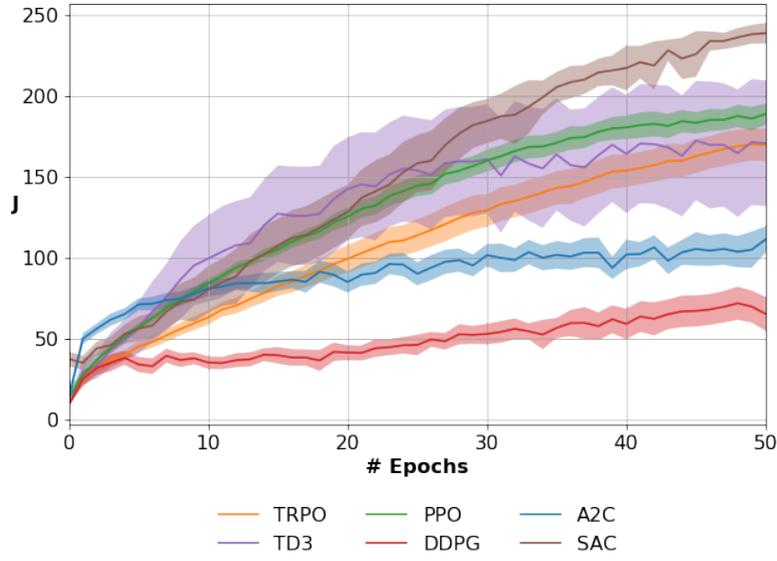
```

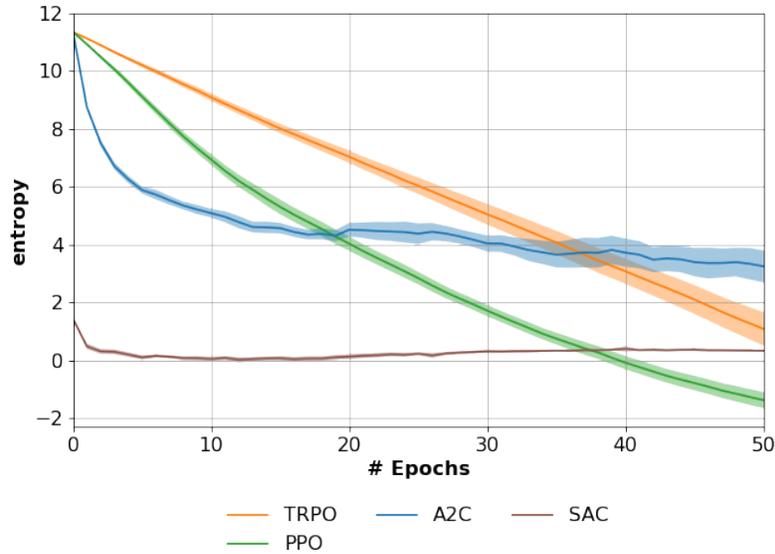
actor_lr: 0.0003
batch_size: 64
critic_fit_params: null
critic_lr: 0.0003
critic_network: TRPONetwork
eps: 0.2
lam: 0.95
n_epochs_policy: 4
n_features: 32

```

(continues on next page)

```
n_steps_per_fit: 3000
preprocessors: null
SAC:
  actor_lr: 0.0001
  actor_network: SACActorNetwork
  batch_size: 256
  critic_lr: 0.0003
  critic_network: SACCriticNetwork
  initial_replay_size: 5000
  lr_alpha: 0.0003
  max_replay_size: 500000
  n_features: 256
  preprocessors: null
  target_entropy: null
  tau: 0.005
  warmup_transitions: 10000
TD3:
  actor_lr: 0.001
  actor_network: TD3ActorNetwork
  batch_size: 100
  critic_lr: 0.001
  critic_network: TD3CriticNetwork
  initial_replay_size: 10000
  max_replay_size: 1000000
  n_features:
    - 400
    - 300
  tau: 0.005
TRPO:
  batch_size: 64
  cg_damping: 0.01
  cg_residual_tol: 1.0e-10
  critic_fit_params: null
  critic_lr: 0.003
  critic_network: TRPONetwork
  ent_coeff: 0.0
  lam: 0.95
  max_kl: 0.01
  n_epochs_cg: 100
  n_epochs_line_search: 10
  n_features: 32
  n_steps_per_fit: 3000
  preprocessors: null
```





### 3.3 Value-Based Benchmarks

We provide the benchmarks for the following Finite Temporal-Difference algorithms:

- **SARSA**
- **QLearning**
- **SpeedyQLearning**
- **WeightedQLearning**
- **DoubleQLearning**
- **SARSALambda**
- **QLambda**

We provide the benchmarks for the following Continuous state Temporal-Difference algorithms:

- **SARSALambdaContinuous**
- **TrueOnlineSARSALambda**

We provide the benchmarks for the following DQN algorithms:

- **DQN**
- **PrioritizedDQN**
- **DoubleDQN**
- **AveragedDQN**
- **DuelingDQN**
- **MaxminDQN**
- **CategoricalDQN**
- **NoisyDQN**

We consider the following environments in the benchmark

### 3.3.1 Finite State Environment Benchmark

Run Parameters	
n_runs	25
n_epochs	100
n_steps	100
n_steps_test	1000

#### GridWorld

```

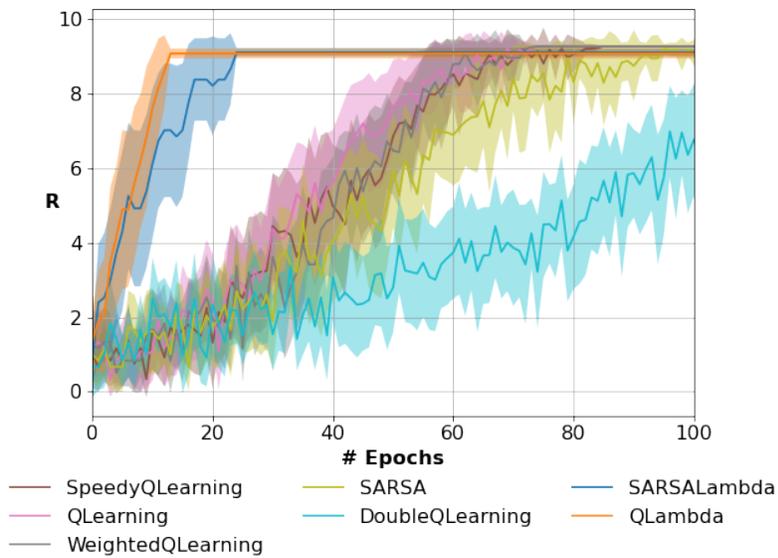
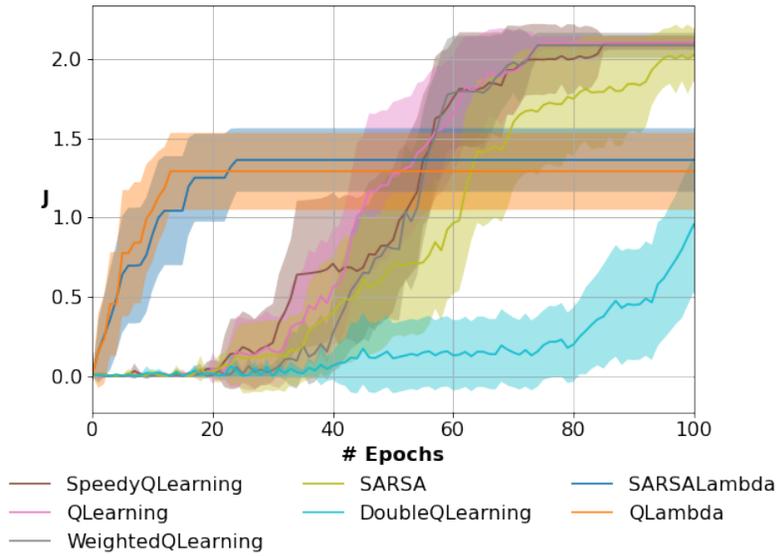
DoubleQLearning:
  decay_eps: 0.5
  decay_lr: 0.8
  epsilon: ExponentialParameter
  epsilon_test: 0.0
  learning_rate: ExponentialParameter
QLambda:
  decay_eps: 0.5
  decay_lr: 0.8
  epsilon: ExponentialParameter
  epsilon_test: 0.0
  lambda_coeff: 0.9
  learning_rate: ExponentialParameter
  trace: replacing
QLearning:
  decay_eps: 0.5
  decay_lr: 0.8
  epsilon: ExponentialParameter
  epsilon_test: 0.0
  learning_rate: ExponentialParameter
SARSA:
  decay_eps: 0.5
  decay_lr: 0.8
  epsilon: ExponentialParameter
  epsilon_test: 0.0
  learning_rate: ExponentialParameter
SARSLambda:
  decay_eps: 0.5
  decay_lr: 0.8
  epsilon: ExponentialParameter
  epsilon_test: 0.0
  lambda_coeff: 0.9
  learning_rate: ExponentialParameter
  trace: replacing
SpeedyQLearning:
  decay_eps: 0.5
  decay_lr: 0.8
  epsilon: ExponentialParameter
  epsilon_test: 0.0
  learning_rate: ExponentialParameter
WeightedQLearning:
  decay_eps: 0.5

```

(continues on next page)

(continued from previous page)

```
decay_lr: 0.8  
epsilon: ExponentialParameter  
epsilon_test: 0.0  
learning_rate: ExponentialParameter  
precision: 1000  
sampling: true
```



### 3.3.2 Gym Environments Benchmarks

Run Parameters	
n_runs	25
n_epochs	100
n_steps	1000
n_episodes_test	10

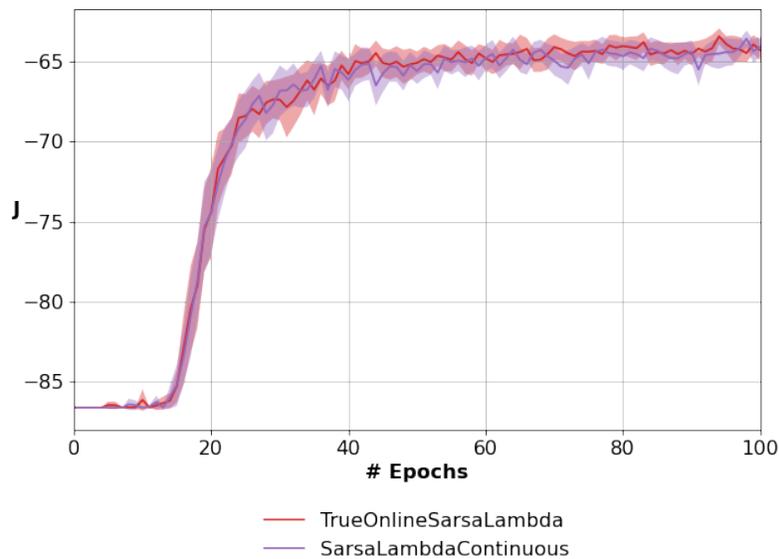
#### MountainCar-v0

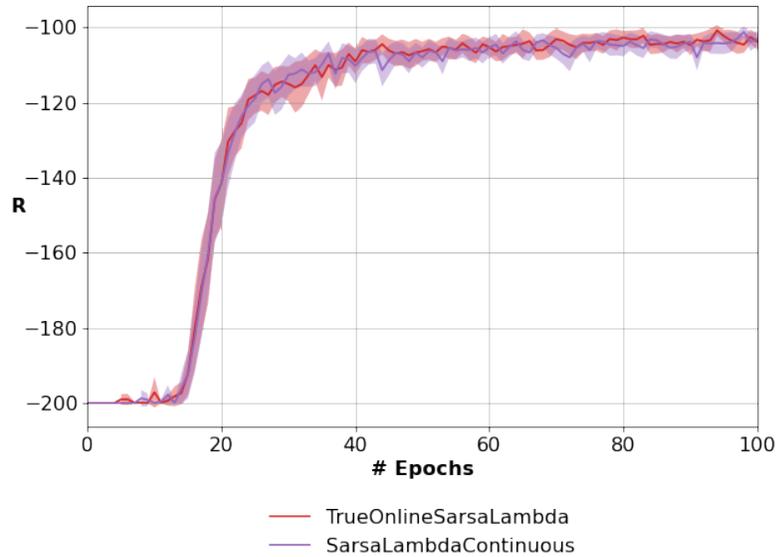
##### SarsaLambdaContinuous:

```
alpha: 0.1
epsilon: 0
epsilon_test: 0.0
lambda_coeff: 0.9
n_tiles: 10
n_tilings: 10
```

##### TrueOnlineSarsaLambda:

```
alpha: 0.1
epsilon: 0
epsilon_test: 0.0
lambda_coeff: 0.9
n_tiles: 10
n_tilings: 10
```





### 3.3.3 Atari Environment Benchmark

Run Parameters	
n_runs	5
n_epochs	200
n_steps	250000
n_episodes_test	125000

#### BreakoutDeterministic-v4

```

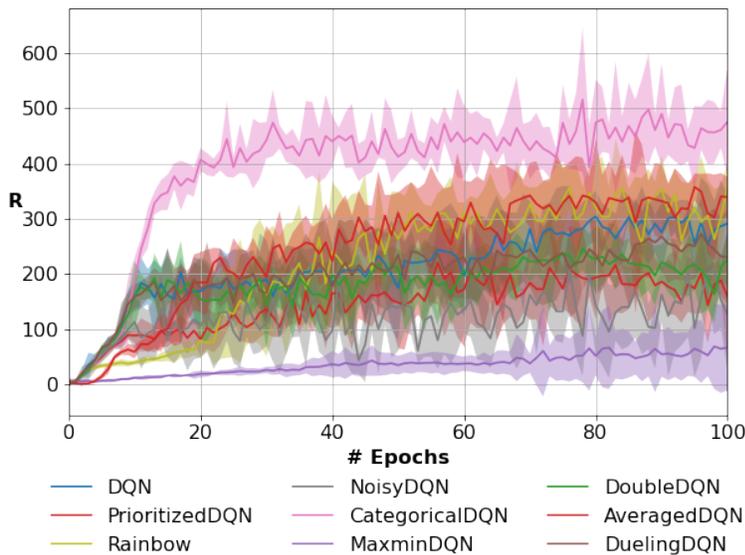
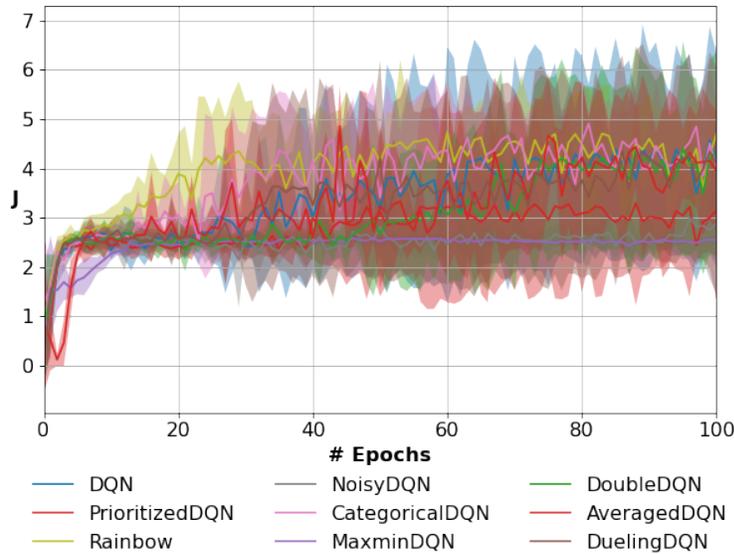
AveragedDQN:
  batch_size: 32
  initial_replay_size: 50000
  lr: 0.0001
  max_replay_size: 1000000
  n_approximators: 10
  n_steps_per_fit: 4
  network: DQNNetwork
  target_update_frequency: 2500
CategoricalDQN:
  batch_size: 32
  initial_replay_size: 50000
  lr: 0.0001
  max_replay_size: 1000000
  n_atoms: 51
  n_features: 512
  n_steps_per_fit: 4
  network: DQNFeatureNetwork
  target_update_frequency: 2500
  v_max: 10
  v_min: -10
DQN:

```

(continues on next page)

(continued from previous page)

```
batch_size: 32
initial_replay_size: 50000
lr: 0.0001
max_replay_size: 1000000
n_steps_per_fit: 4
network: DQNNetwork
target_update_frequency: 2500
DoubleDQN:
batch_size: 32
initial_replay_size: 50000
lr: 0.0001
max_replay_size: 1000000
n_steps_per_fit: 4
network: DQNNetwork
target_update_frequency: 2500
DuelingDQN:
batch_size: 32
initial_replay_size: 50000
lr: 0.0001
max_replay_size: 1000000
n_features: 512
n_steps_per_fit: 4
network: DQNFeatureNetwork
target_update_frequency: 2500
MaxminDQN:
batch_size: 32
initial_replay_size: 50000
lr: 0.0001
max_replay_size: 1000000
n_approximators: 3
n_steps_per_fit: 4
network: DQNNetwork
target_update_frequency: 2500
NoisyDQN:
batch_size: 32
initial_replay_size: 50000
lr: 0.0001
max_replay_size: 1000000
n_features: 512
n_steps_per_fit: 4
network: DQNFeatureNetwork
target_update_frequency: 2500
PrioritizedDQN:
batch_size: 32
initial_replay_size: 50000
lr: 0.0001
max_replay_size: 1000000
n_steps_per_fit: 4
network: DQNNetwork
target_update_frequency: 2500
```



## 3.4 Core functionality

### 3.4.1 Suite

**class BenchmarkSuite**

*log\_dir=None, log\_id=None, use\_timestamp=True, parallel=None, slurm=None* Bases: object

Class to orchestrate the execution of multiple experiments.

**\_\_init\_\_**

*log\_dir=None, log\_id=None, use\_timestamp=True, parallel=None, slurm=None* Constructor.

**Parameters**

- **log\_dir** (*str*) – path to the log directory (Default: `./logs` or `/work/scratch/$USER`)

- **log\_id** (*str*) – log id (Default: benchmark[\_YYYY-mm-dd-HH-MM-SS])
- **use\_timestamp** (*bool*) – select if a timestamp should be appended to the log id
- **parallel** (*dict, None*) – parameters that are passed to the run\_parallel method of the experiment
- **slurm** (*dict, None*) – parameters that are passed to the run\_slurm method of the experiment

**add\_experiments**

*environment\_name, environment\_builder\_params, agent\_names\_list, agent\_builders\_params, \*\*run\_params* Add a set of experiments for the same environment to the suite.

**Parameters**

- **environment\_name** (*str*) – name of the environment for the experiment (E.g. Gym.Pendulum-v0);
- **environment\_builder\_params** (*dict*) – parameters for the environment builder;
- **agent\_names\_list** (*list*) – list of names of the agents for the experiments;
- **agent\_builders\_params** (*list*) – list of dictionaries containing the parameters for the agent builder;
- **run\_params** – Parameters that are passed to the run method of the experiment.

**add\_experiments\_sweeps**

*environment\_name, environment\_builder\_params, agent\_names\_list, agent\_builders\_params, sweeps\_list, \*\*run\_params* Add a set of experiments sweeps for the same environment to the suite.

**Parameters**

- **environment\_name** (*str*) – name of the environment for the experiment (E.g. Gym.Pendulum-v0);
- **environment\_builder\_params** (*dict*) – parameters for the environment builder;
- **agent\_names\_list** (*list*) – list of names of the agents for the experiments;
- **agent\_builders\_params** (*list*) – list of dictionaries containing the parameters for the agent builder;
- **sweeps\_list** (*list*) – list of dictionaries containing the parameter sweep to be executed;
- **run\_params** – Parameters that are passed to the run method of the experiment.

**add\_environment**

*environment\_name, environment\_builder\_params, \*\*run\_params* Add an environment to the benchmarking suite.

**Parameters**

- **environment\_name** (*str*) – name of the environment for the experiment (E.g. Gym.Pendulum-v0);
- **environment\_builder\_params** (*dict*) – parameters for the environment builder;
- **run\_params** – Parameters that are passed to the run method of the experiment.

**add\_agent**

*environment\_name, agent\_name, agent\_params* Add an agent to the benchmarking suite.

**Parameters**

- **environment\_name** (*str*) – name of the environment for the experiment (E.g. Gym.Pendulum-v0);
- **agent\_name** (*str*) – name of the agent for the experiments;
- **agent\_params** (*list*) – dictionary containing the parameters for the agent builder.

**add\_sweep**

*environment\_name, agent\_name, agent\_params, sweep\_dict* Add an agent sweep to the benchmarking suite.

**Parameters**

- **environment\_name** (*str*) – name of the environment for the experiment (E.g. Gym.Pendulum-v0);
- **agent\_name** (*str*) – name of the agent for the experiments;
- **agent\_params** (*list*) – dictionary containing the parameters for the agent builder;
- **sweep\_dict** (*dict*) – dictionary with the sweep configurations.

**run**

*exec\_type='sequential'* Run all experiments in the suite.

**print\_experiments**

Print the experiments in the suite.

**save\_parameters**

Save the experiment parameters in yaml files inside the parameters folder

**save\_plots**

*\*\*plot\_params* Save the result plots to the log directory.

**Parameters *\*\*plot\_params*** – parameters to be passed to the suite visualizer.

**show\_plots**

*\*\*plot\_params* Display the result plots.

**Parameters *\*\*plot\_params*** – parameters to be passed to the suite visualizer.

## 3.4.2 Experiment

**class BenchmarkExperiment**

*agent\_builder, env\_builder, logger* Bases: object

Class to create and run an experiment using MushroomRL

**\_\_init\_\_**

*agent\_builder, env\_builder, logger* Constructor.

**Parameters**

- **agent\_builder** (*AgentBuilder*) – instance of a specific agent builder;
- **env\_builder** (*EnvironmentBuilder*) – instance of an environment builder;
- **logger** (*BenchmarkLogger*) – instance of a benchmark logger.

**run**

*exec\_type='sequential', \*\*run\_params* Execute the experiment.

**Parameters**

- **exec\_type** (*str*, 'sequential') – type of executing the experiment [sequential|parallel|slurm];
- **\*\*run\_params** – parameters for the selected execution type.

**run\_sequential**

*n\_runs, n\_runs\_completed=0, save\_plot=True, \*\*run\_params* Execute the experiment sequential.

**Parameters**

- **n\_runs** (*int*) – number of total runs of the experiment;
- **n\_runs\_completed** (*int*, 0) – number of completed runs of the experiment;
- **save\_plot** (*bool*, *True*) – select if a plot of the experiment should be saved to the log directory;
- **\*\*run\_params** – parameters for executing a benchmark run.

**run\_parallel**

*n\_runs, n\_runs\_completed=0, threading=False, save\_plot=True, max\_concurrent\_runs=None, \*\*run\_params* Execute the experiment in parallel threads.

**Parameters**

- **n\_runs** (*int*) – number of total runs of the experiment;
- **n\_runs\_completed** (*int*, 0) – number of completed runs of the experiment;
- **threading** (*bool*, *False*) – select to use threads instead of processes;
- **save\_plot** (*bool*, *True*) – select if a plot of the experiment should be saved to the log directory;
- **max\_concurrent\_runs** (*int*, -1) – maximum number of concurrent runs. By default it uses the number of cores;
- **\*\*run\_params** – parameters for executing a benchmark run.

**run\_slurm**

*n\_runs, n\_runs\_completed=0, aggregation\_job=True, aggregate\_hours=3, aggregate\_minutes=0, aggregate\_seconds=0, only\_print=False, \*\*run\_params* Execute the experiment with SLURM.

**Parameters**

- **n\_runs** (*int*) – number of total runs of the experiment;
- **n\_runs\_completed** (*int*, 0) – number of completed runs of the experiment;
- **aggregation\_job** (*bool*, *True*) – select if an aggregation job should be scheduled;
- **aggregate\_hours** (*int*, 3) – maximum number of hours for the aggregation job;
- **aggregate\_minutes** (*int*, 0) – maximum number of minutes for the aggregation job;
- **aggregate\_seconds** (*int*, 0) – maximum number of seconds for the aggregation job;
- **only\_print** (*bool*, *False*) – if True, don't launch the benchmarks, only print the submitted commands to the terminal;
- **\*\*run\_params** – parameters for executing a benchmark run.

**reset**

Reset the internal state of the experiment.

**resume**

*logger* Resume an experiment from disk

**start\_timer**

Start the timer.

**stop\_timer**

Stop the timer.

**save\_builders**

Save agent and environment builder to the log directory.

**extend\_and\_save\_J**

*J* Extend *J* with another datapoint and save the current state to the log directory.

**extend\_and\_save\_R**

*R* Extend *R* with another datapoint and save the current state to the log directory.

**extend\_and\_save\_V**

*V* Extend *V* with another datapoint and save the current state to the log directory.

**extend\_and\_save\_entropy**

*entropy* Extend entropy with another datapoint and save the current state to the log directory.

**set\_and\_save\_config**

*\*\*settings* Save the experiment configuration to the log directory.

**set\_and\_save\_stats**

*\*\*info* Save the run statistics to the log directory.

**save\_plot**

Save the result plot to the log directory.

**show\_plot**

Display the result plot.

### 3.4.3 Logger

**class BenchmarkLogger**

*log\_dir=None, log\_id=None, use\_timestamp=True* Bases: `mushroom_rl.core.logger.console_logger.ConsoleLogger`

Class to handle all interactions with the log directory.

**\_\_init\_\_**

*log\_dir=None, log\_id=None, use\_timestamp=True* Constructor.

**Parameters**

- **log\_dir** (*str, None*) – path to the log directory, if not specified defaults to `./logs` or to `/work/scratch/$USER` if the second directory exists;
- **log\_id** (*str, None*) – log id, if not specified defaults to: `benchmark[_YY-mm-ddTHH:MM:SS.zzz]`;
- **use\_timestamp** (*bool, True*) – select if a timestamp should be appended to the log id.

**set\_log\_dir**

*log\_dir* Set the directory for logging.

**Parameters** **log\_dir** (*str*) – path of the directory.

**get\_log\_dir**

**Returns** The path of the logging directory.

**set\_log\_id**

*log\_id, use\_timestamp=True* Set the id of the logged folder.

**Parameters**

- **log\_id** (*str*) – id of the logged folder;
- **use\_timestamp** (*bool, True*) – whether to use the timestamp or not.

**get\_log\_id**

**Returns** The id of the logged folder.

**get\_path**

*filename=""* Get the path of the given file. If no filename is given, it returns the path of the logging folder.

**Parameters** **filename** (*str, ''*) – the name of the file.

**Returns** The complete path of the logged file.

**get\_params\_path**

*filename=""* Get the path of the parameters of the given file. If no filename is given, it returns the path of the parameters folder.

**Parameters** **filename** (*str, ''*) – the name of the file.

**Returns** The complete path of the logged file.

**get\_figure\_path**

*filename="", subfolder=None* Get the path of the figures of the given file. If no filename is given, it returns the path of the figures folder.

**Parameters**

- **filename** (*str, ''*) – the name of the file;
- **subfolder** (*None*) – the name of a subfolder to add.

**Returns** The complete path of the logged file.

**save\_J**

*J* Save the log of the cumulative discounted reward.

**load\_J**

**Returns** The log of the cumulative discounted reward.

**save\_R**

*R* Save the log of the cumulative reward.

**load\_R**

**Returns** The log of the cumulative reward.

**save\_V**

*V* Save the log of the value function.

**load\_V**

**Returns** The log of the value function.

**save\_entropy**

*entropy* Save the log of the entropy function.

**load\_entropy**

**Returns** The log of the entropy function.

**exists\_policy\_entropy**

**Returns** True if the log of the entropy exists, False otherwise.

**exists\_value\_function**

**Returns** True if the log of the value function exists, False otherwise.

**save\_best\_agent**

*agent* Save the best agent in the respective path.

**Parameters** **agent** (*object*) – the agent to save.

**save\_last\_agent**

*agent* Save the last agent in the respective path.

**Parameters** **agent** (*object*) – the agent to save.

**exists\_best\_agent**

**Returns** True if the entropy file exists, False otherwise.

**load\_best\_agent**

**Returns** The best agent.

**load\_last\_agent**

**Returns** The last agent.

**save\_environment\_builder**

*env\_builder* Save the environment builder using the respective path.

**Parameters** **env\_builder** (*str*) – the environment builder to save.

**load\_environment\_builder**

**Returns** The environment builder.

**save\_agent\_builder**

*agent\_builder* Save the agent builder using the respective path.

**Parameters** **agent\_builder** (*str*) – the agent builder to save.

**load\_agent\_builder**

**Returns** The agent builder.

**save\_config**

*config* Save the config file using the respective path.

**Parameters** **config** (*str*) – the config file to save.

**load\_config**

**Returns** The config file.

**exists\_stats**

**Returns** True if the entropy file exists, False otherwise.

**save\_stats**

*stats* Save the statistic file using the respective path.

**Parameters** **stats** (*str*) – the statistics file to save.

**load\_stats**

**Returns** The statistics file.

**save\_params**

*env, params* Save the parameters file.

**Parameters**

- **env** (*str*) – the environment used;
- **params** (*str*) – the parameters file to save.

**save\_figure**

*figure, filename, subfolder=None, as\_pdf=False, transparent=True* Save the figure file using the respective path.

**Parameters**

- **figure** (*object*) – the figure to save;
- **filename** (*str*) – the name of the figure;
- **subfolder** (*str, None*) – optional subfolder where to save the figure;
- **as\_pdf** (*bool, False*) – whether to save the figure in PDF or not;
- **transparent** (*bool, True*) – whether the figure should be transparent or not.

**classmethod from\_path**

*path* Method to create a BenchmarkLogger from a path.

### 3.4.4 Visualizer

**class BenchmarkVisualizer**

*logger, data=None, has\_entropy=None, has\_value=None, id=1* Bases: `object`

Class to handle all visualizations of the experiment.

**plot\_counter** = 0

**\_\_init\_\_**

*logger, data=None, has\_entropy=None, has\_value=None, id=1* Constructor.

**Parameters**

- **logger** (`BenchmarkLogger`) – logger to be used;
- **data** (*dict, None*) – dictionary with data points for visualization;
- **has\_entropy** (*bool, None*) – select if entropy is available for the algorithm.

**is\_data\_persisted**

Check if data was passed as dictionary or should be read from log directory.

**get\_J**

Get J from dictionary or log directory.

**get\_R**

Get R from dictionary or log directory.

**get\_V**

Get V from dictionary or log directory.

**get\_entropy**

Get entropy from dictionary or log directory.

**get\_report**

Create report plot with matplotlib.

**save\_report**

*file\_name='report\_plot'* Method to save an image of a report of the training metrics from a performed experiment.

**show\_report**

Method to show a report of the training metrics from a performed experiment.

**show\_agent**

*episodes=5, mdp\_render=False* Method to run and visualize the best builders in the environment.

**classmethod from\_path**

*path* Method to create a BenchmarkVisualizer from a path.

**class BenchmarkSuiteVisualizer**

*logger, is\_sweep, color\_cycle=None, y\_limit=None, legend=None* Bases: `object`

Class to handle visualization of a benchmark suite.

**plot\_counter = 0**

**\_\_init\_\_**

*logger, is\_sweep, color\_cycle=None, y\_limit=None, legend=None* Constructor.

**Parameters**

- **logger** (`BenchmarkLogger`) – logger to be used;
- **is\_sweep** (`bool`) – whether the benchmark is a parameter sweep.
- **color\_cycle** (`dict, None`) – dictionary with colors to be used for each algorithm;
- **y\_limit** (`dict, None`) – dictionary with environment specific plot limits.
- **legend** (`dict, None`) – dictionary with environment specific legend parameters.

**get\_report**

*env, data\_type, selected\_alg=None* Create report plot with matplotlib.

**get\_boxplot**

*env, metric\_type, data\_type, selected\_alg=None* Create boxplot with matplotlib for a given metric.

**Parameters**

- **env** (`str`) – The environment name;
- **metric\_type** (`str`) – The metric to compute.

**Returns** A figure with the desired boxplot of the given metric.

**save\_reports**

*as\_pdf=True, transparent=True, alg\_sweep=False* Method to save an image of a report of the training metrics from a performed experiment.

**Parameters**

- **as\_pdf** (`bool, True`) – whether to save the reports as pdf files or png;
- **transparent** (`bool, True`) – If true, the figure background is transparent and not white;
- **alg\_sweep** (`bool, False`) – If true, the method will generate a separate figure for each algorithm sweep.

**save\_boxplots**

*as\_pdf=True, transparent=True, alg\_sweep=False* Method to save an image of a report of the training metrics from a performed experiment.

**Parameters**

- **as\_pdf** (*bool, True*) – whether to save the reports as pdf files or png;
- **transparent** (*bool, True*) – If true, the figure background is transparent and not white;
- **alg\_sweep** (*bool, False*) – If true, thw method will generate a separate figure for each algorithm sweep.

**show\_reports**

*boxplots=True, alg\_sweep=False* Method to show a report of the training metrics from a performend experiment.

**Parameters** **alg\_sweep** (*bool, False*) – If true, thw method will generate a separate figure for each algorithm sweep.

## 3.5 Builders

**class EnvironmentBuilder**

*env\_name, env\_params* Bases: object

Class to spawn instances of a MushroomRL environment

**\_\_init\_\_**

*env\_name, env\_params* Constructor

**Parameters**

- **env\_name** – name of the environment to build;
- **env\_params** – required parameters to build the specified environment.

**build**

Build and return an environment

**static set\_eval\_mode**

*env, eval* Make changes to the environment for evaluation mode.

**Parameters**

- **env** (*Environment*) – the environment to change;
- **eval** (*bool*) – flag for activating evaluation mode.

**copy**

Create a deepcopy of the environment\_builder and return it

**class AgentBuilder**

*n\_steps\_per\_fit=None, n\_episodes\_per\_fit=None, compute\_policy\_entropy=True, compute\_entropy\_with\_states=False, compute\_value\_function=True, preprocessors=None* Bases: object

Base class to spawn instances of a MushroomRL agent

**\_\_init\_\_**

*n\_steps\_per\_fit=None, n\_episodes\_per\_fit=None, compute\_policy\_entropy=True, compute\_entropy\_with\_states=False, compute\_value\_function=True, preprocessors=None* Initialize AgentBuilder

**get\_fit\_params**

Get `n_steps_per_fit` and `n_episodes_per_fit` for the specific AgentBuilder

**set\_preprocessors**

*preprocessors* Set preprocessor for the specific AgentBuilder

**Parameters** `preprocessors` – list of preprocessor classes.

**get\_preprocessors**

Get preprocessors for the specific AgentBuilder

**copy**

Create a deepcopy of the AgentBuilder and return it

**build**

*mdp\_info* Build and return the AgentBuilder

**Parameters** `mdp_info` (*MDPInfo*) – information about the environment.

**compute\_Q**

*agent, states* Compute the Q Value for an AgentBuilder

**Parameters**

- **agent** (*Agent*) – the considered agent;
- **states** (*np.ndarray*) – the set of states over which we need to compute the Q function.

**set\_eval\_mode**

*agent, eval* Set the eval mode for the agent. This function can be overwritten by any agent builder to setup specific evaluation mode for the agent.

**Parameters**

- **agent** (*Agent*) – the considered agent;
- **eval** (*bool*) – whether to set eval mode (true) or learn mode.

**classmethod default**

*get\_default\_dict=False, \*\*kwargs* Create a default initialization for the specific AgentBuilder and return it

## 3.5.1 Policy Search Builders

### Policy Gradient

**class PolicyGradientBuilder**

*n\_episodes\_per\_fit, optimizer, \*\*kwargs* Bases: *mushroom\_rl\_benchmark.builders.agent\_builder.AgentBuilder*

AgentBuilder for Policy Gradient Methods. The current builder uses a state dependant gaussian with diagonal standard deviation and linear mean.

**\_\_init\_\_**

*n\_episodes\_per\_fit, optimizer, \*\*kwargs* Constructor.

**Parameters**

- **optimizer** (*Optimizer*) – optimizer to be used by the policy gradient algorithm;
- **\*\*kwargs** – others algorithms parameters.

**build**

*mdp\_info* Build and return the AgentBuilder

**Parameters** *mdp\_info* (*MDPInfo*) – information about the environment.

**classmethod default**

*n\_episodes\_per\_fit=25, alpha=0.01, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

**compute\_Q**

*agent, states* Compute the Q Value for an AgentBuilder

**Parameters**

- **agent** (*Agent*) – the considered agent;
- **states** (*np.ndarray*) – the set of states over which we need to compute the Q function.

**class REINFORCEBuilder**

*n\_episodes\_per\_fit, optimizer, \*\*kwargs* Bases: *mushroom\_rl\_benchmark.builders.policy\_search.policy\_gradient.PolicyGradientBuilder*

**alg\_class**

alias of *mushroom\_rl.algorithms.policy\_search.policy\_gradient.reinforce.REINFORCE*

**class GPOMDPBuilder**

*n\_episodes\_per\_fit, optimizer, \*\*kwargs* Bases: *mushroom\_rl\_benchmark.builders.policy\_search.policy\_gradient.PolicyGradientBuilder*

**alg\_class**

alias of *mushroom\_rl.algorithms.policy\_search.policy\_gradient.gpomdp.GPOMDP*

**class eNACBuilder**

*n\_episodes\_per\_fit, optimizer, \*\*kwargs* Bases: *mushroom\_rl\_benchmark.builders.policy\_search.policy\_gradient.PolicyGradientBuilder*

**alg\_class**

alias of *mushroom\_rl.algorithms.policy\_search.policy\_gradient.enac.eNAC*

**Black-Box optimization****class BBOBuilder**

*n\_episodes\_per\_fit, \*\*kwargs* Bases: *mushroom\_rl\_benchmark.builders.agent\_builder.AgentBuilder*

AgentBuilder for Black Box optimization methods. The current builder uses a simple deterministic linear policy and gaussian Diagonal distribution.

**\_\_init\_\_**

*n\_episodes\_per\_fit, \*\*kwargs* Constructor.

**Parameters**

- **optimizer** (*Optimizer*) – optimizer to be used by the policy gradient algorithm;
- **\*\*kwargs** – others algorithms parameters.

**build**

*mdp\_info* Build and return the AgentBuilder

Parameters `mdp_info` (*MDPInfo*) – information about the environment.

**classmethod default**

`n_episodes_per_fit=25, alpha=0.01, get_default_dict=False` Create a default initialization for the specific AgentBuilder and return it

**compute\_Q**

`agent, states` Compute the Q Value for an AgentBuilder

**Parameters**

- **agent** (*Agent*) – the considered agent;
- **states** (*np.ndarray*) – the set of states over which we need to compute the Q function.

**class PGPEBuilder**

`n_episodes_per_fit, optimizer` Bases: `mushroom_rl_benchmark.builders.policy_search.black_box_optimization.BBOBuilder`

**alg\_class**

alias of `mushroom_rl.algorithms.policy_search.black_box_optimization.pgpe.PGPE`

**\_\_init\_\_**

`n_episodes_per_fit, optimizer` Constructor.

**Parameters**

- **optimizer** (*Optimizer*) – optimizer to be used by the policy gradient algorithm;
- **\*\*kwargs** – others algorithms parameters.

**classmethod default**

`n_episodes_per_fit=25, alpha=0.3, get_default_dict=False` Create a default initialization for the specific AgentBuilder and return it

**class RWRBuilder**

`n_episodes_per_fit, beta` Bases: `mushroom_rl_benchmark.builders.policy_search.black_box_optimization.BBOBuilder`

**alg\_class**

alias of `mushroom_rl.algorithms.policy_search.black_box_optimization.rwr.RWR`

**\_\_init\_\_**

`n_episodes_per_fit, beta` Constructor.

**Parameters**

- **optimizer** (*Optimizer*) – optimizer to be used by the policy gradient algorithm;
- **\*\*kwargs** – others algorithms parameters.

**classmethod default**

`n_episodes_per_fit=25, beta=0.01, get_default_dict=False` Create a default initialization for the specific AgentBuilder and return it

**class REPSBuilder**

`n_episodes_per_fit, eps` Bases: `mushroom_rl_benchmark.builders.policy_search.black_box_optimization.BBOBuilder`

**alg\_class**

alias of `mushroom_rl.algorithms.policy_search.black_box_optimization.reps.REPS`

**\_\_init\_\_**

`n_episodes_per_fit, eps` Constructor.

**Parameters**

- **optimizer** (*Optimizer*) – optimizer to be used by the policy gradient algorithm;
- **\*\*kwargs** – others algorithms parameters.

**classmethod default**

`n_episodes_per_fit=25, eps=0.05, get_default_dict=False` Create a default initialization for the specific AgentBuilder and return it

**class ConstrainedREPSBuilder**

`n_episodes_per_fit, eps, kappa` Bases: `mushroom_rl_benchmark.builders.policy_search.black_box_optimization.BBOBuilder`

**alg\_class**

alias of `mushroom_rl.algorithms.policy_search.black_box_optimization.constrained_reps.ConstrainedREPS`

**\_\_init\_\_**

`n_episodes_per_fit, eps, kappa` Constructor.

**Parameters**

- **optimizer** (*Optimizer*) – optimizer to be used by the policy gradient algorithm;
- **\*\*kwargs** – others algorithms parameters.

**classmethod default**

`n_episodes_per_fit=25, eps=0.05, kappa=0.01, get_default_dict=False` Create a default initialization for the specific AgentBuilder and return it

## 3.5.2 Value Based Builders

### Temporal Difference

**class TDFiniteBuilder**

`learning_rate, epsilon, epsilon_test, **alg_params` Bases: `mushroom_rl_benchmark.builders.agent_builder.AgentBuilder`

AgentBuilder for a generic TD algorithm (for finite states).

**\_\_init\_\_**

`learning_rate, epsilon, epsilon_test, **alg_params` Constructor.

**Parameters**

- **epsilon** (*Parameter*) – exploration coefficient for learning;
- **epsilon\_test** (*Parameter*) – exploration coefficient for test.

**build**

`mdp_info` Build and return the AgentBuilder

**Parameters** `mdp_info` (*MDPInfo*) – information about the environment.

**compute\_Q**

*agent, states* Compute the Q Value for an AgentBuilder

**Parameters**

- **agent** (*Agent*) – the considered agent;
- **states** (*np.ndarray*) – the set of states over which we need to compute the Q function.

**set\_eval\_mode**

*agent, eval* Set the eval mode for the agent. This function can be overwritten by any agent builder to setup specific evaluation mode for the agent.

**Parameters**

- **agent** (*Agent*) – the considered agent;
- **eval** (*bool*) – whether to set eval mode (true) or learn mode.

**classmethod default**

*learning\_rate=0.9, epsilon=0.1, decay\_lr=0.0, decay\_eps=0.0, epsilon\_test=0.0, get\_default\_dict=False*  
Create a default initialization for the specific AgentBuilder and return it

**class QLearningBuilder**

*learning\_rate, epsilon, epsilon\_test* Bases: *mushroom\_rl\_benchmark.builders.value.td.td\_finite.TDFiniteBuilder*

**alg\_class**

alias of *mushroom\_rl.algorithms.value.td.q\_learning.QLearning*

**\_\_init\_\_**

*learning\_rate, epsilon, epsilon\_test* Constructor.

**Parameters**

- **epsilon** (*Parameter*) – exploration coefficient for learning;
- **epsilon\_test** (*Parameter*) – exploration coefficient for test.

**class SARSABuilder**

*learning\_rate, epsilon, epsilon\_test* Bases: *mushroom\_rl\_benchmark.builders.value.td.td\_finite.TDFiniteBuilder*

**alg\_class**

alias of *mushroom\_rl.algorithms.value.td.sarsa.SARSA*

**\_\_init\_\_**

*learning\_rate, epsilon, epsilon\_test* Constructor.

**Parameters**

- **epsilon** (*Parameter*) – exploration coefficient for learning;
- **epsilon\_test** (*Parameter*) – exploration coefficient for test.

**class SpeedyQLearningBuilder**

*learning\_rate, epsilon, epsilon\_test* Bases: *mushroom\_rl\_benchmark.builders.value.td.td\_finite.TDFiniteBuilder*

**alg\_class**

alias of *mushroom\_rl.algorithms.value.td.speedy\_q\_learning.SpeedyQLearning*

**\_\_init\_\_**

*learning\_rate, epsilon, epsilon\_test* Constructor.

**Parameters**

- **epsilon** (*Parameter*) – exploration coefficient for learning;
- **epsilon\_test** (*Parameter*) – exploration coefficient for test.

**class DoubleQLearningBuilder**

*learning\_rate, epsilon, epsilon\_test* Bases: *mushroom\_rl\_benchmark.builders.value.td.td\_finite.TDFiniteBuilder*

**alg\_class**

alias of *mushroom\_rl.algorithms.value.td.double\_q\_learning.DoubleQLearning*

**\_\_init\_\_**

*learning\_rate, epsilon, epsilon\_test* Constructor.

**Parameters**

- **epsilon** (*Parameter*) – exploration coefficient for learning;
- **epsilon\_test** (*Parameter*) – exploration coefficient for test.

**compute\_Q**

*agent, states* Compute the Q Value for an AgentBuilder

**Parameters**

- **agent** (*Agent*) – the considered agent;
- **states** (*np.ndarray*) – the set of states over which we need to compute the Q function.

**class WeightedQLearningBuilder**

*learning\_rate, epsilon, epsilon\_test, sampling, precision* Bases: *mushroom\_rl\_benchmark.builders.value.td.td\_finite.TDFiniteBuilder*

**alg\_class**

alias of *mushroom\_rl.algorithms.value.td.weighted\_q\_learning.WeightedQLearning*

**\_\_init\_\_**

*learning\_rate, epsilon, epsilon\_test, sampling, precision* Constructor.

**Parameters**

- **sampling** (*bool, True*) – use the approximated version to speed up the computation;
- **precision** (*int, 1000*) – number of samples to use in the approximated version.

**classmethod default**

*learning\_rate=0.9, epsilon=0.1, decay\_lr=0.0, decay\_eps=0.0, epsilon\_test=0.0, sampling=True, precision=1000, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

**class TDTraceBuilder**

*learning\_rate, epsilon, epsilon\_test, lambda\_coeff, trace* Bases: *mushroom\_rl\_benchmark.builders.value.td.td\_finite.TDFiniteBuilder*

Builder for TD algorithms with eligibility traces and finite states.

**\_\_init\_\_**

*learning\_rate, epsilon, epsilon\_test, lambda\_coeff, trace* Constructor.

*lambda\_coeff* ([float, Parameter]): eligibility trace coefficient; *trace* (str): type of eligibility trace to use.

**classmethod default**

*learning\_rate=0.9, epsilon=0.1, decay\_lr=0.0, decay\_eps=0.0, epsilon\_test=0.0, lambda\_coeff=0.9, trace='replacing', get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

**class SarsaLambdaBuilder**

*learning\_rate, epsilon, epsilon\_test, lambda\_coeff, trace* Bases: *mushroom\_rl\_benchmark.builders.value.td.td\_trace.TDTraceBuilder*

**alg\_class**

alias of *mushroom\_rl.algorithms.value.td.sarsa\_lambda.SarsaLambda*

**class QLambdaBuilder**

*learning\_rate, epsilon, epsilon\_test, lambda\_coeff, trace* Bases: *mushroom\_rl\_benchmark.builders.value.td.td\_trace.TDTraceBuilder*

**alg\_class**

alias of *mushroom\_rl.algorithms.value.td.q\_lambda.QLambda*

**class SarsaLambdaContinuousBuilder**

*policy, approximator, learning\_rate, lambda\_coeff, epsilon, epsilon\_test, n\_tilings, n\_tiles* Bases: *mushroom\_rl\_benchmark.builders.value.td.td\_continuous.TDContinuousBuilder*

AgentBuilder for Sarsa(Lambda) Continuous. Using tiles as function approximator.

**\_\_init\_\_**

*policy, approximator, learning\_rate, lambda\_coeff, epsilon, epsilon\_test, n\_tilings, n\_tiles* Constructor.

**Parameters approximator** (*class*) – Q-function approximator.

**build**

*mdp\_info* Build and return the AgentBuilder

**Parameters mdp\_info** (*MDPInfo*) – information about the environment.

**classmethod default**

*alpha=0.1, lambda\_coeff=0.9, epsilon=0.0, decay\_eps=0.0, epsilon\_test=0.0, n\_tilings=10, n\_tiles=10, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

**class TrueOnlineSarsaLambdaBuilder**

*policy, learning\_rate, lambda\_coeff, epsilon, epsilon\_test, n\_tilings, n\_tiles* Bases: *mushroom\_rl\_benchmark.builders.value.td.td\_continuous.TDContinuousBuilder*

AgentBuilder for True Online Sarsa(Lambda) Continuous. Using tiles as function approximator.

**\_\_init\_\_**

*policy, learning\_rate, lambda\_coeff, epsilon, epsilon\_test, n\_tilings, n\_tiles* Constructor.

**build**

*mdp\_info* Build and return the AgentBuilder

**Parameters mdp\_info** (*MDPInfo*) – information about the environment.

**classmethod default**

*alpha=0.1, lambda\_coeff=0.9, epsilon=0.0, decay\_eps=0.0, epsilon\_test=0.0, n\_tilings=10, n\_tiles=10, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

## DQN

**class DQNBuilder**

*policy, approximator, approximator\_params, alg\_params, n\_steps\_per\_fit=1* Bases: *mushroom\_rl\_benchmark.builders.agent\_builder.AgentBuilder*

AgentBuilder for Deep Q-Network (DQN).

**`__init__`**

*policy, approximator, approximator\_params, alg\_params, n\_steps\_per\_fit=1* Constructor.

**Parameters**

- **policy** (*Policy*) – policy class;
- **approximator** (*dict*) – Q-function approximator;
- **approximator\_params** (*dict*) – parameters of the Q-function approximator;
- **alg\_params** (*dict*) – parameters for the algorithm;
- **n\_steps\_per\_fit** (*int, 1*) – number of steps per fit.

**build**

*mdp\_info* Build and return the AgentBuilder

**Parameters** **mdp\_info** (*MDPInfo*) – information about the environment.

**compute\_Q**

*agent, states* Compute the Q Value for an AgentBuilder

**Parameters**

- **agent** (*Agent*) – the considered agent;
- **states** (*np.ndarray*) – the set of states over which we need to compute the Q function.

**set\_eval\_mode**

*agent, eval* Set the eval mode for the agent. This function can be overwritten by any agent builder to setup specific evaluation mode for the agent.

**Parameters**

- **agent** (*Agent*) – the considered agent;
- **eval** (*bool*) – whether to set eval mode (true) or learn mode.

**classmethod default**

*lr=0.0001, network=<class 'mushroom\_rl\_benchmark.builders.network.dqn\_network.DQNNetwork'>, initial\_replay\_size=50000, max\_replay\_size=1000000, batch\_size=32, target\_update\_frequency=2500, n\_steps\_per\_fit=1, use\_cuda=False, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

**class DoubleDQNBuilder**

*policy, approximator, approximator\_params, alg\_params, n\_steps\_per\_fit=1* Bases: [mushroom\\_rl\\_benchmark.builders.value.dqn.dqn.DQNBuilder](#)

**build**

*mdp\_info* Build and return the AgentBuilder

**Parameters** **mdp\_info** (*MDPInfo*) – information about the environment.

**class AveragedDQNBuilder**

*policy, approximator, approximator\_params, alg\_params, n\_steps\_per\_fit=1* Bases: [mushroom\\_rl\\_benchmark.builders.value.dqn.dqn.DQNBuilder](#)

**build**

*mdp\_info* Build and return the AgentBuilder

**Parameters** **mdp\_info** (*MDPInfo*) – information about the environment.

**classmethod default**

*lr=0.0001, network=<class 'mushroom\_rl\_benchmark.builders.network.dqn\_network.DQNNetwork'>, initial\_replay\_size=50000, max\_replay\_size=1000000, batch\_size=32, target\_update\_frequency=2500, n\_steps\_per\_fit=1, n\_approximators=10, use\_cuda=False, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

**class PrioritizedDQNBuilder**

*policy, approximator, approximator\_params, alg\_params, n\_steps\_per\_fit=1* Bases: *mushroom\_rl\_benchmark.builders.value.dqn.dqn.DQNBuilder*

**build**

*mdp\_info* Build and return the AgentBuilder

**Parameters** *mdp\_info* (*MDPInfo*) – information about the environment.

**classmethod default**

*lr=0.0001, network=<class 'mushroom\_rl\_benchmark.builders.network.dqn\_network.DQNNetwork'>, initial\_replay\_size=50000, max\_replay\_size=1000000, batch\_size=32, target\_update\_frequency=2500, n\_steps\_per\_fit=1, use\_cuda=False, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

**class DuelingDQNBuilder**

*policy, approximator, approximator\_params, alg\_params, n\_steps\_per\_fit=1* Bases: *mushroom\_rl\_benchmark.builders.value.dqn.dqn.DQNBuilder*

**build**

*mdp\_info* Build and return the AgentBuilder

**Parameters** *mdp\_info* (*MDPInfo*) – information about the environment.

**classmethod default**

*lr=0.0001, network=<class 'mushroom\_rl\_benchmark.builders.network.dqn\_network.DQNFeatureNetwork'>, initial\_replay\_size=50000, max\_replay\_size=1000000, batch\_size=32, target\_update\_frequency=2500, n\_features=512, n\_steps\_per\_fit=1, use\_cuda=False, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

**class MaxminDQNBuilder**

*policy, approximator, approximator\_params, alg\_params, n\_steps\_per\_fit=1* Bases: *mushroom\_rl\_benchmark.builders.value.dqn.dqn.DQNBuilder*

**build**

*mdp\_info* Build and return the AgentBuilder

**Parameters** *mdp\_info* (*MDPInfo*) – information about the environment.

**classmethod default**

*lr=0.0001, network=<class 'mushroom\_rl\_benchmark.builders.network.dqn\_network.DQNNetwork'>, initial\_replay\_size=50000, max\_replay\_size=1000000, batch\_size=32, target\_update\_frequency=2500, n\_steps\_per\_fit=1, n\_approximators=3, use\_cuda=False, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

**class NoisyDQNBuilder**

*policy, approximator, approximator\_params, alg\_params, n\_steps\_per\_fit=1* Bases: *mushroom\_rl\_benchmark.builders.value.dqn.dqn.DQNBuilder*

**build**

*mdp\_info* Build and return the AgentBuilder

**Parameters** *mdp\_info* (*MDPInfo*) – information about the environment.

**classmethod default**

*lr=0.0001, network=<class 'mushroom\_rl\_benchmark.builders.network.dqn\_network.DQNFeatureNetwork'>,*

*initial\_replay\_size=50000, max\_replay\_size=1000000, batch\_size=32, target\_update\_frequency=2500, n\_features=512, n\_steps\_per\_fit=1, use\_cuda=False, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

#### class CategoricalDQNBuilder

*policy, approximator, approximator\_params, alg\_params, n\_steps\_per\_fit=1* Bases: *mushroom\_rl\_benchmark.builders.value.dqn.dqn.DQNBuilder*

#### build

*mdp\_info* Build and return the AgentBuilder

**Parameters** *mdp\_info* (*MDPInfo*) – information about the environment.

#### classmethod default

*lr=0.0001, network=<class 'mushroom\_rl\_benchmark.builders.network.dqn\_network.DQNFeatureNetwork'>, initial\_replay\_size=50000, max\_replay\_size=1000000, batch\_size=32, target\_update\_frequency=2500, n\_features=512, n\_steps\_per\_fit=1, v\_min=-10, v\_max=10, n\_atoms=51, use\_cuda=False, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

### 3.5.3 Actor Critic Builders

#### Classic AC

#### class StochasticACBuilder

*std\_0, alpha\_theta, alpha\_v, lambda\_par, n\_tilings, n\_tiles, \*\*kwargs* Bases: *mushroom\_rl\_benchmark.builders.agent\_builder.AgentBuilder*

Builder for the stochastic actor critic algorithm. Using linear approximator with tiles for mean, standard deviation and value function approximator. The value function approximator also uses a bias term.

#### \_\_init\_\_

*std\_0, alpha\_theta, alpha\_v, lambda\_par, n\_tilings, n\_tiles, \*\*kwargs* Constructor.

#### Parameters

- **std\_0** (*float*) – initial standard deviation;
- **alpha\_theta** (*Parameter*) – Learning rate for the policy;
- **alpha\_v** (*Parameter*) – Learning rate for the value function;
- **n\_tilings** (*int*) – number of tilings to be used as approximator;
- **n\_tiles** (*int*) – number of tiles for each state space dimension.

#### build

*mdp\_info* Build and return the AgentBuilder

**Parameters** *mdp\_info* (*MDPInfo*) – information about the environment.

#### classmethod default

*std\_0=1.0, alpha\_theta=0.001, alpha\_v=0.1, lambda\_par=0.9, n\_tilings=10, n\_tiles=11, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

#### compute\_Q

*agent, states* Compute the Q Value for an AgentBuilder

#### Parameters

- **agent** (*Agent*) – the considered agent;

- **states** (*np.ndarray*) – the set of states over which we need to compute the Q function.

**class COPDAC\_QBuilder**

*std\_exp, std\_eval, alpha\_theta, alpha\_omega, alpha\_v, n\_tilings, n\_tiles, \*\*kwargs* Bases: *mushroom\_rl\_benchmark.builders.agent\_builder.AgentBuilder*

Builder for the COPDAQ\_Q actor critic algorithm. Using linear approximator with tiles for the mean and value function approximator.

**\_\_init\_\_**

*std\_exp, std\_eval, alpha\_theta, alpha\_omega, alpha\_v, n\_tilings, n\_tiles, \*\*kwargs* Constructor.

**Parameters**

- **std\_exp** (*float*) – exploration standard deviation;
- **std\_eval** (*float*) – evaluation standard deviation;
- **alpha\_theta** (*Parameter*) – Learning rate for the policy;
- **alpha\_omega** (*Parameter*) – Learning rate for the
- **alpha\_v** (*Parameter*) – Learning rate for the value function;
- **n\_tilings** (*int*) – number of tilings to be used as approximator;
- **n\_tiles** (*int*) – number of tiles for each state space dimension.

**build**

*mdp\_info* Build and return the AgentBuilder

**Parameters** *mdp\_info* (*MDPInfo*) – information about the environment.

**set\_eval\_mode**

*agent, eval* Set the eval mode for the agent. This function can be overwritten by any agent builder to setup specific evaluation mode for the agent.

**Parameters**

- **agent** (*Agent*) – the considered agent;
- **eval** (*bool*) – whether to set eval mode (true) or learn mode.

**classmethod default**

*std\_exp=0.1, std\_eval=0.001, alpha\_theta=0.005, alpha\_omega=0.5, alpha\_v=0.5, n\_tilings=10, n\_tiles=11, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

**compute\_Q**

*agent, states* Compute the Q Value for an AgentBuilder

**Parameters**

- **agent** (*Agent*) – the considered agent;
- **states** (*np.ndarray*) – the set of states over which we need to compute the Q function.

## Deep AC

**class A2CBuilder**

*policy\_params, actor\_optimizer, critic\_params, alg\_params, n\_steps\_per\_fit=5, preprocessors=None* Bases: *mushroom\_rl\_benchmark.builders.agent\_builder.AgentBuilder*

AgentBuilder for Advantage Actor Critic algorithm (A2C)

`__init__`

`policy_params, actor_optimizer, critic_params, alg_params, n_steps_per_fit=5, preprocessors=None`  
 Constructor.

**Parameters**

- **policy\_params** (*dict*) – parameters for the policy;
- **actor\_optimizer** (*dict*) – parameters for the actor optimizer;
- **critic\_params** (*dict*) – parameters for the critic;
- **alg\_params** (*dict*) – parameters for the algorithm;
- **n\_steps\_per\_fit** (*int, 5*) – number of steps per fit;
- **preprocessors** (*list, None*) – list of preprocessors.

**build**

`mdp_info` Build and return the AgentBuilder

**Parameters** `mdp_info` (*MDPInfo*) – information about the environment.

**compute\_Q**

`agent, states` Compute the Q Value for an AgentBuilder

**Parameters**

- **agent** (*Agent*) – the considered agent;
- **states** (*np.ndarray*) – the set of states over which we need to compute the Q function.

**classmethod default**

`actor_lr=0.0007, critic_lr=0.0007, eps_actor=0.003, eps_critic=1e-05, batch_size=64, max_grad_norm=0.5, ent_coeff=0.01, critic_network=<class 'mushroom_rl_benchmark.builders.network.a2c_network.A2CNetwork'>, n_features=64, preprocessors=None, use_cuda=False, get_default_dict=False` Create a default initialization for the specific AgentBuilder and return it

**class DDPGBuilder**

`policy_class, policy_params, actor_params, actor_optimizer, critic_params, alg_params, n_steps_per_fit=1`

Bases: `mushroom_rl_benchmark.builders.agent_builder.AgentBuilder`

AgentBuilder for Deep Deterministic Policy Gradient algorithm (DDPG)

`__init__`

`policy_class, policy_params, actor_params, actor_optimizer, critic_params, alg_params, n_steps_per_fit=1` Constructor.

**Parameters**

- **policy\_class** (*Policy*) – policy class;
- **policy\_params** (*dict*) – parameters for the policy;
- **actor\_params** (*dict*) – parameters for the actor;
- **actor\_optimizer** (*dict*) – parameters for the actor optimizer;
- **critic\_params** (*dict*) – parameters for the critic;
- **alg\_params** (*dict*) – parameters for the algorithm;
- **n\_steps\_per\_fit** (*int, 1*) – number of steps per fit.

**build**

*mdp\_info* Build and return the AgentBuilder

**Parameters** *mdp\_info* (*MDPInfo*) – information about the environment.

**compute\_Q**

*agent, states* Compute the Q Value for an AgentBuilder

**Parameters**

- **agent** (*Agent*) – the considered agent;
- **states** (*np.ndarray*) – the set of states over which we need to compute the Q function.

**classmethod default**

*actor\_lr=0.0001, actor\_network=<class 'mushroom\_rl\_benchmark.builders.network.ddpg\_network.DDPGActorNetwork'>, critic\_lr=0.001, critic\_network=<class 'mushroom\_rl\_benchmark.builders.network.ddpg\_network.DDPGCriticNetwork'>, initial\_replay\_size=500, max\_replay\_size=50000, batch\_size=64, n\_features=[80, 80], tau=0.001, use\_cuda=False, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

**class PPOBuilder**

*policy\_params, actor\_optimizer, critic\_params, alg\_params, n\_steps\_per\_fit=3000, preprocessors=None*  
Bases: *mushroom\_rl\_benchmark.builders.agent\_builder.AgentBuilder*

AgentBuilder for Proximal Policy Optimization algorithm (PPO)

**\_\_init\_\_**

*policy\_params, actor\_optimizer, critic\_params, alg\_params, n\_steps\_per\_fit=3000, preprocessors=None*  
Constructor.

**Parameters**

- **policy\_params** (*dict*) – parameters for the policy;
- **actor\_optimizer** (*dict*) – parameters for the actor optimizer;
- **critic\_params** (*dict*) – parameters for the critic;
- **alg\_params** (*dict*) – parameters for the algorithm;
- **n\_steps\_per\_fit** (*int, 3000*) – number of steps per fit;
- **preprocessors** (*list, None*) – list of preprocessors.

**build**

*mdp\_info* Build and return the AgentBuilder

**Parameters** *mdp\_info* (*MDPInfo*) – information about the environment.

**compute\_Q**

*agent, states* Compute the Q Value for an AgentBuilder

**Parameters**

- **agent** (*Agent*) – the considered agent;
- **states** (*np.ndarray*) – the set of states over which we need to compute the Q function.

**classmethod default**

*eps=0.2, ent\_coeff=0.0, n\_epochs\_policy=4, actor\_lr=0.0003, critic\_lr=0.0003, critic\_fit\_params=None, critic\_network=<class 'mushroom\_rl\_benchmark.builders.network.trpo\_network.TRPONetwork'>, lam=0.95, batch\_size=64, n\_features=32, n\_steps\_per\_fit=3000, std\_0=1.0, preprocessors=None,*

*use\_cuda=False, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

### class SACBuilder

*actor\_mu\_params, actor\_sigma\_params, actor\_optimizer, critic\_params, alg\_params, n\_q\_samples=100, n\_steps\_per\_fit=1, preprocessors=None* Bases: *mushroom\_rl\_benchmark.builders.agent\_builder.AgentBuilder*

AgentBuilder Soft Actor-Critic algorithm (SAC)

#### `__init__`

*actor\_mu\_params, actor\_sigma\_params, actor\_optimizer, critic\_params, alg\_params, n\_q\_samples=100, n\_steps\_per\_fit=1, preprocessors=None* Constructor.

#### Parameters

- **actor\_mu\_params** (*dict*) – parameters for actor mu;
- **actor\_sigma\_params** (*dict*) – parameters for actor sigma;
- **actor\_optimizer** (*dict*) – parameters for the actor optimizer;
- **critic\_params** (*dict*) – parameters for the critic;
- **alg\_params** (*dict*) – parameters for the algorithm;
- **n\_q\_samples** (*int, 100*) – number of samples to compute value function;
- **n\_steps\_per\_fit** (*int, 1*) – number of steps per fit;
- **preprocessors** (*list, None*) – list of preprocessors.

#### `build`

*mdp\_info* Build and return the AgentBuilder

Parameters **mdp\_info** (*MDPInfo*) – information about the environment.

#### `compute_Q`

*agent, states* Compute the Q Value for an AgentBuilder

#### Parameters

- **agent** (*Agent*) – the considered agent;
- **states** (*np.ndarray*) – the set of states over which we need to compute the Q function.

#### `classmethod default`

*actor\_lr=0.0003, actor\_network=<class 'mushroom\_rl\_benchmark.builders.network.sac\_network.SACActorNetwork'>, critic\_lr=0.0003, critic\_network=<class 'mushroom\_rl\_benchmark.builders.network.sac\_network.SACriticNetwork'>, initial\_replay\_size=64, max\_replay\_size=50000, n\_features=64, warmup\_transitions=100, batch\_size=64, tau=0.005, lr\_alpha=0.003, preprocessors=None, target\_entropy=None, use\_cuda=False, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

### class TD3Builder

*policy\_class, policy\_params, actor\_params, actor\_optimizer, critic\_params, alg\_params, n\_steps\_per\_fit=1* Bases: *mushroom\_rl\_benchmark.builders.agent\_builder.AgentBuilder*

AgentBuilder for Twin Delayed DDPG algorithm (TD3)

#### `__init__`

*policy\_class, policy\_params, actor\_params, actor\_optimizer, critic\_params, alg\_params, n\_steps\_per\_fit=1* Constructor.

**Parameters**

- **policy\_class** (*Policy*) – policy class;
- **policy\_params** (*dict*) – parameters for the policy;
- **actor\_params** (*dict*) – parameters for the actor;
- **actor\_optimizer** (*dict*) – parameters for the actor optimizer;
- **critic\_params** (*dict*) – parameters for the critic;
- **alg\_params** (*dict*) – parameters for the algorithm;
- **n\_steps\_per\_fit** (*int*, 1) – number of steps per fit.

**build**

*mdp\_info* Build and return the AgentBuilder

**Parameters** *mdp\_info* (*MDPInfo*) – information about the environment.

**compute\_Q**

*agent, states* Compute the Q Value for an AgentBuilder

**Parameters**

- **agent** (*Agent*) – the considered agent;
- **states** (*np.ndarray*) – the set of states over which we need to compute the Q function.

**classmethod default**

*actor\_lr=0.0001, actor\_network=<class 'mushroom\_rl\_benchmark.builders.network.td3\_network.TD3ActorNetwork'>, critic\_lr=0.001, critic\_network=<class 'mushroom\_rl\_benchmark.builders.network.td3\_network.TD3CriticNetwork'>, initial\_replay\_size=500, max\_replay\_size=50000, batch\_size=64, n\_features=[80, 80], tau=0.001, use\_cuda=False, get\_default\_dict=False* Create a default initialization for the specific AgentBuilder and return it

**class TRPOBuilder**

*policy\_params, critic\_params, alg\_params, n\_steps\_per\_fit=3000, preprocessors=None* Bases: *mushroom\_rl\_benchmark.builders.agent\_builder.AgentBuilder*

AgentBuilder for Trust Region Policy optimization algorithm (TRPO)

**\_\_init\_\_**

*policy\_params, critic\_params, alg\_params, n\_steps\_per\_fit=3000, preprocessors=None* Constructor.

**Parameters**

- **policy\_params** (*dict*) – parameters for the policy;
- **critic\_params** (*dict*) – parameters for the critic;
- **alg\_params** (*dict*) – parameters for the algorithm;
- **n\_steps\_per\_fit** (*int*, 3000) – number of steps per fit;
- **preprocessors** (*list*, *None*) – list of preprocessors.

**build**

*mdp\_info* Build and return the AgentBuilder

**Parameters** *mdp\_info* (*MDPInfo*) – information about the environment.

**compute\_Q**

*agent, states* Compute the Q Value for an AgentBuilder

### Parameters

- **agent** (*Agent*) – the considered agent;
- **states** (*np.ndarray*) – the set of states over which we need to compute the Q function.

#### classmethod default

*critic\_lr=0.0003, critic\_network=<class 'mushroom\_rl\_benchmark.builders.network.trpo\_network.TRPONetwork'>, max\_kl=0.01, ent\_coeff=0.0, lam=0.95, batch\_size=64, n\_features=32, critic\_fit\_params=None, n\_steps\_per\_fit=3000, n\_epochs\_line\_search=10, n\_epochs\_cg=100, cg\_damping=0.01, cg\_residual\_tol=1e-10, std\_0=1.0, preprocessors=None, use\_cuda=False, get\_default\_dict=False*  
 Create a default initialization for the specific AgentBuilder and return it

## 3.6 Networks

### class A2CNetwork

*input\_shape, output\_shape, n\_features, \*\*kwargs* Bases: `torch.nn.modules.module.Module`

#### `__init__`

*input\_shape, output\_shape, n\_features, \*\*kwargs* Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

#### forward

*state, \*\*kwargs* Defines the computation performed at every call.

Should be overridden by all subclasses.

---

**Note:** Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

---

### class DDPGCriticNetwork

*input\_shape, output\_shape, n\_features, \*\*kwargs* Bases: `torch.nn.modules.module.Module`

#### `__init__`

*input\_shape, output\_shape, n\_features, \*\*kwargs* Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

#### forward

*state, action* Defines the computation performed at every call.

Should be overridden by all subclasses.

---

**Note:** Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

---

### class DDPGActorNetwork

*input\_shape, output\_shape, \*\*kwargs* Bases: `torch.nn.modules.module.Module`

#### `__init__`

*input\_shape, output\_shape, \*\*kwargs* Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

**forward**

*state* Defines the computation performed at every call.

Should be overridden by all subclasses.

---

**Note:** Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

---

**class SACCriticNetwork**

*input\_shape, output\_shape, n\_features, \*\*kwargs* Bases: `torch.nn.modules.module.Module`

**\_\_init\_\_**

*input\_shape, output\_shape, n\_features, \*\*kwargs* Initializes internal `Module` state, shared by both `nn.Module` and `ScriptModule`.

**forward**

*state, action* Defines the computation performed at every call.

Should be overridden by all subclasses.

---

**Note:** Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

---

**class SACActorNetwork**

*input\_shape, output\_shape, n\_features, \*\*kwargs* Bases: `torch.nn.modules.module.Module`

**\_\_init\_\_**

*input\_shape, output\_shape, n\_features, \*\*kwargs* Initializes internal `Module` state, shared by both `nn.Module` and `ScriptModule`.

**forward**

*state* Defines the computation performed at every call.

Should be overridden by all subclasses.

---

**Note:** Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

---

**class TD3CriticNetwork**

*input\_shape, output\_shape, n\_features, \*\*kwargs* Bases: `torch.nn.modules.module.Module`

**\_\_init\_\_**

*input\_shape, output\_shape, n\_features, \*\*kwargs* Initializes internal `Module` state, shared by both `nn.Module` and `ScriptModule`.

**forward**

*state, action* Defines the computation performed at every call.

Should be overridden by all subclasses.

---

**Note:** Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks

---

while the latter silently ignores them.

---

**class TD3ActorNetwork**

*input\_shape, output\_shape, n\_features, \*\*kwargs* Bases: `torch.nn.modules.module.Module`

**\_\_init\_\_**

*input\_shape, output\_shape, n\_features, \*\*kwargs* Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

**forward**

*state* Defines the computation performed at every call.

Should be overridden by all subclasses.

---

**Note:** Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

---

**class TRPONetwork**

*input\_shape, output\_shape, n\_features, \*\*kwargs* Bases: `torch.nn.modules.module.Module`

**\_\_init\_\_**

*input\_shape, output\_shape, n\_features, \*\*kwargs* Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

**forward**

*state, \*\*kwargs* Defines the computation performed at every call.

Should be overridden by all subclasses.

---

**Note:** Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

---

## 3.7 Experiment

**exec\_run**

*agent\_builder, env\_builder, n\_epochs, n\_steps=None, n\_episodes=None, n\_steps\_test=None, n\_episodes\_test=None, seed=None, save\_agent=False, quiet=True, \*\*kwargs* Function that handles the execution of an experiment run.

**Parameters**

- **agent\_builder** (`AgentBuilder`) – agent builder to spawn an agent;
- **env\_builder** (`EnvironmentBuilder`) – environment builder to spawn an environment;
- **n\_epochs** (*int*) – number of epochs;
- **n\_steps** (*int, None*) – number of steps per epoch;
- **n\_episodes** (*int, None*) – number of episodes per epoch;
- **n\_steps\_test** (*int, None*) – number of steps for testing;

- **n\_episodes\_test** (*int*, *None*) – number of episodes for testing;
- **seed** (*int*, *None*) – the seed;
- **save\_agent** (*bool*, *False*) – select if the agent should be logged or not;
- **quiet** (*bool*, *True*) – select if run should print execution information.

**compute\_metrics**

*core*, *eval\_params*, *agent\_builder*, *env\_builder* Function to compute the metrics.

**Parameters**

- **eval\_params** (*dict*) – parameters for running the evaluation;
- **agent\_builder** (*AgentBuilder*) – the agent builder;
- **env\_builder** (*EnvironmentBuilder*) – environment builder to spawn an environment;

**print\_metrics**

*logger*, *epoch*, *J*, *R*, *V*, *E* Function that pretty prints the metrics on the standard output.

**Parameters**

- **logger** (*Logger*) – MushroomRL logger object;
- **epoch** (*int*) – the current epoch;
- **J** (*float*) – the current value of J;
- **R** (*float*) – the current value of R;
- **V** (*float*) – the current value of V;
- **E** (*float*) – the current value of E (Set None if not defined).

### 3.7.1 Slurm utilities

**aggregate\_results**

*res\_dir*, *res\_id*, *console\_log\_dir=None* Function to aggregate the benchmark results from running in SLURM mode.

**Parameters**

- **res\_dir** (*str*) – path to the result directory;
- **res\_id** (*str*) – log id of the result directory;
- **console\_log\_dir** (*str*, *None*) – path to be used to log console.

**make\_arguments**

*\*\*params* Create a script argument string from dictionary

**read\_arguments\_run**

*arg\_string=None* Parse the arguments for the run script.

**Parameters** **arg\_string** (*str*, *None*) – pass the argument string.

**read\_arguments\_aggregate**

*arg\_string=None* Parse the arguments for the aggregate script.

**Parameters** **arg\_string** (*str*, *None*) – pass the argument string.

**create\_slurm\_script**

*slurm\_path, slurm\_script\_name='slurm.sh', \*\*slurm\_params* Function to create a slurm script in a specific directory

**Parameters**

- **slurm\_path** (*str*) – path to locate the slurm script;
- **slurm\_script\_name** (*str, slurm.sh*) – name of the slurm script;
- **\*\*slurm\_params** – parameters for generating the slurm file content.

**Returns** The path to the slurm script.

**generate\_slurm**

*exp\_name, exp\_dir\_slurm, python\_file, gres=None, project\_name=None, n\_exp=1, max\_concurrent\_runs=None, memory=2000, hours=24, minutes=0, seconds=0* Function to generate the slurm file content.

**Parameters**

- **exp\_name** (*str*) – name of the experiment;
- **exp\_dir\_slurm** (*str*) – directory where the slurm log files are located;
- **python\_file** (*str*) – path to the python file that should be executed;
- **gres** (*str, None*) – request cluster resources. E.g. to add a GPU in the IAS cluster specify `gres='gpu:rtx2080:1'`;
- **project\_name** (*str, None*) – name of the slurm project;
- **n\_exp** (*int, 1*) – number of experiments in the slurm array;
- **max\_concurrent\_runs** (*int, None*) – maximum number of runs that should be executed in parallel on the SLURM cluster;
- **memory** (*int, 2000*) – memory limit in mega bytes (MB) for the slurm jobs;
- **hours** (*int, 24*) – maximum number of execution hours for the slurm jobs;
- **minutes** (*int, 0*) – maximum number of execution minutes for the slurm jobs;
- **seconds** (*int, 0*) – maximum number of execution seconds for the slurm jobs.

**Returns** The slurm script as string.

**to\_duration**

*hours, minutes, seconds*

## 3.8 Utils

**get\_init\_states**

*dataset* Get the initial states of a MushroomRL dataset

**Parameters** **dataset** (*Dataset*) – a MushroomRL dataset.

**extract\_arguments**

*args, method* Extract the arguments from a dictionary that fit to a methods parameters.

**Parameters**

- **args** (*dict*) – dictionary of arguments;
- **method** (*function*) – method for which the arguments should be extracted.

**object\_to\_primitive**

*obj* Converts an object into a string using the class name

**Parameters** *obj* – the object to convert.

**Returns** A string representing the object.

**dictionary\_to\_primitive**

*data* Function that converts a dictionary by transforming any objects inside into strings

**Parameters** *data* (*dict*) – the dictionary to convert.

**Returns** The converted dictionary.

**get\_mean\_and\_confidence**

*data* Compute the mean and 95% confidence interval

**Parameters** *data* (*np.ndarray*) – Array of experiment data of shape (n\_runs, n\_epochs).

**Returns** The mean of the dataset at each epoch along with the confidence interval.

**plot\_mean\_conf**

*data, ax, color='blue', line='-', facecolor=None, alpha=0.4, label=None* Method to plot mean and confidence interval for data on pyplot axes.

**build\_sweep\_list**

*algs, sweep\_conf, base\_name='c\_'* Build the sweep list, from a compact dictionary specification, for every considered algorithm.

**Parameters**

- **algs** (*list*) – list of algorithms to be considered;
- **sweep\_conf** (*dict*) – dictionary with a compact sweep configuration for every algorithm;
- **base\_name** (*str*, *'c\_'*) – base name for the sweep configuration.

**Returns** The sweep list to be used with the suite.

**build\_sweep\_dict**

*base\_name='c\_', \*\*kwargs* Build the sweep dictionary, from a set of variable specifications.

**Parameters**

- **base\_name** (*str*, *'c\_'*) – base name for the sweep configuration;
- **\*\*kwargs** – the parameter specifications for the sweep.

**Returns** The sweep dictionary, where the key is the sweep name and the value is a dictionary with the sweep parameters.

**generate\_sweep**

*base\_name='c\_', \*\*kwargs* Generator that returns tuples with sweep name and parameters

**Parameters**

- **base\_name** (*str*, *'c\_'*) – base name for the sweep configuration;
- **\*\*kwargs** – the parameter specifications for the sweep.

**generate\_sweep\_params**

*\*\*kwargs* Generator that returns sweep parameters

**Parameters** **\*\*kwargs** – the parameter specifications for the sweep.

### m

`mushroom_rl_benchmark.builders.actor_critic.classic_actor_critic.copdac_g`, 66  
`mushroom_rl_benchmark.builders.actor_critic.classic_actor_critic.stochastic_ac`, 65  
`mushroom_rl_benchmark.builders.actor_critic.deep_actor_critic.a2c`, 66  
`mushroom_rl_benchmark.builders.actor_critic.deep_actor_critic.ddpg`, 67  
`mushroom_rl_benchmark.builders.actor_critic.deep_actor_critic.ppo`, 68  
`mushroom_rl_benchmark.builders.actor_critic.deep_actor_critic.sac`, 69  
`mushroom_rl_benchmark.builders.actor_critic.deep_actor_critic.td3`, 69  
`mushroom_rl_benchmark.builders.actor_critic.deep_actor_critic.trpo`, 70  
`mushroom_rl_benchmark.builders.agent_builder`, 55  
`mushroom_rl_benchmark.builders.environment_builder`, 55  
`mushroom_rl_benchmark.builders.network.a2c_network`, 71  
`mushroom_rl_benchmark.builders.network.ddpg_network`, 71  
`mushroom_rl_benchmark.builders.network.sac_network`, 72  
`mushroom_rl_benchmark.builders.network.td3_network`, 72  
`mushroom_rl_benchmark.builders.network.trpo_network`, 73  
`mushroom_rl_benchmark.builders.policy_search.black_box_optimization`, 57  
`mushroom_rl_benchmark.builders.policy_search.policy_gradient`, 56  
`mushroom_rl_benchmark.builders.value.dqn.averaged_dqn`, 63  
`mushroom_rl_benchmark.builders.value.dqn.categorical_dqn`, 65  
`mushroom_rl_benchmark.builders.value.dqn.double_dqn`, 63  
`mushroom_rl_benchmark.builders.value.dqn.dqn`, 62  
`mushroom_rl_benchmark.builders.value.dqn.dueling_dqn`, 64  
`mushroom_rl_benchmark.builders.value.dqn.maxmin_dqn`, 64  
`mushroom_rl_benchmark.builders.value.dqn.noisy_dqn`, 64  
`mushroom_rl_benchmark.builders.value.dqn.prioritized_dqn`, 64  
`mushroom_rl_benchmark.builders.value.td.sarsa_lambda_continuous`, 62  
`mushroom_rl_benchmark.builders.value.td.td_finite`, 59  
`mushroom_rl_benchmark.builders.value.td.td_trace`, 61  
`mushroom_rl_benchmark.builders.value.td.true_online_sarsa_lambda`, 62  
`mushroom_rl_benchmark.core.experiment`, 48  
`mushroom_rl_benchmark.core.logger`, 50  
`mushroom_rl_benchmark.core.suite`, 46  
`mushroom_rl_benchmark.core.suite_visualizer`, 54

`mushroom_rl_benchmark.core.visualizer`, 53  
`mushroom_rl_benchmark.experiment.run`, 73  
`mushroom_rl_benchmark.experiment.slurm.aggregate_results`, 74  
`mushroom_rl_benchmark.experiment.slurm.arguments`, 74  
`mushroom_rl_benchmark.experiment.slurm.run_script`, 74  
`mushroom_rl_benchmark.experiment.slurm.slurm_script`, 74  
`mushroom_rl_benchmark.utils.plot`, 76  
`mushroom_rl_benchmark.utils.primitive`, 75  
`mushroom_rl_benchmark.utils.sweep`, 76  
`mushroom_rl_benchmark.utils.utils`, 75

## Symbols

\_\_init\_\_ () (*A2CBuilder* method), 67  
 \_\_init\_\_ () (*A2CNetwork* method), 71  
 \_\_init\_\_ () (*AgentBuilder* method), 55  
 \_\_init\_\_ () (*BBOBuilder* method), 57  
 \_\_init\_\_ () (*BenchmarkExperiment* method), 48  
 \_\_init\_\_ () (*BenchmarkLogger* method), 50  
 \_\_init\_\_ () (*BenchmarkSuite* method), 46  
 \_\_init\_\_ () (*BenchmarkSuiteVisualizer* method), 54  
 \_\_init\_\_ () (*BenchmarkVisualizer* method), 53  
 \_\_init\_\_ () (*COPDAC\_QBuilder* method), 66  
 \_\_init\_\_ () (*ConstrainedREPSBuilder* method), 59  
 \_\_init\_\_ () (*DDPGActorNetwork* method), 71  
 \_\_init\_\_ () (*DDPGBuilder* method), 67  
 \_\_init\_\_ () (*DDPGCriticNetwork* method), 71  
 \_\_init\_\_ () (*DQNBuilder* method), 63  
 \_\_init\_\_ () (*DoubleQLearningBuilder* method), 61  
 \_\_init\_\_ () (*EnvironmentBuilder* method), 55  
 \_\_init\_\_ () (*PGPEBuilder* method), 58  
 \_\_init\_\_ () (*PPOBuilder* method), 68  
 \_\_init\_\_ () (*PolicyGradientBuilder* method), 56  
 \_\_init\_\_ () (*QLearningBuilder* method), 60  
 \_\_init\_\_ () (*REPSBuilder* method), 59  
 \_\_init\_\_ () (*RWRBuilder* method), 58  
 \_\_init\_\_ () (*SACActorNetwork* method), 72  
 \_\_init\_\_ () (*SACBuilder* method), 69  
 \_\_init\_\_ () (*SACCriticNetwork* method), 72  
 \_\_init\_\_ () (*SARSABuilder* method), 60  
 \_\_init\_\_ () (*SarsaLambdaContinuousBuilder* method), 62  
 \_\_init\_\_ () (*SpeedyQLearningBuilder* method), 60  
 \_\_init\_\_ () (*StochasticACBuilder* method), 65  
 \_\_init\_\_ () (*TD3ActorNetwork* method), 73  
 \_\_init\_\_ () (*TD3Builder* method), 69  
 \_\_init\_\_ () (*TD3CriticNetwork* method), 72  
 \_\_init\_\_ () (*TDFiniteBuilder* method), 59  
 \_\_init\_\_ () (*TDTraceBuilder* method), 61  
 \_\_init\_\_ () (*TRPOBuilder* method), 70  
 \_\_init\_\_ () (*TRPONetwork* method), 73  
 \_\_init\_\_ () (*TrueOnlineSarsaLambdaBuilder* method), 62  
 \_\_init\_\_ () (*WeightedQLearningBuilder* method), 61

## A

*A2CBuilder* (class in *mushroom\_rl\_benchmark.builders.actor\_critic.deep\_actor\_critic.a2c*), 66  
*A2CNetwork* (class in *mushroom\_rl\_benchmark.builders.network.a2c\_network*), 71

`add_agent()` (*BenchmarkSuite method*), 47  
`add_environment()` (*BenchmarkSuite method*), 47  
`add_experiments()` (*BenchmarkSuite method*), 47  
`add_experiments_sweeps()` (*BenchmarkSuite method*), 47  
`add_sweep()` (*BenchmarkSuite method*), 48  
`AgentBuilder` (*class in mushroom\_rl\_benchmark.builders.agent\_builder*), 55  
`aggregate_results()` (*in module mushroom\_rl\_benchmark.experiment.slurm.aggregate\_results*), 74  
`alg_class` (*ConstrainedREPSBuilder attribute*), 59  
`alg_class` (*DoubleQLearningBuilder attribute*), 61  
`alg_class` (*eNACBuilder attribute*), 57  
`alg_class` (*GPOMDPBuilder attribute*), 57  
`alg_class` (*PGPEBuilder attribute*), 58  
`alg_class` (*QLambdaBuilder attribute*), 62  
`alg_class` (*QLearningBuilder attribute*), 60  
`alg_class` (*REINFORCEBuilder attribute*), 57  
`alg_class` (*REPSBuilder attribute*), 58  
`alg_class` (*RWRBuilder attribute*), 58  
`alg_class` (*SARSABuilder attribute*), 60  
`alg_class` (*SARSALambdaBuilder attribute*), 62  
`alg_class` (*SpeedyQLearningBuilder attribute*), 60  
`alg_class` (*WeightedQLearningBuilder attribute*), 61  
`AveragedDQNBuilder` (*class in mushroom\_rl\_benchmark.builders.value.dqn.averaged\_dqn*), 63

## B

`BBOBuilder` (*class in mushroom\_rl\_benchmark.builders.policy\_search.black\_box\_optimization*), 57  
`BenchmarkExperiment` (*class in mushroom\_rl\_benchmark.core.experiment*), 48  
`BenchmarkLogger` (*class in mushroom\_rl\_benchmark.core.logger*), 50  
`BenchmarkSuite` (*class in mushroom\_rl\_benchmark.core.suite*), 46  
`BenchmarkSuiteVisualizer` (*class in mushroom\_rl\_benchmark.core.suite\_visualizer*), 54  
`BenchmarkVisualizer` (*class in mushroom\_rl\_benchmark.core.visualizer*), 53  
`build()` (*A2CBuilder method*), 67  
`build()` (*AgentBuilder method*), 56  
`build()` (*AveragedDQNBuilder method*), 63  
`build()` (*BBOBuilder method*), 57  
`build()` (*CategoricalDQNBuilder method*), 65  
`build()` (*COPDAC\_QBuilder method*), 66  
`build()` (*DDPGBuilder method*), 67  
`build()` (*DoubleDQNBuilder method*), 63  
`build()` (*DQNBuilder method*), 63  
`build()` (*DuelingDQNBuilder method*), 64  
`build()` (*EnvironmentBuilder method*), 55  
`build()` (*MaxminDQNBuilder method*), 64  
`build()` (*NoisyDQNBuilder method*), 64  
`build()` (*PolicyGradientBuilder method*), 56  
`build()` (*PPOBuilder method*), 68  
`build()` (*PrioritizedDQNBuilder method*), 64  
`build()` (*SACBuilder method*), 69  
`build()` (*SarsaLambdaContinuousBuilder method*), 62  
`build()` (*StochasticACBuilder method*), 65  
`build()` (*TD3Builder method*), 70  
`build()` (*TDFiniteBuilder method*), 59  
`build()` (*TRPOBuilder method*), 70  
`build()` (*TrueOnlineSarsaLambdaBuilder method*), 62  
`build_sweep_dict()` (*in module mushroom\_rl\_benchmark.utils.sweep*), 76  
`build_sweep_list()` (*in module mushroom\_rl\_benchmark.utils.sweep*), 76

## C

`CategoricalDQNBuilder` (*class in mushroom\_rl\_benchmark.builders.value.dqn.categorical\_dqn*), 65  
`compute_metrics()` (*in module mushroom\_rl\_benchmark.experiment.run*), 74  
`compute_Q()` (*A2CBuilder method*), 67  
`compute_Q()` (*AgentBuilder method*), 56  
`compute_Q()` (*BBOBuilder method*), 58  
`compute_Q()` (*COPDAC\_QBuilder method*), 66  
`compute_Q()` (*DDPGBuilder method*), 68  
`compute_Q()` (*DoubleQLearningBuilder method*), 61  
`compute_Q()` (*DQNBuilder method*), 63  
`compute_Q()` (*PolicyGradientBuilder method*), 57

compute\_Q() (*PPOBuilder method*), 68  
 compute\_Q() (*SACBuilder method*), 69  
 compute\_Q() (*StochasticACBuilder method*), 65  
 compute\_Q() (*TD3Builder method*), 70  
 compute\_Q() (*TDFiniteBuilder method*), 59  
 compute\_Q() (*TRPOBuilder method*), 70  
 ConstrainedREPSBuilder (*class in mushroom\_rl\_benchmark.builders.policy\_search.black\_box\_optimization*), 59  
 COPDAC\_QBuilder (*class in mushroom\_rl\_benchmark.builders.actor\_critic.classic\_actor\_critic.copdac\_q*), 66  
 copy() (*AgentBuilder method*), 56  
 copy() (*EnvironmentBuilder method*), 55  
 create\_slurm\_script() (*in module mushroom\_rl\_benchmark.experiment.slurm.slurm\_script*), 74

## D

DDPGActorNetwork (*class in mushroom\_rl\_benchmark.builders.network.ddpg\_network*), 71  
 DDPGBuilder (*class in mushroom\_rl\_benchmark.builders.actor\_critic.deep\_actor\_critic.ddpg*), 67  
 DDPGCriticNetwork (*class in mushroom\_rl\_benchmark.builders.network.ddpg\_network*), 71  
 default() (*mushroom\_rl\_benchmark.builders.actor\_critic.classic\_actor\_critic.copdac\_q.COPDAC\_QBuilder class method*), 66  
 default() (*mushroom\_rl\_benchmark.builders.actor\_critic.classic\_actor\_critic.stochastic\_ac.StochasticACBuilder class method*), 65  
 default() (*mushroom\_rl\_benchmark.builders.actor\_critic.deep\_actor\_critic.a2c.A2CBuilder class method*), 67  
 default() (*mushroom\_rl\_benchmark.builders.actor\_critic.deep\_actor\_critic.ddpg.DDPGBuilder class method*), 68  
 default() (*mushroom\_rl\_benchmark.builders.actor\_critic.deep\_actor\_critic.ppo.PPOBuilder class method*), 68  
 default() (*mushroom\_rl\_benchmark.builders.actor\_critic.deep\_actor\_critic.sac.SACBuilder class method*), 69  
 default() (*mushroom\_rl\_benchmark.builders.actor\_critic.deep\_actor\_critic.td3.TD3Builder class method*), 70  
 default() (*mushroom\_rl\_benchmark.builders.actor\_critic.deep\_actor\_critic.trpo.TRPOBuilder class method*), 71  
 default() (*mushroom\_rl\_benchmark.builders.agent\_builder.AgentBuilder class method*), 56  
 default() (*mushroom\_rl\_benchmark.builders.policy\_search.black\_box\_optimization.BBOBuilder class method*), 58  
 default() (*mushroom\_rl\_benchmark.builders.policy\_search.black\_box\_optimization.ConstrainedREPSBuilder class method*), 59  
 default() (*mushroom\_rl\_benchmark.builders.policy\_search.black\_box\_optimization.PGPEBuilder class method*), 58  
 default() (*mushroom\_rl\_benchmark.builders.policy\_search.black\_box\_optimization.REPSBuilder class method*), 59  
 default() (*mushroom\_rl\_benchmark.builders.policy\_search.black\_box\_optimization.RWRBuilder class method*), 58  
 default() (*mushroom\_rl\_benchmark.builders.policy\_search.policy\_gradient.PolicyGradientBuilder class method*), 57  
 default() (*mushroom\_rl\_benchmark.builders.value.dqn.averaged\_dqn.AveragedDQNBuilder class method*), 63  
 default() (*mushroom\_rl\_benchmark.builders.value.dqn.categorical\_dqn.CategoricalDQNBuilder class method*), 65  
 default() (*mushroom\_rl\_benchmark.builders.value.dqn.dqn.DQNBuilder class method*), 63  
 default() (*mushroom\_rl\_benchmark.builders.value.dqn.dueling\_dqn.DuelingDQNBuilder class method*), 64  
 default() (*mushroom\_rl\_benchmark.builders.value.dqn.maxmin\_dqn.MaxminDQNBuilder class method*), 64  
 default() (*mushroom\_rl\_benchmark.builders.value.dqn.noisy\_dqn.NoisyDQNBuilder class method*), 64  
 default() (*mushroom\_rl\_benchmark.builders.value.dqn.prioritized\_dqn.PrioritizedDQNBuilder class method*), 64  
 default() (*mushroom\_rl\_benchmark.builders.value.td.sarsa\_lambda\_continuous.SarsaLambdaContinuousBuilder class method*), 62  
 default() (*mushroom\_rl\_benchmark.builders.value.td.td\_finite.TDFiniteBuilder class method*), 60  
 default() (*mushroom\_rl\_benchmark.builders.value.td.td\_finite.WeightedQLearningBuilder class method*), 61  
 default() (*mushroom\_rl\_benchmark.builders.value.td.td\_trace.TDTraceBuilder class method*), 61  
 default() (*mushroom\_rl\_benchmark.builders.value.td.true\_online\_sarsa\_lambda.TrueOnlineSarsaLambdaBuilder class method*), 62  
 dictionary\_to\_primitive() (*in module mushroom\_rl\_benchmark.utils.primitive*), 76  
 DoubleDQNBuilder (*class in mushroom\_rl\_benchmark.builders.value.dqn.double\_dqn*), 63  
 DoubleQLearningBuilder (*class in mushroom\_rl\_benchmark.builders.value.td.td\_finite*), 61  
 DQNBuilder (*class in mushroom\_rl\_benchmark.builders.value.dqn.dqn*), 62  
 DuelingDQNBuilder (*class in mushroom\_rl\_benchmark.builders.value.dqn.dueling\_dqn*), 64

## E

eNACBuilder (*class in mushroom\_rl\_benchmark.builders.policy\_search.policy\_gradient*), 57  
 EnvironmentBuilder (*class in mushroom\_rl\_benchmark.builders.environment\_builder*), 55  
 exec\_run() (*in module mushroom\_rl\_benchmark.experiment.run*), 73  
 exists\_best\_agent() (*BenchmarkLogger method*), 52  
 exists\_policy\_entropy() (*BenchmarkLogger method*), 52  
 exists\_stats() (*BenchmarkLogger method*), 52  
 exists\_value\_function() (*BenchmarkLogger method*), 52  
 extend\_and\_save\_entropy() (*BenchmarkExperiment method*), 50  
 extend\_and\_save\_J() (*BenchmarkExperiment method*), 50  
 extend\_and\_save\_R() (*BenchmarkExperiment method*), 50  
 extend\_and\_save\_V() (*BenchmarkExperiment method*), 50  
 extract\_arguments() (*in module mushroom\_rl\_benchmark.utils.utils*), 75

## F

forward() (*A2CNetwork method*), 71  
 forward() (*DDPGActorNetwork method*), 71

`forward()` (*DDPGCriticNetwork method*), 71  
`forward()` (*SACActorNetwork method*), 72  
`forward()` (*SACCriticNetwork method*), 72  
`forward()` (*TD3ActorNetwork method*), 73  
`forward()` (*TD3CriticNetwork method*), 72  
`forward()` (*TRPONetwork method*), 73  
`from_path()` (*mushroom\_rl\_benchmark.core.logger.BenchmarkLogger class method*), 53  
`from_path()` (*mushroom\_rl\_benchmark.core.visualizer.BenchmarkVisualizer class method*), 54

## G

`generate_slurm()` (*in module mushroom\_rl\_benchmark.experiment.slurm.slurm\_script*), 75  
`generate_sweep()` (*in module mushroom\_rl\_benchmark.utils.sweep*), 76  
`generate_sweep_params()` (*in module mushroom\_rl\_benchmark.utils.sweep*), 76  
`get_boxplot()` (*BenchmarkSuiteVisualizer method*), 54  
`get_entropy()` (*BenchmarkVisualizer method*), 53  
`get_figure_path()` (*BenchmarkLogger method*), 51  
`get_fit_params()` (*AgentBuilder method*), 55  
`get_init_states()` (*in module mushroom\_rl\_benchmark.utils.utils*), 75  
`get_J()` (*BenchmarkVisualizer method*), 53  
`get_log_dir()` (*BenchmarkLogger method*), 50  
`get_log_id()` (*BenchmarkLogger method*), 51  
`get_mean_and_confidence()` (*in module mushroom\_rl\_benchmark.utils.plot*), 76  
`get_params_path()` (*BenchmarkLogger method*), 51  
`get_path()` (*BenchmarkLogger method*), 51  
`get_preprocessors()` (*AgentBuilder method*), 56  
`get_R()` (*BenchmarkVisualizer method*), 53  
`get_report()` (*BenchmarkSuiteVisualizer method*), 54  
`get_report()` (*BenchmarkVisualizer method*), 54  
`get_V()` (*BenchmarkVisualizer method*), 53  
`GPOMDPBuilder` (*class in mushroom\_rl\_benchmark.builders.policy\_search.policy\_gradient*), 57

## I

`is_data_persisted` (*BenchmarkVisualizer attribute*), 53

## L

`load_agent_builder()` (*BenchmarkLogger method*), 52  
`load_best_agent()` (*BenchmarkLogger method*), 52  
`load_config()` (*BenchmarkLogger method*), 52  
`load_entropy()` (*BenchmarkLogger method*), 51  
`load_environment_builder()` (*BenchmarkLogger method*), 52  
`load_J()` (*BenchmarkLogger method*), 51  
`load_last_agent()` (*BenchmarkLogger method*), 52  
`load_R()` (*BenchmarkLogger method*), 51  
`load_stats()` (*BenchmarkLogger method*), 52  
`load_V()` (*BenchmarkLogger method*), 51

## M

`make_arguments()` (*in module mushroom\_rl\_benchmark.experiment.slurm.arguments*), 74  
`MaxminDQNBuilder` (*class in mushroom\_rl\_benchmark.builders.value.dqn.maxmin\_dqn*), 64  
`mushroom_rl_benchmark.builders.actor_critic.classic_actor_critic.copdac_q` (*module*), 66  
`mushroom_rl_benchmark.builders.actor_critic.classic_actor_critic.stochastic_ac` (*module*), 65  
`mushroom_rl_benchmark.builders.actor_critic.deep_actor_critic.a2c` (*module*), 66  
`mushroom_rl_benchmark.builders.actor_critic.deep_actor_critic.ddpg` (*module*), 67  
`mushroom_rl_benchmark.builders.actor_critic.deep_actor_critic.ppo` (*module*), 68  
`mushroom_rl_benchmark.builders.actor_critic.deep_actor_critic.sac` (*module*), 69  
`mushroom_rl_benchmark.builders.actor_critic.deep_actor_critic.td3` (*module*), 69  
`mushroom_rl_benchmark.builders.actor_critic.deep_actor_critic.trpo` (*module*), 70  
`mushroom_rl_benchmark.builders.agent_builder` (*module*), 55  
`mushroom_rl_benchmark.builders.environment_builder` (*module*), 55  
`mushroom_rl_benchmark.builders.network.a2c_network` (*module*), 71  
`mushroom_rl_benchmark.builders.network.ddpg_network` (*module*), 71  
`mushroom_rl_benchmark.builders.network.sac_network` (*module*), 72  
`mushroom_rl_benchmark.builders.network.td3_network` (*module*), 72  
`mushroom_rl_benchmark.builders.network.trpo_network` (*module*), 73  
`mushroom_rl_benchmark.builders.policy_search.black_box_optimization` (*module*), 57

mushroom\_rl\_benchmark.builders.policy\_search.policy\_gradient (*module*), 56  
 mushroom\_rl\_benchmark.builders.value.dqn.averaged\_dqn (*module*), 63  
 mushroom\_rl\_benchmark.builders.value.dqn.categorical\_dqn (*module*), 65  
 mushroom\_rl\_benchmark.builders.value.dqn.double\_dqn (*module*), 63  
 mushroom\_rl\_benchmark.builders.value.dqn.dqn (*module*), 62  
 mushroom\_rl\_benchmark.builders.value.dqn.dueling\_dqn (*module*), 64  
 mushroom\_rl\_benchmark.builders.value.dqn.maxmin\_dqn (*module*), 64  
 mushroom\_rl\_benchmark.builders.value.dqn.noisy\_dqn (*module*), 64  
 mushroom\_rl\_benchmark.builders.value.dqn.prioritized\_dqn (*module*), 64  
 mushroom\_rl\_benchmark.builders.value.td.sarsa\_lambda\_continuous (*module*), 62  
 mushroom\_rl\_benchmark.builders.value.td.td\_finite (*module*), 59  
 mushroom\_rl\_benchmark.builders.value.td.td\_trace (*module*), 61  
 mushroom\_rl\_benchmark.builders.value.td.true\_online\_sarsa\_lambda (*module*), 62  
 mushroom\_rl\_benchmark.core.experiment (*module*), 48  
 mushroom\_rl\_benchmark.core.logger (*module*), 50  
 mushroom\_rl\_benchmark.core.suite (*module*), 46  
 mushroom\_rl\_benchmark.core.suite\_visualizer (*module*), 54  
 mushroom\_rl\_benchmark.core.visualizer (*module*), 53  
 mushroom\_rl\_benchmark.experiment.run (*module*), 73  
 mushroom\_rl\_benchmark.experiment.slurm.aggregate\_results (*module*), 74  
 mushroom\_rl\_benchmark.experiment.slurm.arguments (*module*), 74  
 mushroom\_rl\_benchmark.experiment.slurm.run\_script (*module*), 74  
 mushroom\_rl\_benchmark.experiment.slurm.slurm\_script (*module*), 74  
 mushroom\_rl\_benchmark.utils.plot (*module*), 76  
 mushroom\_rl\_benchmark.utils.primitive (*module*), 75  
 mushroom\_rl\_benchmark.utils.sweep (*module*), 76  
 mushroom\_rl\_benchmark.utils.utils (*module*), 75

## N

NoisyDQNBuilder (*class in mushroom\_rl\_benchmark.builders.value.dqn.noisy\_dqn*), 64

## O

object\_to\_primitive () (*in module mushroom\_rl\_benchmark.utils.primitive*), 75

## P

PGPEBuilder (*class in mushroom\_rl\_benchmark.builders.policy\_search.black\_box\_optimization*), 58  
 plot\_counter (*BenchmarkSuiteVisualizer attribute*), 54  
 plot\_counter (*BenchmarkVisualizer attribute*), 53  
 plot\_mean\_conf () (*in module mushroom\_rl\_benchmark.utils.plot*), 76  
 PolicyGradientBuilder (*class in mushroom\_rl\_benchmark.builders.policy\_search.policy\_gradient*), 56  
 PPOBuilder (*class in mushroom\_rl\_benchmark.builders.actor\_critic.deep\_actor\_critic.ppo*), 68  
 print\_experiments () (*BenchmarkSuite method*), 48  
 print\_metrics () (*in module mushroom\_rl\_benchmark.experiment.run*), 74  
 PrioritizedDQNBuilder (*class in mushroom\_rl\_benchmark.builders.value.dqn.prioritized\_dqn*), 64

## Q

QLambdaBuilder (*class in mushroom\_rl\_benchmark.builders.value.td.td\_trace*), 62  
 QLearningBuilder (*class in mushroom\_rl\_benchmark.builders.value.td.td\_finite*), 60

## R

read\_arguments\_aggregate () (*in module mushroom\_rl\_benchmark.experiment.slurm.arguments*), 74  
 read\_arguments\_run () (*in module mushroom\_rl\_benchmark.experiment.slurm.arguments*), 74  
 REINFORCEBuilder (*class in mushroom\_rl\_benchmark.builders.policy\_search.policy\_gradient*), 57  
 REPSBuilder (*class in mushroom\_rl\_benchmark.builders.policy\_search.black\_box\_optimization*), 58  
 reset () (*BenchmarkExperiment method*), 49  
 resume () (*BenchmarkExperiment method*), 49  
 run () (*BenchmarkExperiment method*), 48  
 run () (*BenchmarkSuite method*), 48  
 run\_parallel () (*BenchmarkExperiment method*), 49  
 run\_sequential () (*BenchmarkExperiment method*), 49  
 run\_slurm () (*BenchmarkExperiment method*), 49  
 RWRBuilder (*class in mushroom\_rl\_benchmark.builders.policy\_search.black\_box\_optimization*), 58

## S

SACAActorNetwork (*class in mushroom\_rl\_benchmark.builders.network.sac\_network*), 72

SACBuilder (class in mushroom\_rl\_benchmark.builders.actor\_critic.deep\_actor\_critic.sac), 69  
SACCriticNetwork (class in mushroom\_rl\_benchmark.builders.network.sac\_network), 72  
SARSABuilder (class in mushroom\_rl\_benchmark.builders.value.td.td\_finite), 60  
SARSALambdaBuilder (class in mushroom\_rl\_benchmark.builders.value.td.td\_trace), 62  
SarsaLambdaContinuousBuilder (class in mushroom\_rl\_benchmark.builders.value.td.sarsa\_lambda\_continuous), 62  
save\_agent\_builder() (BenchmarkLogger method), 52  
save\_best\_agent() (BenchmarkLogger method), 52  
save\_boxplots() (BenchmarkSuiteVisualizer method), 54  
save\_builders() (BenchmarkExperiment method), 50  
save\_config() (BenchmarkLogger method), 52  
save\_entropy() (BenchmarkLogger method), 51  
save\_environment\_builder() (BenchmarkLogger method), 52  
save\_figure() (BenchmarkLogger method), 53  
save\_J() (BenchmarkLogger method), 51  
save\_last\_agent() (BenchmarkLogger method), 52  
save\_parameters() (BenchmarkSuite method), 48  
save\_params() (BenchmarkLogger method), 53  
save\_plot() (BenchmarkExperiment method), 50  
save\_plots() (BenchmarkSuite method), 48  
save\_R() (BenchmarkLogger method), 51  
save\_report() (BenchmarkVisualizer method), 54  
save\_reports() (BenchmarkSuiteVisualizer method), 54  
save\_stats() (BenchmarkLogger method), 52  
save\_V() (BenchmarkLogger method), 51  
set\_and\_save\_config() (BenchmarkExperiment method), 50  
set\_and\_save\_stats() (BenchmarkExperiment method), 50  
set\_eval\_mode() (AgentBuilder method), 56  
set\_eval\_mode() (COPDAC\_QBuilder method), 66  
set\_eval\_mode() (DQNBuilder method), 63  
set\_eval\_mode() (EnvironmentBuilder static method), 55  
set\_eval\_mode() (TDFiniteBuilder method), 60  
set\_log\_dir() (BenchmarkLogger method), 50  
set\_log\_id() (BenchmarkLogger method), 51  
set\_preprocessors() (AgentBuilder method), 56  
show\_agent() (BenchmarkVisualizer method), 54  
show\_plot() (BenchmarkExperiment method), 50  
show\_plots() (BenchmarkSuite method), 48  
show\_report() (BenchmarkVisualizer method), 54  
show\_reports() (BenchmarkSuiteVisualizer method), 55  
SpeedyQLearningBuilder (class in mushroom\_rl\_benchmark.builders.value.td.td\_finite), 60  
start\_timer() (BenchmarkExperiment method), 50  
StochasticACBuilder (class in mushroom\_rl\_benchmark.builders.actor\_critic.classic\_actor\_critic.stochastic\_ac), 65  
stop\_timer() (BenchmarkExperiment method), 50

## T

TD3ActorNetwork (class in mushroom\_rl\_benchmark.builders.network.td3\_network), 73  
TD3Builder (class in mushroom\_rl\_benchmark.builders.actor\_critic.deep\_actor\_critic.td3), 69  
TD3CriticNetwork (class in mushroom\_rl\_benchmark.builders.network.td3\_network), 72  
TDFiniteBuilder (class in mushroom\_rl\_benchmark.builders.value.td.td\_finite), 59  
TDTraceBuilder (class in mushroom\_rl\_benchmark.builders.value.td.td\_trace), 61  
to\_duration() (in module mushroom\_rl\_benchmark.experiment.slurm.slurm\_script), 75  
TRPOBuilder (class in mushroom\_rl\_benchmark.builders.actor\_critic.deep\_actor\_critic.trpo), 70  
TRPONetwork (class in mushroom\_rl\_benchmark.builders.network.trpo\_network), 73  
TrueOnlineSarsaLambdaBuilder (class in mushroom\_rl\_benchmark.builders.value.td.true\_online\_sarsa\_lambda), 62

## W

WeightedQLearningBuilder (class in mushroom\_rl\_benchmark.builders.value.td.td\_finite), 61