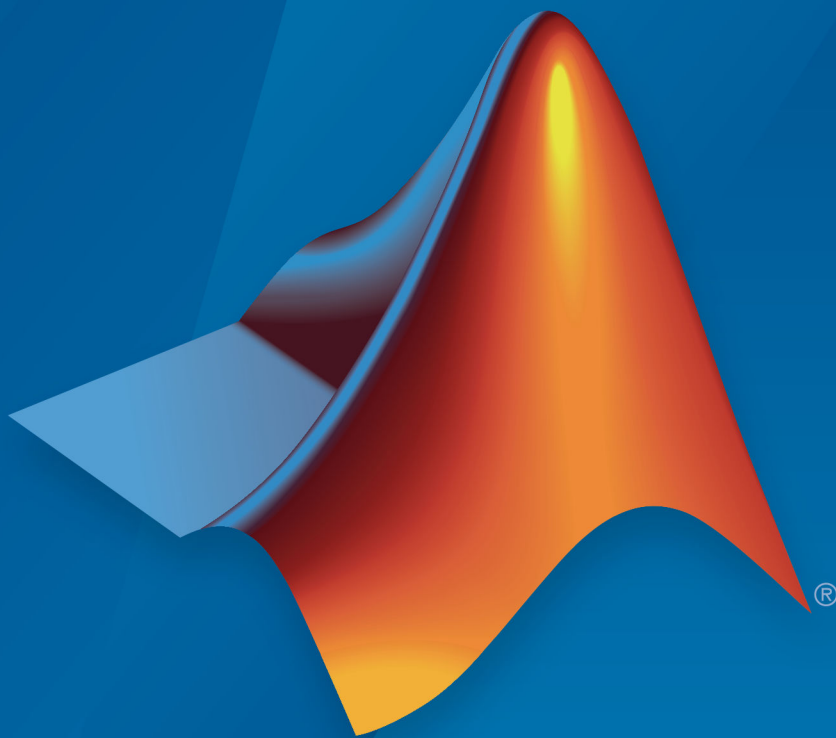


Reinforcement Learning Toolbox™ Release Notes



MATLAB®

How to Contact MathWorks



Latest news: www.mathworks.com
Sales and services: www.mathworks.com/sales_and_services
User community: www.mathworks.com/matlabcentral
Technical support: www.mathworks.com/support/contact_us



Phone: 508-647-7000



The MathWorks, Inc.
1 Apple Hill Drive
Natick, MA 01760-2098

Reinforcement Learning Toolbox™ Release Notes

© COPYRIGHT 2019 by The MathWorks, Inc.

The software described in this document is furnished under a license agreement. The software may be used or copied only under the terms of the license agreement. No part of this manual may be photocopied or reproduced in any form without prior written consent from The MathWorks, Inc.

FEDERAL ACQUISITION: This provision applies to all acquisitions of the Program and Documentation by, for, or through the federal government of the United States. By accepting delivery of the Program or Documentation, the government hereby agrees that this software or documentation qualifies as commercial computer software or commercial computer software documentation as such terms are used or defined in FAR 12.212, DFARS Part 227.72, and DFARS 252.227-7014. Accordingly, the terms and conditions of this Agreement and only those rights specified in this Agreement, shall pertain to and govern the use, modification, reproduction, release, performance, display, and disclosure of the Program and Documentation by the federal government (or other entity acquiring for or through the federal government) and shall supersede any conflicting contractual terms or conditions. If this License fails to meet the government's needs or is inconsistent in any respect with federal procurement law, the government agrees to return the Program and Documentation, unused, to The MathWorks, Inc.

Trademarks

MATLAB and Simulink are registered trademarks of The MathWorks, Inc. See www.mathworks.com/trademarks for a list of additional trademarks. Other product or brand names may be trademarks or registered trademarks of their respective holders.

Patents

MathWorks products are protected by one or more U.S. patents. Please see www.mathworks.com/patents for more information.

R2019b

Parallel Agent Simulation: Verify trained policies by running multiple agent simulations in parallel	1-2
PPO Agent: Train policies using proximal policy optimization algorithm for improved training stability	1-2
New Examples: Train reinforcement learning policies for applications such as robotics, automated driving, and control design	1-2

R2019a

Reinforcement Learning Algorithms: Train policies using DQN, DDPG, A2C, and other algorithms	2-2
Environment Modeling: Create MATLAB and Simulink environment models and provide observation and reward signals for training policies	2-2
Policy and Value Function Representation: Parameterize policies using deep neural networks, linear basis functions, and look-up tables	2-3
Interoperability: Import policies from Keras and the ONNX model format	2-3

Training Acceleration: Parallelize environment simulations and gradient calculations on GPUs and multicore CPUs for policy training	2-3
Code Generation: Deploy trained policies to embedded devices through automatic code generation for CPUs and GPUs ...	2-4
Reference Examples: Implement controllers using reinforcement learning for automated driving and robotics applications	2-4

R2019b

Version: 1.1

New Features

Bug Fixes

Parallel Agent Simulation: Verify trained policies by running multiple agent simulations in parallel

You can now run multiple agent simulations in parallel. If you have Parallel Computing Toolbox™ software, you can run parallel simulations on multicore computers. If you have MATLAB® Parallel Server™ software, you can run parallel simulations on computer clusters or cloud resources. For more information, see `rlSimulationOptions`.

PPO Agent: Train policies using proximal policy optimization algorithm for improved training stability

You can now train policies using proximal policy optimization (PPO). This algorithm is a type of policy gradient training that alternates between sampling data through environmental interaction and optimizing a clipped surrogate objective function using stochastic gradient descent. The clipped surrogate objective function improves training stability by limiting the size of the policy change at each step.

For more information on PPO agents, see “Proximal Policy Optimization Agents”.

New Examples: Train reinforcement learning policies for applications such as robotics, automated driving, and control design

The following new examples show how to train policies for robotics, automated driving, and control design:

- “Quadruped Robot Locomotion Using DDPG Agent”
- “Imitate MPC Controller for Lane Keep Assist”

R2019a

Version: 1.0

New Features

Reinforcement Learning Algorithms: Train policies using DQN, DDPG, A2C, and other algorithms

Using Reinforcement Learning Toolbox™ software, you can train policies using several standard reinforcement learning algorithms. You can create agents to train policies for the following:

- Q-learning
- SARSA
- Deep Q-networks (DQN)
- Deep deterministic policy gradients (DDPG)
- Policy gradient (PG)
- Advantage actor-critic (A2C)

You can also train policies using other algorithms by creating a custom agent.

For more information on creating and training agents, see Reinforcement Learning Agents and Train Reinforcement Learning Agents.

Environment Modeling: Create MATLAB and Simulink environment models and provide observation and reward signals for training policies

In a reinforcement learning scenario, the environment models the dynamics and system behavior with which the agent interacts. To define an environment model, you specify the following:

- Action and observation signals that the agent uses to interact with the environment.
- Reward signal that the agent uses to measure its success.
- Environment dynamic behavior.

You can model your environment using MATLAB and Simulink®. For more information, see Create MATLAB Environments for Reinforcement Learning and Create Simulink Environments for Reinforcement Learning

Policy and Value Function Representation: Parameterize policies using deep neural networks, linear basis functions, and look-up tables

Reinforcement Learning Toolbox software provides objects for actor and critic representations. The actor represents the policy that selects the action to take. The critic represents the value function that estimates the value of the current policy. Depending on your application and selected agent, you can define policy and value functions using deep neural networks, linear basis functions, or look-up tables. For more information, see [Create Policy and Value Function Representations](#).

Interoperability: Import policies from Keras and the ONNX model format

You can import existing deep neural network policies and value functions from other deep learning frameworks, such as Keras and the ONNX™ format. For more information, see [Import Policy and Value Function Representations](#).

Training Acceleration: Parallelize environment simulations and gradient calculations on GPUs and multicore CPUs for policy training

You can accelerate policy training by running parallel training simulations. If you have:

- Parallel Computing Toolbox software, you can run parallel simulations on multicore computers
- MATLAB Parallel Server software, you can run parallel simulations on computer clusters or cloud resources

You can also speed up deep neural network training and inference with high-performance NVIDIA® GPUs.

For more information, see [Train Reinforcement Learning Agents](#).

Code Generation: Deploy trained policies to embedded devices through automatic code generation for CPUs and GPUs

Once you have trained your reinforcement learning policy, you can generate code for policy deployment. You can generate optimized CUDA[®] code using GPU Coder[™] and C/C++ code using MATLAB Coder[™].

You can deploy trained policies as C/C++ shared libraries, Microsoft[®] .NET Framework assemblies, Java[®] classes, and Python[®] packages.

For more information, see [Deploy Trained Reinforcement Learning Policies](#).

Reference Examples: Implement controllers using reinforcement learning for automated driving and robotics applications

This release includes the following examples on training reinforcement learning policies for robotics and automated driving applications:

- Train DDPG Agent to Control Flying Robot
- Train Biped Robot to Walk Using DDPG Agent
- Train DQN Agent for Lane Keeping Assist
- Train DDPG Agent for Adaptive Cruise Control
- Train DDPG Agent for Path Following Control