



Intelligent Agent Development with Deep Reinforcement Learning and OpenAI Gym



Intelligent Agent Development with Deep Reinforcement Learning and OpenAI Gym

Course Overview:

The course is an immersive, hands-on training designed for professionals who wish to build AI systems using OpenAI Gym and deep reinforcement learning techniques. Based on the comprehensive book *Hands-On Intelligent Agents with OpenAI Gym*, this course offers a step-by-step practical journey through developing intelligent agents that solve real-world tasks such as game playing, robotics simulation, and autonomous driving. Key topics include Q-learning, Deep Q-Learning, experience replay, actor-critic methods, and environment customisation. Covering essential platforms like PyTorch, TensorBoard, CARLA, Roboschool, Gym-Retro, and MuJoCo, participants will acquire the skills to implement agents for both discrete and continuous action spaces.

Target Audience:

- AI/ML Engineers and Developers
- Robotics Engineers
- Data Scientists interested in RL
- Software Engineers exploring AI agents
- Game Developers

Targeted Organizational Departments:

- AI Research and Development Units
- Robotics and Automation Teams
- Innovation Labs
- Software Engineering Departments
- Simulation and Gaming Divisions

Targeted Industries:

- Automotive autonomous vehicles
- Robotics and Industrial Automation
- Gaming and Simulation
- Aerospace and Defence
- Healthcare Tech for training intelligent diagnostics agents

Course Offerings:

By the end of this course, participants will be able to:

- Set up and use OpenAI Gym and custom environments
- Apply Q-learning and Deep Q-learning using PyTorch
- Train agents using experience replay and epsilon-greedy policies
- Customise gym environments, including CARLA and MuJoCo
- Visualise training progress with TensorBoard
- Understand and apply policy gradients, actor-critic, PPO, and Rainbow RL
- Build and test agents on Atari games and Gym-Retro environments
- Monitor and optimise performance with reward shaping and preprocessing techniques

Training Methodology:

This course employs an applied, project-based methodology combining theoretical foundations with real-world practice. Learners will engage in interactive tutorials, group-based agent-building exercises, live demonstrations, and guided reinforcement learning projects. Emphasis is placed on practical implementation using PyTorch, JSON config files, CUDA acceleration, and OpenAI Gym. Case studies on Mountain Car, Cart Pole, Atari games, and CARLA simulations will illustrate key learning principles. Feedback sessions, breakout discussions, and reflective reviews ensure retention and hands-on mastery.

Course Toolbox:

- OpenAI Gym Environments Library
- PyTorch Deep Learning Framework
- Conda and CUDA Setup Guides
- TensorBoard for monitoring
- JSON templates for hyperparameters
- Atari and Gym-Retro emulators
- CARLA autonomous driving simulator
- Sample agent architectures DQN, PPO, DDPG, Rainbow
- Pre-built notebooks and implementation guides

Course Agenda:

Day 1: Foundations of Intelligent Agents & Reinforcement Learning

- **Topic 1:** Introduction to Intelligent Agents and Learning Environments
- **Topic 2:** Exploring the Capabilities and Interface of OpenAI Gym
- **Topic 3:** Categories of Gym Tasks: From Classic Control to Robotics
- **Topic 4:** Setting Up Your Python, Conda, CUDA, and PyTorch Environments
- **Topic 5:** Deep Dive into Reinforcement Learning and MDPs
- **Topic 6:** Understanding the Policy, Value Functions, and Exploration Strategies
- **Reflection & Review:** Fundamentals of AI Agents and Environment Interaction



Day 2: Hands-On with Q-Learning and Deep Q-Learning

- **Topic 1:** Solving the Mountain Car Problem Using Q-Learning
- **Topic 2:** Implementing Q-learning with NumPy and Hyperparameter Tuning
- **Topic 3:** Transition to Deep Q-Learning using PyTorch
- **Topic 4:** Applying Experience Replay and Epsilon-Greedy Policies
- **Topic 5:** Stabilizing Learning with Target Networks
- **Topic 6:** Visualizing Agent Performance Using TensorBoard
- **Reflection & Review:** Comparing Traditional and Deep Q-Learning Approaches

Day 3: Custom Environments and Real-World Applications

- **Topic 1:** Creating Custom Gym Environments with Templates and Registration
- **Topic 2:** Building the CARLA Driving Simulator as a Gym-Compatible Environment
- **Topic 3:** Implementing Reset, Step Functions, and Sensor Integration
- **Topic 4:** Managing Discrete vs Continuous Action Spaces in CARLA
- **Topic 5:** Real-Time Testing and Visualization of Simulation-Based Environments
- **Topic 6:** Techniques for Accessing and Using Camera/Sensor Data
- **Reflection & Review:** Environment Design for Reinforcement Learning Agents

Day 4: Advanced Agents with Actor-Critic Algorithms

- **Topic 1:** Fundamentals of Policy Gradients and Actor-Critic Architectures
- **Topic 2:** Implementing n-Step Advantage Actor-Critic Algorithms
- **Topic 3:** Designing Actor and Critic Networks for Autonomous Agents
- **Topic 4:** Logging, Monitoring, and Saving Model Progress
- **Topic 5:** Training Actor-Critic Agents in the CARLA Simulator
- **Topic 6:** Exploring Synchronous vs Asynchronous Implementations
- **Reflection & Review:** From Theory to Practice in Actor-Critic Training

Day 5: The Learning Landscape - PPO, DDPG, Rainbow & Beyond

- **Topic 1:** Proximal Policy Optimization PPO – Concepts and Use Cases
- **Topic 2:** Deep Deterministic Policy Gradient DDPG and Continuous Control
- **Topic 3:** The Rainbow Algorithm: Integrating Value-Based Enhancements
- **Topic 4:** Implementing Prioritized Replay, Dueling Nets, and Distributional RL
- **Topic 5:** Roboschool, Gym-Retro, DeepMind Lab, and StarCraft II Environments
- **Topic 6:** Comparative Insights Across Algorithms and Environment Suites
- **Reflection & Review:** Capstone Discussion on Agent Development and Deployment

FAQ:



What specific qualifications or prerequisites are needed for participants before enrolling in the course?

A working knowledge of Python and basic understanding of machine learning principles is recommended. Familiarity with NumPy and neural networks will help but is not mandatory.

How long is each day's session, and is there a total number of hours required for the entire course?

Each day's session is generally structured to last around 4-5 hours, with breaks and interactive activities included. The total course duration spans five days, approximately 20-25 hours of instruction.

Why does Deep Q-Learning use a target network and experience replay?

Target networks stabilise learning by keeping a fixed Q-target during updates. Experience replay improves sample efficiency and breaks temporal correlations in the training data, which helps avoid divergence in Q-learning.

How This Course is Different from Other Intelligent Agent Development Courses:

Unlike generic AI courses, this program is uniquely grounded in the proven methodologies and real-world examples from the *Hands-On Intelligent Agents with OpenAI Gym* book. It emphasises practical, code-level implementations of OpenAI Gym tutorial-based environments like Mountain Car and Cart Pole, uses PyTorch RL agent implementation techniques, and incorporates TensorBoard for reinforcement learning progress visualisation. By covering a diverse algorithm landscape, including Rainbow RL, PPO, and DDPG, it ensures a holistic skill set.



Training Course Categories

Agile PM and Project Management Training Courses

Certified Courses By International Bodies

Communication and Public Relations Training Courses

Data Analytics Training and Data Science Courses

Environment & Sustainability Training Courses

Finance and Accounting Training Courses

Governance, Risk and Compliance Training Courses

Human Resources Training and Development Courses

IT Security Training & IT Training Courses

Leadership and Management Training Courses

Legal Training, Procurement and Contracting Courses

Maintenance Training and Engineering Training Courses

Marketing, Customer Relations, and Sales Courses

Occupational Health, Safety and Security Training Courses

Personal & Self-Development Training Courses

Quality and Operations Management Training Courses

Secretarial and Administration Training Courses



Training Cities

Accra Ghana	Al Jubail Saudi Arabia	Amman Jordan
Amsterdam Netherlands	Athens Greece	Baku Azerbaijan
Bangkok Thailand	Barcelona Spain	Berlin Germany
Cairo Egypt	Cape town South Africa	Casablanca Morocco
Chicago USA	Doha Qatar	Dubai UAE
Frankfurt Germany	Geneva Switzerland	Istanbul Turkey
Jakarta Indonesia	Johannesburg South Africa	Kuala Lumpur Malaysia
Kuwait Kuwait	Langkawi Malaysia	Lisbon Portugal
London UK	Madrid Spain	Manama Bahrain
Marbella Spain	Milan Italy	Montreux Switzerland
Munich Germany	Muscat Oman	Nairobi Kenya
New York USA	Paris France	Phuket Thailand
Porto Portugal	Prague Czech Republic	Riyadh Saudi Arabia



Training Cities

Rome

Italy

San Diego

USA

Seoul

South Korea

Sharm El-Sheikh

Egypt

Singapore

Singapore

Tashkent

Uzbekistan

Tbilisi

Georgia

Tokyo

Japan

Toronto

Canada

Trabzon

Turkey

Vienna

Austria

Zanzibar

Tanzania

Zoom

Online Training