Advanced Reinforcement Learning for Autonomous Systems Training Course



30 Nov - 04 Dec 2026

Boston

7500 € (Per Person)

Ref: #AI3523_173003

Course Introduction / Overview:

This intensive training course provides a comprehensive exploration of advanced reinforcement learning (RL) techniques specifically tailored for the development of autonomous systems. In an era where intelligent automation is revolutionizing industries, mastering the principles of how agents learn to make optimal decisions is paramount. This program moves beyond theoretical concepts to focus on the practical implementation of cutting-edge algorithms that power everything from self-driving cars to sophisticated robotic manipulators. As detailed by pioneers like Richard S. Sutton in his seminal work, "Reinforcement Learning: An Introduction," the field offers a powerful framework for solving complex sequential decision-making problems. BIG BEN Training Center has designed this curriculum to bridge the gap between academic research and real-world engineering challenges. Participants will delve into deep reinforcement learning, policy optimization, and model-based methods, gaining the skills to build, train, and deploy robust autonomous agents capable of operating effectively in dynamic and uncertain environments. The course emphasizes a hands-on approach, ensuring that learners can confidently apply these advanced AI control systems to drive innovation within their organizations.

Target Audience / This training course is suitable for:

- AI and Machine Learning Engineers.
- Robotics Engineers and Specialists.
- Autonomous Systems Developers.
- Control Systems Engineers.
- Research Scientists in AI and Robotics.
- Software Developers working on autonomous vehicles or drones.
- Data Scientists transitioning into AI-driven control systems.
- Technical Project Managers overseeing AI and automation projects.
- Academics and postgraduate students specializing in artificial intelligence.

Target Sectors and Industries:

- Automotive and Autonomous Vehicles.
- Aerospace and Defense.
- Robotics and Industrial Automation.
- Logistics and Supply Chain Management.
- Manufacturing and Smart Factories.
- Drones and Unmanned Aerial Vehicle (UAV) Technology.
- Healthcare Technology and Medical Robotics.
- Government agencies involved in technology and defense research.
- Energy and Utilities for automated inspection and maintenance.

Target Organizations Departments:

- Research and Development (R&D).
- Engineering and Product Development.
- Innovation and Technology Labs.
- Software Engineering and IT.
- Data Science and Analytics.
- Operations and Automation.
- Quality Assurance and Testing for Autonomous Systems.
- Strategic Planning and Technology Integration.

Course Offerings:

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

- Master the core principles of Markov Decision Processes (MDPs) and Bellman equations.
- Implement foundational value-based methods like Q-learning and Deep Q-Networks (DQN).
- Develop and train agents using advanced policy gradient and actor-critic algorithms.
- Apply state-of-the-art algorithms such as PPO and SAC to complex control tasks.
- Design reward functions that effectively guide agent behavior in autonomous systems.
- Build and test RL agents in simulated environments for robotics and navigation.
- Understand the challenges and techniques for transferring learned policies from simulation to reality.
- Analyze the principles of safe reinforcement learning to ensure system reliability.
- Evaluate the ethical implications and considerations in deploying autonomous RL agents.
- Architect complete reinforcement learning pipelines for industrial applications.

Course Methodology:

The training methodology at BIG BEN Training Center is designed to be immersive, practical, and highly interactive, ensuring participants gain tangible skills. This course moves beyond traditional lectures by integrating hands-on coding labs where participants will implement reinforcement learning algorithms from the ground up. Each module is supported by real-world case studies, such as training a virtual drone for autonomous navigation or teaching a robotic arm to perform a manipulation task. We foster a collaborative learning environment through group projects and peer-to-peer feedback sessions, allowing participants to tackle complex problems collectively and learn from diverse perspectives. Expert instructors facilitate discussions, provide personalized guidance, and ensure that theoretical concepts are always linked to practical application. The curriculum incorporates a blend of individual exercises, team-based challenges, and interactive Q&A sessions to cater to various learning styles. This approach guarantees that participants not only understand the theory behind advanced reinforcement learning but also develop the confidence and competence to apply it to solve real-world challenges in autonomous systems.

Course Agenda (Course Units):

Unit One Foundations of Reinforcement Learning for Autonomy

- Introduction to Autonomous Systems and Decision-Making.
- The Reinforcement Learning Framework and Key Terminology.
- Modeling Problems with Markov Decision Processes (MDPs).
- The Bellman Equations for Value and Policy Evaluation.
- Dynamic Programming Approaches with Policy and Value Iteration.
- Monte Carlo Methods for Model-Free Prediction and Control.
- Temporal-Difference (TD) Learning, SARSA, and Q-Learning.

Unit Two Deep Reinforcement Learning with Value-Based Methods

- Limitations of Traditional RL and the Need for Deep Learning.
- Introduction to Neural Networks as Function Approximators.
- Deep Q-Networks (DQN) for High-Dimensional State Spaces.
- Improving DQN with Double DQN and Dueling Architectures.
- Experience Replay and its Role in Stabilizing Training.
- Applying DQN to Classic Control and Autonomous Navigation Problems.
- Analysis of Value-Based Methods' Strengths and Weaknesses.

Unit Three Advanced Policy-Based and Actor-Critic Methods

- Introduction to Policy Gradient Methods.
- The REINFORCE Algorithm and its Implementation.
- Actor-Critic Architectures for Stable and Efficient Learning.
- Advantage Actor-Critic (A2C) and Asynchronous Advantage Actor-Critic (A3C).
- Deep Deterministic Policy Gradient (DDPG) for Continuous Control.
- Soft Actor-Critic (SAC) for Maximum Entropy RL.
- Proximal Policy Optimization (PPO) as a State-of-the-Art Algorithm.

Unit Four Practical Implementation for Autonomous Systems

- Setting up Simulation Environments for RL (e.g., Gym, MuJoCo).
- Designing Effective Reward Functions for Complex Tasks.
- Application in Autonomous Navigation and Pathfinding.
- Application in Robotic Manipulation and Grasping.
- The Challenge of Simulation-to-Reality (Sim-to-Real) Transfer.
- Techniques for Bridging the Reality Gap, including Domain Randomization.
- Introduction to Multi-Agent Reinforcement Learning (MARL) for cooperative systems.

Unit Five Safety, Ethics, and Future Directions in RL

- The Importance of Safety and Reliability in Autonomous Systems.
- Introduction to Safe Reinforcement Learning Techniques.
- Constrained MDPs and Shielding Mechanisms.
- Explain ability and Interpretability of RL Policies.
- Ethical Considerations and Bias in Reward Design.
- Inverse Reinforcement Learning (IRL) for Learning from Demonstration.
- Current Research Trends and the Future of RL in Autonomy.

FAQ:

Qualifications required for registering to this course?

There are no requirements.

How long is each daily session, and what is the total number of training hours for the course?

This training course spans five days, with daily sessions ranging between 4 to 5 hours, including breaks and interactive activities, bringing the total duration to 20 - 25 training hours.

Something to think about:

As autonomous systems become more integrated into society, how can we design reinforcement learning reward functions that align with complex human values and prevent unintended negative consequences?

What unique qualities does this course offer compared to other courses?

This course distinguishes itself by moving beyond foundational theory to focus squarely on the application of advanced, state-of-the-art reinforcement learning algorithms to real-world autonomous systems. While many programs cover Q-learning and basic policy gradients, our curriculum dedicates significant time to modern techniques like Proximal Policy Optimization (PPO) and Soft Actor-Critic (SAC), which are critical for solving complex, continuous control problems in robotics and autonomous navigation. A key differentiator is the dedicated unit on the simulation-to-reality (sim-to-real) gap, a crucial and often-overlooked challenge in deploying RL agents. We provide practical strategies for training agents in simulation and successfully transferring their learned skills to physical hardware. Furthermore, the course integrates a vital discussion on safety and ethics, equipping participants not just with technical skills but also with the critical thinking needed to build responsible and reliable autonomous systems. The emphasis is on a holistic, implementation-focused learning experience that prepares engineers and scientists to tackle the entire lifecycle of an RL project, from problem formulation and algorithm selection to deployment and ethical consideration.