Training Course on Reinforcement Learning for Control Systems

Professional Development Training Course Outline

Category	Engineering	**Duration**	10 Days
Base Fee	$3,000 USD	**Delivery Mode**	Online / On-site Hybrid

Course Introduction

Introduction

This advanced training course provides a comprehensive deep dive into Reinforcement Learning (RL) for Control Systems, equipping participants with the cutting-edge skills to design, implement, and optimize intelligent, autonomous control solutions. The curriculum focuses on bridging the gap between classical control theory and the powerful paradigm of learning from interaction, enabling systems to discover optimal control policies through trial and error. Attendees will gain hands-on expertise with various RL algorithms, from Q-learning and SARSA to policy gradient methods and Deep Reinforcement Learning (DRL) architectures like DQN and PPO. Training Course on Reinforcement Learning for Control Systems is essential for control engineers, robotics specialists, and automation professionals seeking to develop adaptive, robust, and self-optimizing controllers for complex, dynamic, and often unknown environments in applications such as robotics, aerospace, industrial automation, and smart grids.

The program emphasizes practical implementation and real-world problem-solving, exploring trending topics such as model-free vs. model-based RL, multi-agent reinforcement learning (MARL), transfer learning in RL, safety-aware reinforcement learning, and the deployment of RL agents on embedded platforms for real-time control. Participants will delve into the intricacies of reward function design, exploration-exploitation trade-offs, and addressing sample efficiency challenges inherent in RL. By the end of this course, attendees will possess the expertise to architect, train, and deploy sophisticated RL-based control systems, driving unprecedented levels of autonomy, adaptability, and performance across diverse applications, from intelligent manufacturing to autonomous vehicles. This training is indispensable for professionals aiming to be at the forefront of the next generation of intelligent and autonomous control systems.

Programme Curriculum

Training Course on Reinforcement Learning for Control Systems

Introduction

This advanced training course provides a comprehensive deep dive into Reinforcement Learning (RL) for Control Systems, equipping participants with the cutting-edge skills to design, implement, and optimize intelligent, autonomous control solutions. The curriculum focuses on bridging the gap between classical control theory and the powerful paradigm of learning from interaction, enabling systems to discover optimal control policies through trial and error. Attendees will gain hands-on expertise with various RL algorithms, from Q-learning and SARSA to policy gradient methods and Deep Reinforcement Learning (DRL) architectures like DQN and PPO. Training Course on Reinforcement Learning for Control Systems is essential for control engineers, robotics specialists, and automation professionals seeking to develop adaptive, robust, and self-optimizing controllers for complex, dynamic, and often unknown environments in applications such as robotics, aerospace, industrial automation, and smart grids.

The program emphasizes practical implementation and real-world problem-solving, exploring trending topics such as model-free vs. model-based RL, multi-agent reinforcement learning (MARL), transfer learning in RL, safety-aware reinforcement learning, and the deployment of RL agents on embedded platforms for real-time control. Participants will delve into the intricacies of reward function design, exploration-exploitation trade-offs, and addressing sample efficiency challenges inherent in RL. By the end of this course, attendees will possess the expertise to architect, train, and deploy sophisticated RL-based control systems, driving unprecedented levels of autonomy, adaptability, and performance across diverse applications, from intelligent manufacturing to autonomous vehicles. This training is indispensable for professionals aiming to be at the forefront of the next generation of intelligent and autonomous control systems.

Course duration

10 Days

Course Objectives

1. Understand the fundamental concepts of Reinforcement Learning (RL), including agents, environments, states, actions, and rewards.
2. Formulate control problems as Markov Decision Processes (MDPs).
3. Implement value-based RL algorithms (Q-Learning, SARSA) for discrete control problems.
4. Apply policy-based RL methods (Policy Gradients, Actor-Critic) for continuous control tasks.
5. Develop Deep Reinforcement Learning (DRL) agents using architectures like DQN, DDPG, and PPO.
6. Design effective reward functions for various control system objectives.
7. Manage the exploration-exploitation trade-off in RL agent training.
8. Implement model-based RL techniques for improved sample efficiency.
9. Apply RL for system identification and adaptive control.
10. Explore Multi-Agent Reinforcement Learning (MARL) for cooperative and competitive control.
11. Understand safety and stability considerations when deploying RL in critical control systems.
12. Optimize and deploy RL agents on real-time embedded platforms.
13. Integrate transfer learning and curriculum learning to accelerate RL training in control.

Organizational Benefits

1. Development of highly autonomous and self-optimizing control systems, reducing manual tuning.
2. Improved adaptability and robustness of systems to unknown or changing environments.
3. Faster prototyping and deployment of advanced, data-driven control strategies.
4. Reduced energy consumption through optimally learned control policies.
5. Enhanced performance and efficiency in complex, non-linear control tasks.
6. Greater innovation capacity in robotics, automation, and intelligent systems.
7. Ability to tackle problems intractable by classical control methods.
8. Competitive advantage by leveraging cutting-edge AI for intelligent control.
9. Reduced operational costs through automated decision-making and optimization.
10. Development of resilient systems capable of learning and adapting to faults or disturbances.

Target Participants

- Control Engineers
- Robotics Engineers
- Automation Engineers
- AI/ML Engineers with a focus on control applications
- System Architects
- Researchers in Control Systems and AI
- Electrical and Mechanical Engineers involved in dynamic system design

Course Outline

Module 1: Introduction to Control and Reinforcement Learning

- **Classical Control Review:** PID, State-Space, Linear vs. Non-linear Systems.
- **Challenges of Complex Control:** Model uncertainty, dynamic environments.
- **What is Reinforcement Learning?** Agent-environment interaction, reward hypothesis.
- **Key RL Components:** States, Actions, Rewards, Policy, Value Function.
- **Case Study:** Formulating a simple inverted pendulum balancing problem as an RL task.

Module 2: Markov Decision Processes (MDPs)

- **Defining MDPs:** States, actions, transition probabilities, rewards.
- **Bellman Equations:** Optimality equations for value functions.
- **Solving MDPs:** Policy Iteration and Value Iteration algorithms.
- **Dynamic Programming:** Tabular methods for small MDPs.
- **Case Study:** Solving a simple grid-world navigation problem using Value Iteration.

Module 3: Model-Free Value-Based RL

- **Monte Carlo Methods:** Learning from episodes, first-visit, every-visit MC.
- **Temporal Difference (TD) Learning:** TD(0), TD(?) for bootstrapping.
- **SARSA Algorithm:** On-policy TD control.
- **Q-Learning Algorithm:** Off-policy TD control.
- **Case Study:** Training an agent to solve the classic Cliff Walking problem using Q-Learning.

Module 4: Policy Gradient Methods for Continuous Control

- **Policy-Based vs. Value-Based RL:** Direct policy optimization.
- **Policy Gradient Theorem:** Derivation and understanding.
- **REINFORCE Algorithm:** Monte Carlo policy gradient.
- **Actor-Critic Methods:** Combining value and policy learning.
- **Case Study:** Implementing REINFORCE to control a continuous action space (e.g., a simple robotic joint).

Module 5: Deep Reinforcement Learning (DRL) Fundamentals

- **Function Approximation:** Using neural networks for value/policy functions.
- **Experience Replay:** Breaking correlations in training data.
- **Target Networks:** Stabilizing Q-learning with deep neural networks.
- **Deep Q-Networks (DQN):** The first breakthrough in DRL.
- **Case Study:** Training a DQN agent to play a simple Atari game (e.g., Pong or Breakout).

Module 6: Advanced Deep Reinforcement Learning Architectures

- **Continuous Control with DRL:** DDPG (Deep Deterministic Policy Gradient).
- **Proximal Policy Optimization (PPO):** State-of-the-art policy gradient method.
- **Asynchronous Advantage Actor-Critic (A2C/A3C):** Parallel training.
- **Evolution Strategies (ES):** Alternative to gradient-based methods.
- **Case Study:** Applying PPO to a continuous robotic arm control problem in a simulation environment.

Module 7: Reward Function Design and Shaping

- **Importance of Reward Design:** Guiding agent behavior.
- **Sparse vs. Dense Rewards:** Challenges and strategies.
- **Reward Shaping:** Modifying rewards to accelerate learning.
- **Curiosity and Intrinsic Motivation:** Exploration strategies.
- **Case Study:** Designing an effective reward function for an autonomous drone navigation task.

Module 8: Exploration vs. Exploitation Strategies

- **Greedy vs. Epsilon-Greedy:** Simple exploration.
- **Upper Confidence Bound (UCB) and Thompson Sampling:** More sophisticated strategies.
- **Noisy Networks and Parameter Space Noise:** Exploring with neural network weights.
- **Intrinsic Motivation:** Pseudo-counts, prediction errors for novel states.
- **Case Study:** Comparing the impact of different exploration strategies on the learning speed of an RL agent.

Module 9: Model-Based Reinforcement Learning

- **Learning a World Model:** Predicting next state and reward.
- **Planning with a Learned Model:** Monte Carlo Tree Search (MCTS), Value Iteration.
- **Dyna-Q Architecture:** Integrating planning and learning.
- **Benefits:** Sample efficiency, handling complex environments.
- **Case Study:** Building a simple model of a dynamic system and using it to improve RL agent performance.

Module 10: Reinforcement Learning for Adaptive Control

- **Adaptive Control Concepts:** Handling uncertainties and varying dynamics.

- **RL for Parameter Adaptation:** Tuning controller gains online.
- **Learning Controllers for Unknown Systems:** Direct and indirect adaptive RL.
- **Robustness in Adaptive RL:** Handling noise and disturbances.
- **Case Study:** Developing an RL-based adaptive controller for a robotic manipulator with unknown payload changes.

Module 11: Multi-Agent Reinforcement Learning (MARL)

- **Cooperative vs. Competitive MARL:** Different interaction types.
- **Centralized vs. Decentralized MARL:** Architectures.
- **Challenges in MARL:** Non-stationarity, credit assignment problem.
- **MARL Algorithms:** MADDPG, QMIX.
- **Case Study:** Training multiple RL agents to cooperate in a simulated traffic light control scenario.

Module 12: Safety and Stability in RL Control Systems

- **Safety Constraints:** Incorporating constraints into RL (e.g., control barrier functions).
- **Risk-Aware RL:** Optimizing for risk-averse policies.
- **Formal Verification for RL:** Emerging techniques.
- **Human-in-the-Loop RL:** Operator intervention and supervision.
- **Case Study:** Discussing methods to ensure an RL-controlled robotic system avoids collisions.

Module 13: Transfer Learning and Curriculum Learning in RL

- **Transfer Learning in RL:** Leveraging pre-trained policies or models.
- **Domain Randomization:** Training in simulation for real-world transfer.
- **Curriculum Learning:** Gradually increasing task difficulty.
- **Meta-RL:** Learning to learn across different tasks.
- **Case Study:** Training an RL agent in a simplified simulation and transferring its knowledge to a more complex environment.

Module 14: Real-Time Deployment of RL Agents

- **Hardware for RL Inference:** GPUs, NPUs, embedded processors.
- **Model Optimization for Deployment:** Quantization, pruning.
- **Real-Time Operating Systems (RTOS) and RL:** Latency considerations.
- **Interfacing RL Agents with Physical Systems:** Sensors, actuators.
- **Case Study:** Discussing the challenges and solutions for deploying an RL agent on an actual robotic platform for real-time control.

Module 15: Future

Upcoming Scheduled Sessions

Start Date	End Date	Location	Price
23 Nov 2026	04 Dec 2026	Online	$2,000

Start Date	End Date	Location	Price
23 Nov 2026	04 Dec 2026	Physical	$3,000
23 Nov 2026	04 Dec 2026	Physical	$6,000
23 Nov 2026	04 Dec 2026	Physical	$6,000
23 Nov 2026	04 Dec 2026	Physical	$12,000
23 Nov 2026	04 Dec 2026	Physical	$9,000
23 Nov 2026	04 Dec 2026	Physical	$12,000
23 Nov 2026	04 Dec 2026	Physical	$12,000

Ready to enroll or request a customized program? Book online or contact us:

Register Online Now

Email: info@fineskilltrainingcenter.com | Website: www.fineskilltrainingcenter.com