

Training Course on Reinforcement Learning for Decision Making

Professional Development Training Course Outline

Category	Data Science	Duration	10 Days
Base Fee	\$3,000 USD	Delivery Mode	Online / On-site Hybrid

Course Introduction

Training Course on Reinforcement Learning for Decision Making provides a comprehensive deep dive into **Reinforcement Learning (RL)**, a cutting-edge paradigm revolutionizing **AI-driven decision-making** in dynamic and uncertain environments. Participants will gain a robust theoretical foundation in core RL concepts, from **Markov Decision Processes (MDPs)** and **value functions** to advanced **Deep Reinforcement Learning (DRL)** algorithms. Through hands-on exercises and real-world **case studies**, attendees will master the practical application of RL for **optimizing complex systems**, enabling intelligent agents to learn optimal strategies through **trial and error** and maximize long-term rewards.

The course emphasizes the strategic importance of RL in today's data-driven landscape, covering its transformative impact across diverse sectors like **autonomous systems, finance, robotics, healthcare, and smart infrastructure**. We will explore how RL agents can tackle challenges like **resource allocation, predictive maintenance, personalized recommendations, and strategic game theory**, equipping participants with the skills to design, implement, and evaluate RL solutions for **complex decision-making problems**. This program is ideal for professionals seeking to leverage the power of **adaptive AI** to drive innovation and achieve superior outcomes in increasingly intricate operational landscapes.

Programme Curriculum

Training Course on Reinforcement Learning for Decision Making: Theory and Applications in Complex Environments

Introduction

Training Course on Reinforcement Learning for Decision Making provides a comprehensive deep dive into **Reinforcement Learning (RL)**, a cutting-edge paradigm revolutionizing **AI-driven decision-making** in dynamic and uncertain environments. Participants will gain a robust theoretical

foundation in core RL concepts, from **Markov Decision Processes (MDPs)** and **value functions** to advanced **Deep Reinforcement Learning (DRL)** algorithms. Through hands-on exercises and real-world **case studies**, attendees will master the practical application of RL for **optimizing complex systems**, enabling intelligent agents to learn optimal strategies through **trial and error** and maximize long-term rewards.

The course emphasizes the strategic importance of RL in today's data-driven landscape, covering its transformative impact across diverse sectors like **autonomous systems, finance, robotics, healthcare, and smart infrastructure**. We will explore how RL agents can tackle challenges like **resource allocation, predictive maintenance, personalized recommendations, and strategic game theory**, equipping participants with the skills to design, implement, and evaluate RL solutions for **complex decision-making problems**. This program is ideal for professionals seeking to leverage the power of **adaptive AI** to drive innovation and achieve superior outcomes in increasingly intricate operational landscapes.

Course Duration

10 days

Course Objectives

1. Comprehend the foundational principles of Reinforcement Learning, including agents, environments, states, actions, rewards, and policies.
2. Accurately model real-world sequential decision-making challenges using the Markov Decision Process (MDP) framework.
3. Apply Value Iteration and Policy Iteration for solving MDPs with known models, understanding their computational efficiency.
4. Gain proficiency in Monte Carlo (MC) and Temporal Difference (TD) learning methods (e.g., Q-learning, SARSA) for learning optimal policies without explicit environmental models.
5. Analyze and implement strategies to balance exploring new actions with exploiting known optimal actions in uncertain environments.
6. Grasp the integration of deep neural networks with RL for handling high-dimensional state and action spaces.
7. Learn about REINFORCE, Actor-Critic, and Proximal Policy Optimization (PPO) for directly optimizing policy functions.
8. Understand the complexities and approaches for decision-making in environments with multiple interacting intelligent agents.
9. Explore techniques like Experience Replay, Prioritized Experience Replay, and Dueling DQN for enhanced learning stability and efficiency.
10. Learn robust metrics and methodologies for assessing the effectiveness and generalization of trained RL agents.
11. Analyze and design RL solutions for practical scenarios in **autonomous vehicles, algorithmic trading, robotics control, and intelligent resource management**.
12. Discuss the ethical implications of autonomous decision-making and explore methods for interpreting RL agent behavior.
13. Gain insights into cutting-edge topics like **Hierarchical Reinforcement Learning (HRL), Inverse Reinforcement Learning (IRL), and Meta-Reinforcement Learning**.

Organizational Benefits

- Empower organizations to make optimal, data-driven decisions in dynamic and uncertain business environments.

- Automate complex processes, optimize resource allocation, and improve system performance across various domains.
- Foster a culture of advanced AI adoption, enabling the development of intelligent products and services.
- Optimize energy consumption, predict equipment failures, and mitigate risks through intelligent autonomous systems.
- Develop highly personalized recommendation engines and adaptive user interfaces, leading to increased customer satisfaction and engagement.
- Leverage RL's adaptive learning capabilities to solve previously intractable optimization and control problems.
- Build systems that can continuously learn and adapt to changing market conditions, unforeseen events, and evolving user behaviors.

Target Audience

1. Data Scientists & Machine Learning Engineers
2. AI Researchers & Developers.
3. Robotics Engineers.
4. Quantitative Analysts & Traders.
5. Operations Research Analysts.
6. Software Architects & System Designers.
7. Academics & Students.
8. Product Managers & Business Leaders.

Course Outline

Module 1: Introduction to Reinforcement Learning Fundamentals

- What is Reinforcement Learning? Core concepts (agent, environment, state, action, reward, policy).
- RL vs. Supervised vs. Unsupervised Learning.
- The Reinforcement Learning Problem Formulation: Goals and Challenges.
- Key elements of an RL system and the agent-environment interface.
- **Case Study:** Simple Gridworld Navigation – training an agent to find the optimal path to a goal.

Module 2: Markov Decision Processes (MDPs)

- Defining MDPs: States, Actions, Transition Probabilities, Reward Function.
- The Bellman Equation and Optimal Value Functions (V^* and Q^*).
- Discount Factor and its importance in long-term rewards.
- Solving MDPs: Prediction and Control.
- **Case Study:** Inventory Management – modeling stock levels and demand as an MDP to optimize reordering.

Module 3: Dynamic Programming for MDPs

- Policy Evaluation: Iteratively computing value functions for a given policy.
- Policy Improvement Theorem and Policy Iteration Algorithm.
- Value Iteration Algorithm: Finding the optimal policy and value function.
- Comparison of Policy Iteration and Value Iteration, convergence properties.
- **Case Study:** Automated Factory Control – optimizing production line sequencing to maximize throughput given machine states.

Module 4: Monte Carlo Methods for Prediction

- Learning from Experience: Introduction to Model-Free Reinforcement Learning.
- First-Visit and Every-Visit Monte Carlo Prediction.
- Estimating State-Value and Action-Value Functions.
- Advantages and disadvantages of Monte Carlo methods.
- **Case Study:** BlackJack Game – learning optimal strategy by playing many hands and averaging returns.

Module 5: Monte Carlo Control & Exploration

- Monte Carlo ES (Exploring Starts) for On-Policy Control.
- On-policy vs. Off-policy learning.
- ϵ -Greedy policies for balancing exploration and exploitation.
- Limitations of Monte Carlo Control in complex environments.
- **Case Study:** Robotic Arm Manipulation – learning a sequence of movements to grasp an object, with initial random movements for exploration.

Module 6: Temporal Difference (TD) Learning: TD(0)

- Introduction to Temporal Difference Learning: Learning from bootstrapped estimates.
- TD(0) for Value Prediction: Updating value estimates based on immediate next state.
- Advantages of TD learning over Monte Carlo methods (online learning, less variance).
- Relationship between TD(0) and Dynamic Programming.
- **Case Study:** Predicting Football Game Outcomes – updating win probabilities incrementally after each play.

Module 7: TD Control: SARSA & Q-Learning

- SARSA (State-Action-Reward-State-Action): On-policy TD control.
- Q-Learning: Off-policy TD control, learning the optimal action-value function.
- Comparison of SARSA and Q-Learning, and their convergence properties.
- Practical considerations for implementing TD control algorithms.
- **Case Study:** Self-Driving Car Navigation – learning optimal turning and acceleration actions in a simulated environment based on real-time feedback.

Module 8: Function Approximation in RL

- The Need for Function Approximation: Handling large or continuous state/action spaces.
- Linear Function Approximation: Feature engineering and parameter updates.
- Introduction to Neural Networks as Function Approximators.
- Challenges of combining RL with function approximation (e.g., instability).
- **Case Study:** Robotic Locomotion – approximating continuous joint angles and motor torques for smooth movement.

Module 9: Deep Reinforcement Learning (DRL) - Deep Q-Networks (DQN)

- Introduction to Deep Q-Networks (DQN): Stabilizing Q-learning with neural networks.
- Experience Replay: Storing and replaying past experiences to break correlations.
- Target Networks: Stabilizing the training process.
- Variations of DQN: Double DQN, Dueling DQN.
- **Case Study:** Atari Game Play (e.g., Breakout, Space Invaders) – learning to play games directly from pixel inputs.

Module 10: Policy Gradient Methods

- Direct Policy Optimization: Learning policies directly instead of value functions.
- REINFORCE Algorithm: Monte Carlo policy gradient.
- Policy Gradient Theorem and its importance.
- Advantages and disadvantages of policy gradient methods (e.g., continuous action spaces).
- **Case Study:** Robotic Grasping – learning a policy to directly output motor commands for grasping objects.

Module 11: Actor-Critic Methods

- Combining Value-Based and Policy-Based Approaches.
- Actor-Critic Architecture: An actor for policy and a critic for value estimation.
- Advantage Actor-Critic (A2C) and Asynchronous Advantage Actor-Critic (A3C).
- Benefits of Actor-Critic methods (lower variance, faster convergence).
- **Case Study:** Traffic Light Optimization – an agent learning to control traffic flow by adjusting light timings based on traffic density.

Module 12: Advanced DRL Techniques

- Proximal Policy Optimization (PPO): A popular and robust policy gradient algorithm.
- Trust Region Policy Optimization (TRPO): Ensuring stable policy updates.
- Exploration Strategies: Noisy Networks, Parameter Space Noise.
- Curiosity-Driven Exploration.
- **Case Study:** Financial Portfolio Management – an agent learning to allocate assets to maximize returns while managing risk in a dynamic market.

Module 13: Multi-Agent Reinforcement Learning (MARL)

- Challenges of MARL: Non-stationarity, credit assignment, partial observability.
- Cooperative vs. Competitive MARL environments.
- Centralized Training, Decentralized Execution.
- Introduction to MADDPG and other MARL algorithms.
- **Case Study:** Autonomous Drone Swarm Coordination – training multiple drones to collaboratively map an area or perform a task.

Module 14: Practical Considerations & Ethics in RL

- Reward Shaping: Designing effective reward functions.
- Hyperparameter Tuning and Debugging RL Agents.
- Simulation Environments for RL Training.
- Ethical implications of autonomous RL systems (bias, control, accountability).
- **Case Study:** Personalized Healthcare – RL for dynamic treatment recommendations, addressing data privacy and ethical considerations.

Module 15: Advanced Topics & Future Trends in RL

- Hierarchical Reinforcement Learning (HRL): Decomposing complex tasks into sub-tasks.
- Inverse Reinforcement Learning (IRL): Learning reward functions from expert demonstrations.
- Meta-Reinforcement Learning: Learning to learn across different tasks.
- Offline Reinforcement Learning: Learning from static datasets without real-time interaction.
- **Case Study:** Drug Discovery and Material Science – using RL to navigate vast chemical spaces for optimal molecular design.

Training Methodology

This course adopts a highly interactive and hands-on training methodology, blending theoretical concepts with practical application.

- **Interactive Lectures:** Engaging presentations covering core RL theories and algorithms.
- **Code-Along Sessions:** Step-by-step guided coding exercises using Python and popular RL libraries (e.g., OpenAI Gym, Stable Baselines3, Ray RLlib).
- **Practical Labs & Exercises:** Hands-on implementation of RL algorithms on various simulated environments and problem sets.
- **Case Study Analysis:** In-depth discussion and practical application of RL to real-world business and scientific challenges.
- **Group Discussions & Problem Solving:** Collaborative sessions to foster deeper understanding and critical thinking.
- **Mini-Projects:** Opportunities to apply learned concepts to solve smaller, self-contained RL problems.
- **Q&A and Expert Feedback:** Dedicated time for questions and personalized guidance from instructors.

Register as a group from 3 participants for a Discount

Send us an email: info@fineskilltrainingcenter.org or call +254769199797

Certification

Upon successful completion of this training, participants will be issued with a globally- recognized certificate.

Tailor-Made Course

We also offer tailor-made courses based on your needs.

Key Notes

- a. The participant must be conversant with English.
- b. Upon completion of training the participant will be issued with an Authorized Training Certificate
- c. Course duration is flexible and the contents can be modified to fit any number of days.
- d. The course fee includes facilitation training materials, 2 coffee breaks, buffet lunch and A Certificate upon successful completion of Training.
- e. Upcoming Scheduled Sessions

Start Date	End Date	Location	Price
21 Sep 2026	02 Oct 2026	Online	\$2,000

Start Date	End Date	Location	Price
21 Sep 2026	02 Oct 2026	Physical	\$3,000
21 Sep 2026	02 Oct 2026	Physical	\$0
21 Sep 2026	02 Oct 2026	Physical	\$0
21 Sep 2026	02 Oct 2026	Physical	\$0
21 Sep 2026	02 Oct 2026	Physical	\$0
21 Sep 2026	02 Oct 2026	Physical	\$0
21 Sep 2026	02 Oct 2026	Physical	\$0

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

Register Online Now

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