Course Title: Mathematical Foundations of Reinforcement Learning

Lector: Prof. Dr. **Denis Belomestny** (Duisburg-Essen University, HSE University)

Course Overview:

This course offers a deep dive into the mathematical principles that underlie reinforcement learning (RL), bridging the gap between theory and application. Designed for advanced undergraduates, graduate students, and professionals with a background in mathematics and machine learning, the course covers key mathematical concepts that form the foundation of modern RL algorithms.

Course Objectives:

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

1. Understand and apply the mathematical framework of Markov Decision Processes (MDPs).

2. Analyze core concepts such as value functions, policy optimization, and Bellman equations.

3. Explore dynamic programming techniques and their relationship to RL algorithms.

4. Delve into probability theory, stochastic processes, and their relevance to exploration-exploitation trade-offs.

5. Investigate modern RL methods such as Q-learning, policy gradients, and actor-critic algorithms from a theoretical standpoint.

6. Gain insights into convergence properties, sample complexity, and the use of function approximation in RL.

Key Topics:

1. Markov Decision Processes (MDPs):
   - States, actions, rewards, and transitions
   - Policy evaluation and improvement
   - Bellman optimality and Bellman backup operators

2. Value Functions and Optimality:
   - State-value and action-value functions
   - Bellman equations and dynamic programming
   - The role of discount factors in long-term decision-making

3. Exploration and Exploitation:
   - Probability theory and stochastic processes in RL

- Multi-armed bandits and regret minimization
- Epsilon-greedy, UCB, and Thompson sampling strategies

4. Dynamic Programming:
   - Value iteration and policy iteration
   - Convergence properties of DP methods
   - The role of DP in model-based and model-free RL
5. Approximation Methods:
   - Function approximation for large state and action spaces
   - Stochastic gradient descent in RL
   - Temporal Difference (TD) learning and TD(λ)
6. Advanced RL Algorithms:
   - Q-learning and SARSA
   - Policy gradient methods and REINFORCE algorithm
   - Actor-critic algorithms and the A3C model
7. Convergence and Stability:
   - Theoretical analysis of algorithmic convergence
   - Sample complexity and efficiency
   - Stability in RL systems
8. Applications and Case Studies:
   - Practical examples of RL in robotics, gaming, finance, and more
   - Case studies illustrating theoretical concepts in real-world settings

Prerequisites:

- A solid understanding of calculus, linear algebra, and probability theory
- Prior experience with machine learning concepts
- Familiarity with basic optimization techniques

Recommended Materials:

- *"Reinforcement Learning: An Introduction"* by Sutton & Barto
- *"Dynamic Programming and Optimal Control"* by Bertsekas
- Research papers and articles on recent advances in RL theory