

MA6634: REINFORCEMENT LEARNING AND ITS APPLICATIONS IN FINANCE

Effective Term

Semester A 2026/27

Part I Course Overview

Course Title

Reinforcement Learning and Its Applications in Finance

Subject Code

MA - Mathematics

Course Number

6634

Academic Unit

Mathematics (MA)

College/School

College of Science (SI)

Course Duration

One Semester

Credit Units

3

Level

P5, P6 - Postgraduate Degree

Medium of Instruction

English

Medium of Assessment

English

Prerequisites

Nil

Precursors

Nil

Equivalent Courses

Nil

Exclusive Courses

Nil

Part II Course Details

Abstract

This course introduces fundamental mathematical methods for reinforcement learning and its application in financial problems. It helps students to understand the basic concepts and programming tools in reinforcement learning models and develops students' ability to apply the technique to financial applications.

Course Intended Learning Outcomes (CILOs)

CILOs		Weighting (if app.)	DEC-A1	DEC-A2	DEC-A3
1	Explain fundamental concepts and theories of reinforcement learning models include action space, state space, reward, value function, Finite Markov decision processes, policy, and Bellman equation.	30	x	x	
2	Apply classical dynamic programming and reinforcement learning algorithms, include policy iteration, Q learning, policy gradient methods, and deep Q learning.	25	x	x	
3	Program the reinforcement learning algorithms.	20	x	x	x
4	Implement the models in financial applications	25	x	x	

A1: Attitude

Develop an attitude of discovery/innovation/creativity, as demonstrated by students possessing a strong sense of curiosity, asking questions actively, challenging assumptions or engaging in inquiry together with teachers.

A2: Ability

Develop the ability/skill needed to discover/innovate/create, as demonstrated by students possessing critical thinking skills to assess ideas, acquiring research skills, synthesizing knowledge across disciplines or applying academic knowledge to real-life problems.

A3: Accomplishments

Demonstrate accomplishment of discovery/innovation/creativity through producing /constructing creative works/new artefacts, effective solutions to real-life problems or new processes.

Learning and Teaching Activities (LTAs)

LTAs		Brief Description	CILO No.	Hours/week (if applicable)
1	Lectures	Students learn the primary knowledge through the lectures.	1, 2, 3, 4	3 hours/week
2	Assignments	Students review the basic concepts and implement advanced theory for better understanding through take-home assignments.	1, 2, 3, 4	after-class
3	Project	In the project, students apply reinforcement learning models to more complex financial problems and implement computing techniques introduced in this course.	2, 3, 4	after-class

Assessment Tasks / Activities (ATs)

	ATs	CILO No.	Weighting (%)	Remarks ("- " for nil entry)	Allow Use of GenAI?
1	Hand-in assignments	1, 2, 3, 4	20	-	Yes
2	Project	2, 3, 4	20	-	Yes

Continuous Assessment (%)

40

Examination (%)

60

Examination Duration (Hours)

2

Assessment Rubrics (AR)**Assessment Task**

1. Hand-in assignments (for students admitted before Semester A 2022/23 and in Semester A 2024/25 & thereafter)

Criterion

1.1 Comprehensive understanding of the mathematical concept in Markov decision processes and reinforcement learning.

Excellent

(A+, A, A-) Consistently demonstrates a thorough understanding of the Markov decision processes and reinforcement learning and has strong ability to solve complex problems

Good

(B+, B, B-) Adequately demonstrates an understanding of the Markov decision processes and reinforcement learning and has ability to solve complex problems

Fair

(C+, C, C-) Demonstrates some understanding of the Markov decision processes and reinforcement learning and has some ability to solve simple problems

Marginal

(D) Demonstrates limited understanding of the Markov decision processes and reinforcement learning and has limited ability to solve simple problems

Failure

(F) Demonstrates little understanding of the Markov decision processes and reinforcement learning and is unable to solve relevant problems

Assessment Task

1. Hand-in assignments (for students admitted before Semester A 2022/23 and in Semester A 2024/25 & thereafter)

Criterion

1.2 Ability to apply appropriate algorithms to solve MDP control problems.

Excellent

(A+, A, A-) Consistently demonstrates a thorough understanding of reinforcement learning and has strong ability to solve MDP control problems

Good

(B+, B, B-) Adequately demonstrates an understanding of reinforcement learning and has ability to solve MDP control problems

Fair

(C+, C, C-) Demonstrates some understanding of reinforcement learning and has some ability to solve MDP control problems

Marginal

(D) Demonstrates limited understanding of reinforcement learning and has limited ability to solve MDP control problems

Failure

(F) Demonstrates little understanding of reinforcement learning and is unable to solve MDP control problems

Assessment Task

1. Hand-in assignments (for students admitted before Semester A 2022/23 and in Semester A 2024/25 & thereafter)

Criterion

1.3 Ability to relate reinforcement learning to financial applications.

Excellent

(A+, A, A-) Consistently demonstrates a thorough understanding of reinforcement learning and has strong ability to solve financial problems

Good

(B+, B, B-) Adequately demonstrates an understanding of reinforcement learning and has ability to solve financial problems

Fair

(C+, C, C-) Demonstrates some understanding of reinforcement learning and has some ability to solve financial problems

Marginal

(D) Demonstrates limited understanding of reinforcement learning and has limited ability to solve financial problems

Failure

(F) Demonstrates little understanding of reinforcement learning and is unable to solve financial problems

Assessment Task

2. Project (for students admitted before Semester A 2022/23 and in Semester A 2024/25 & thereafter)

Criterion

2.1 Ability to apply the reinforcement learning algorithms covered in the course in financial applications.

Excellent

(A+, A, A-) Demonstrates a comprehensive understanding of reinforcement learning algorithms and strong ability in applying program language to solve financial problems

Good

(B+, B, B-) Adequately demonstrates an understanding of reinforcement learning algorithms and ability in applying program language to solve financial problems

Fair

(C+, C, C-) Demonstrates some understanding of reinforcement learning algorithms and little ability in applying program language to solve financial problems

Marginal

(D) Demonstrates some understanding of reinforcement learning algorithms but cannot apply program language to solve financial problems

Failure

(F) Inappropriately or unable to apply reinforcement learning algorithms to solve problems by programming

Assessment Task

2. Project (for students admitted before Semester A 2022/23 and in Semester A 2024/25 & thereafter)

Criterion

2.2 Ability to program and train a simple reinforcement learning model.

Excellent

(A+, A, A-) Demonstrates a comprehensive understanding of reinforcement learning algorithms and strong ability in applying program language to solve financial problems

Good

(B+, B, B-) Adequately demonstrates an understanding of reinforcement learning algorithms and ability in applying program language to solve financial problems

Fair

(C+, C, C-) Demonstrates some understanding of reinforcement learning algorithms and little ability in applying program language to solve financial problems

Marginal

(D) Demonstrates some understanding of reinforcement learning algorithms but cannot apply program language to solve financial problems

Failure

(F) Inappropriately or unable to apply reinforcement learning model to solve problems by programming

Assessment Task

2. Project (for students admitted before Semester A 2022/23 and in Semester A 2024/25 & thereafter)

Criterion

2.3 Ability to discuss and validate the results with data.

Excellent

(A+, A, A-) Demonstrates a comprehensive understanding of reinforcement learning algorithms and strong ability to evaluate and validate the results with data

Good

(B+, B, B-) Adequately demonstrates an understanding of reinforcement learning algorithms and ability to evaluate and validate the results with data

Fair

(C+, C, C-) Demonstrates some understanding of reinforcement learning algorithms and some ability to evaluate and validate the results with data

Marginal

(D) Demonstrates limited understanding of reinforcement learning algorithms and limited ability to evaluate and validate the results with data

Failure

(F) Inappropriately or unable to evaluate and validate the results with data

Assessment Task

3. Examinations (for students admitted before Semester A 2022/23 and in Semester A 2024/25 & thereafter)

Criterion

3.1 Understanding of the fundamental concepts, principles, and their application scenarios.

Excellent

(A+, A, A-) Consistently demonstrates a thorough understanding of the concepts, principles, and applications of reinforcement learning

Good

(B+, B, B-) Adequately demonstrates an understanding of the concepts, principles, and applications of reinforcement learning

Fair

(C+, C, C-) Demonstrates some understanding of the concepts, principles, and applications of reinforcement learning

Marginal

(D) Demonstrates limited understanding of the concepts, principles, and applications of reinforcement learning

Failure

(F) Demonstrates little understanding of the concepts, principles, and applications of reinforcement learning

Assessment Task

3. Examinations (for students admitted before Semester A 2022/23 and in Semester A 2024/25 & thereafter)

Criterion

3.2 Ability to implement the dynamic programming and reinforcement learning algorithms covered in the course.

Excellent

(A+, A, A-) Consistently demonstrates a thorough understanding of dynamic programming and reinforcement learning algorithms and has strong ability to solve complex problems

Good

(B+, B, B-) Adequately demonstrates an understanding of dynamic programming and reinforcement learning algorithms and has ability to solve complex problems

Fair

(C+, C, C-) Demonstrates some understanding of dynamic programming and reinforcement learning algorithms and has some ability to solve simple problems

Marginal

(D) Demonstrates limited understanding of dynamic programming and reinforcement learning algorithms and has limited ability to solve simple problems

Failure

(F) Demonstrates little understanding of dynamic programming and reinforcement learning algorithms and is unable to solve relevant problems

Assessment Task

3. Examinations (for students admitted before Semester A 2022/23 and in Semester A 2024/25 & thereafter)

Criterion

3.3 Comprehensive problem-solving skills in financial applications underlying dynamic programming and reinforcement learning model.

Excellent

(A+, A, A-) Consistently demonstrates a thorough understanding of dynamic programming and reinforcement learning model and can always apply them to solve financial problems

Good

(B+, B, B-) Adequately demonstrates an understanding of dynamic programming and reinforcement learning model and can usually apply them to solve financial problems

Fair

(C+, C, C-) Demonstrates some understanding of dynamic programming and reinforcement learning model and can sometimes apply them to solve simple financial problems

Marginal

(D) Demonstrates limited understanding of dynamic programming and reinforcement learning model and can seldom apply them to solve simple financial problems

Failure

(F) Demonstrates little understanding of dynamic programming and reinforcement learning model and can rarely or never apply them to solve relevant problems

Assessment Task

1. Hand-in assignments (for students admitted from Semester A 2022/23 to Summer Term 2024)

Criterion

1.1 Comprehensive understanding of the mathematical concept in Markov decision processes and reinforcement learning.

Excellent

(A+, A, A-) Consistently demonstrates a thorough understanding of the Markov decision processes and reinforcement learning and has strong ability to solve complex problems

Good

(B+, B) Adequately demonstrates an understanding of the Markov decision processes and reinforcement learning and has ability to solve complex problems

Marginal

(B-, C+, C) Demonstrates some understanding of the Markov decision processes and reinforcement learning and has limited ability to solve simple problems

Failure

(F) Demonstrates little understanding of the Markov decision processes and reinforcement learning and is unable to solve relevant problems

Assessment Task

1. Hand-in assignments (for students admitted from Semester A 2022/23 to Summer Term 2024)

Criterion

1.2 Ability to apply appropriate algorithms to solve MDP control problems.

Excellent

(A+, A, A-) Consistently demonstrates a thorough understanding of reinforcement learning and has strong ability to solve MDP control problems

Good

(B+, B) Adequately demonstrates an understanding of reinforcement learning and has ability to solve MDP control problems

Marginal

(B-, C+, C) Demonstrates some understanding of reinforcement learning and has limited ability to solve MDP control problems

Failure

(F) Demonstrates little understanding of reinforcement learning and is unable to solve MDP control problems

Assessment Task

1. Hand-in assignments (for students admitted from Semester A 2022/23 to Summer Term 2024)

Criterion

1.3 Ability to relate reinforcement learning to financial applications.

Excellent

(A+, A, A-) Consistently demonstrates a thorough understanding of reinforcement learning and has strong ability to solve financial problems

Good

(B+, B) Adequately demonstrates an understanding of reinforcement learning and has ability to solve financial problems

Marginal

(B-, C+, C) Demonstrates some understanding of reinforcement learning and has limited ability to solve financial problems

Failure

(F) Demonstrates little understanding of reinforcement learning and is unable to solve financial problems

Assessment Task

2. Project (for students admitted from Semester A 2022/23 to Summer Term 2024)

Criterion

2.1 Ability to apply the reinforcement learning algorithms covered in the course in financial applications.

Excellent

(A+, A, A-) Demonstrates a comprehensive understanding of reinforcement learning algorithms and strong ability in applying program language to solve financial problems

Good

(B+, B) Adequately demonstrates an understanding of reinforcement learning algorithms and ability in applying program language to solve financial problems

Marginal

(B-, C+, C) Demonstrates some understanding of reinforcement learning algorithms and little ability in applying program language to solve financial problems

Failure

(F) Inappropriately or unable to apply reinforcement learning algorithms to solve problems by programming

Assessment Task

2. Project (for students admitted from Semester A 2022/23 to Summer Term 2024)

Criterion

2.2 Ability to program and train a simple reinforcement learning model.

Excellent

(A+, A, A-) Demonstrates a comprehensive understanding of reinforcement learning algorithms and strong ability in applying program language to solve financial problems

Good

(B+, B) Adequately demonstrates an understanding of reinforcement learning algorithms and ability in applying program language to solve financial problems

Marginal

(B-, C+, C) Demonstrates some understanding of reinforcement learning algorithms and little ability in applying program language to solve financial problems

Failure

(F) Inappropriately or unable to apply reinforcement learning model to solve problems by programming

Assessment Task

2. Project (for students admitted from Semester A 2022/23 to Summer Term 2024)

Criterion

2.3 Ability to discuss and validate the results with data.

Excellent

(A+, A, A-) Demonstrates a comprehensive understanding of reinforcement learning algorithms and strong ability to evaluate and validate the results with data

Good

(B+, B) Adequately demonstrates an understanding of reinforcement learning algorithms and ability to evaluate and validate the results with data

Marginal

(B-, C+, C) Demonstrates some understanding of reinforcement learning algorithms and limited ability to evaluate and validate the results with data

Failure

(F) Inappropriately or unable to evaluate and validate the results with data

Assessment Task

3. Examinations (for students admitted from Semester A 2022/23 to Summer Term 2024)

Criterion

3.1 Understanding of the fundamental concepts, principles, and their application scenarios.

Excellent

(A+, A, A-) Consistently demonstrates a thorough understanding of the concepts, principles, and applications of reinforcement learning

Good

(B+, B) Adequately demonstrates an understanding of the concepts, principles, and applications of reinforcement learning

Marginal

(B-, C+, C) Demonstrates some understanding of the concepts, principles, and applications of reinforcement learning

Failure

(F) Demonstrates little understanding of the concepts, principles, and applications of reinforcement learning

Assessment Task

3. Examinations (for students admitted from Semester A 2022/23 to Summer Term 2024)

Criterion

3.2 Ability to implement the dynamic programming and reinforcement learning algorithms covered in the course.

Excellent

(A+, A, A-) Consistently demonstrates a thorough understanding of dynamic programming and reinforcement learning algorithms and has strong ability to solve complex problems

Good

(B+, B) Adequately demonstrates an understanding of dynamic programming and reinforcement learning algorithms and has ability to solve complex problems

Marginal

(B-, C+, C) Demonstrates some understanding of dynamic programming and reinforcement learning algorithms and has limited ability to solve simple problems

Failure

(F) Demonstrates little understanding of dynamic programming and reinforcement learning algorithms and is unable to solve relevant problems

Assessment Task

3. Examinations (for students admitted from Semester A 2022/23 to Summer Term 2024)

Criterion

3.3 Comprehensive problem-solving skills in financial applications underlying dynamic programming and reinforcement learning model.

Excellent

(A+, A, A-) Consistently demonstrates a thorough understanding of dynamic programming and reinforcement learning model and can always apply them to solve financial problems

Good

(B+, B) Adequately demonstrates an understanding of dynamic programming and reinforcement learning model and can usually apply them to solve financial problems

Marginal

(B-, C+, C) Demonstrates some understanding of dynamic programming and reinforcement learning model and can sometimes apply them to solve simple financial problems

Failure

(F) Demonstrates little understanding of dynamic programming and reinforcement learning model and can rarely or never apply them to solve relevant problems

Part III Other Information

Keyword Syllabus

Markov Decision Processes. Bellman equation. Dynamic Programming. Q learning and Deep Q learning, Policy gradient methods. Dynamic Asset-Allocation. Optimal Exercise of American Options

Reading List

Compulsory Readings

Title	
1	Course materials provided

Additional Readings

Title	
1	Reinforcement Learning : An Introduction, Sutton & Barto 1st Edition
2	Foundations of Reinforcement Learning with Applications in Finance, Ashwin Rao & Tikhon Jelvis 1st Edition