SDSC4001: FOUNDATION OF REINFORCEMENT LEARNING

Effective Term

Semester A 2023/24

Part I Course Overview

Course Title

Foundation of Reinforcement Learning

Subject Code

SDSC - School of Data Science

Course Number

4001

Academic Unit

School of Data Science (DS)

College/School

School of Data Science (DS)

Course Duration

One Semester

Credit Units

3

Level

B1, B2, B3, B4 - Bachelor's Degree

Medium of Instruction

English

Medium of Assessment

English

Prerequisites

SDSC2002 Convex Optimization or MA3515 Introduction to Optimization
AND
MA2506 Probability and Statistics or MA2510 Probability and Statistics

Precursors

Nil

Equivalent Courses

Nil

Exclusive Courses

Nil

Part II Course Details

Abstract

This advanced elective course introduces the essential elements and mathematical foundations of the modern reinforcement learning: the optimal control theory, including dynamic programming and numerical techniques. It emphasizes both the fundamental theories in control theory and the numerical methods in context of reinforcement learning algorithms. It also equips students with computing algorithms and techniques for applications to some practical problems.

Course Intended Learning Outcomes (CILOs)

	CILOs	Weighting (if app.)	DEC-A1	DEC-A2	DEC-A3
1	Explain clearly basic concepts in reinforcement learning, optimal control and dynamic programming.	10	x		
2	Understand the concept, theory and properties of Markov Decision Process and the fundamentals of optimal control and dynamic programming	25	x	x	
3	Explain and apply the methods and theories of Markov decision process and optimal control and dynamic programming to the reinforcement learning context.	25	x	x	
4	Explain algorithms of reinforcement learning in the context of data science and machine learning.	20		x	x
5	Apply reinforcement learning to formulating and solving real-life problems	20		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.

Teaching and Learning Activities (TLAs)

	TLAs	Brief Description	CILO No.	Hours/week (if applicable)
1	Lecture	Learning through teaching is primarily based on lectures.	1, 2, 3, 4, 5	39 hours in total

2	Assignments	Learning through assignments helps students understand techniques of basic methods as well as their applications.	1, 2, 3, 4	after-class
3	Online tutorials	Learning through online tutorials helps students to solve a range of problems	4	after-class

Assessment Tasks / Activities (ATs)

	ATs	CILO No.	Weighting (%)	Remarks (e.g. Parameter for GenAI use)
1	Test	1, 2, 4	10	Questions are designed for the part of the course to see how well the students have learned basic concepts and their applications in solving problems.
2	Assignments	1, 2, 3, 4	10	The assignments provide students chances to demonstrate their achievements on techniques of dynamic programming and reinforcement learning learned in this course.
3	Project(s)	1, 2, 3, 4	30	The projects provide students chances to explore their interests and focus on the particular problem/ application/theory that they are interested in. Possible topics include state-of-the-art reinforcement learning algorithms and theories, as well as advanced topics in reinforcement learning that are not covered in lectures.

Continuous Assessment (%)

50

Examination (%)

50

Examination Duration (Hours)

2

Additional Information for ATs

Note: To pass the course, apart from obtaining a minimum of 40% in the overall mark, a student must also obtain a minimum mark of 30% in both continuous assessment and examination components.

Assessment Rubrics (AR)

Assessment Task

Test

Criterion

Ability to understand the basic concepts of methods and recognize their applications in solving application problems

Excellent (A+, A, A-)

High

Good (B+, B, B-)

Significant

Fair (C+, C, C-)

Moderate

Marginal (D)

Basic

Failure (F)

Not even reaching marginal levels

Assessment Task

Assignments

Criterion

Ability to apply the techniques in a diversity of problems

Excellent (A+, A, A-)

High

Good (B+, B, B-)

Significant

Fair (C+, C, C-)

Moderate

Marginal (D)

Basic

Failure (F)

Not even reaching marginal levels

Assessment Task

Examination

Criterion
Ability to solve problems of reinforcement learning and Markov decision process with fundamental methods.

Excellent (A+, A, A-)
High

Good (B+, B, B-)
Significant

Fair (C+, C, C-)
Moderate

Marginal (D)
Basic

Failure (F)
Not even reaching marginal levels

Assessment Task
Project(s)

Criterion
Ability to demonstrate students' achievements on techniques learned in this course

Excellent (A+, A, A-)
High

Good (B+, B, B-)
Significant

Fair (C+, C, C-)
Moderate

Marginal (D)
Basic

Failure (F)
Not even reaching marginal levels

Part III Other Information

Keyword Syllabus
Basics of dynamic programming, the shortest path problem, Markov decision processes, value iteration, policy iteration, linear programming, temporal-difference learning, Monte Carlo method, Q-learning, policy gradient, function approximation, bandit problem

Reading List

Compulsory Readings

	Title
1	Lecture note
2	Reinforcement Learning: An Introduction by Richard S. Sutton and Andrew G. Barto, The MIT Press 2017.

Additional Readings

	Title
1	Introduction to Stochastic dynamic programming By Sheldon Ross, 1983.
2	“Optimal Control Theory: An Introduction” (Dover Books on Electrical Engineering), by Donald E. Kirk. 2004.
3	Deterministic and Stochastic Optimal Control by W. Fleming and R. Rishel. Springer. 1975.