City University of Hong Kong
Course Syllabus

offered by School of Data Science
with effect from Semester A 2024/25

Part I Course Overview

Course Title:	Dynamic Programming and Reinforcement Learning
Course Code:	SDSC6007
Course Duration:	One Semester
Credit Units:	3
Level:	P6
Medium of Instruction:	English
Medium of Assessment:	English
Prerequisites: *(Course Code and Title)*	Nil
Precursors: *(Course Code and Title)*	Nil
Equivalent Courses: *(Course Code and Title)*	Nil
Exclusive Courses: *(Course Code and Title)*	Nil

Part II Course Details

1. Abstract

The course introduces Dynamic Programming - the basic models and solution techniques for problems of sequential decision making under uncertainty, and Reinforcement Learning - a framework for learning through an autonomous agent's trial and error interaction with the world to make near optimal decisions.

The course will cover the following foundational materials related to dynamic programming and reinforcement learning, including Markov decision processes, value functions, Monte Carlo estimation, dynamic programming, temporal difference learning, and function approximation. The objective of this course is to help students develop intuitive understandings of these advanced optimization and learning methods and algorithms, familiarize with the mathematical theories of these methods and algorithms, and be able to apply Dynamic Programming and Reinforcement Learning techniques to solve real-world problems.

2. Course Intended Learning Outcomes (CILOs)

(CILOs state what the student is expected to be able to do at the end of the course according to a given standard of performance.)

No.	CILOs	Weighting (if applicable)	Discovery-enriched curriculum related learning outcomes (please tick where appropriate)		
			A1	*A2*	*A3*
1.	Explain the basic ideas, concepts and principles of dynamic programming and reinforcement learning	10%	✓	✓	
2.	Explain the theories and solution methodologies for optimal control, dynamic programming and reinforcement learning	60%	✓	✓	✓
3.	Formulate various problems to be dynamic programming/reinforcement learning problems and implement the correct solvers to solve them	15%		✓	
4.	Model the applications of dynamic programming and reinforcement learning in real world	15%	✓	✓	✓
		100%			

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.

3. Learning and Teaching Activities (LTAs)

(LTAs designed to facilitate students' achievement of the CILOs.)

LTA	Brief Description	CILO No.				Hours/week (if applicable)
		1	2	3	4	
Lectures	Students will engage in lectures that introduce key knowledge points of dynamic programming and reinforcement learning methods covered in this course	✓	✓	✓	✓	29 hours/sem
Tutorial Sessions	Students will participate in tutorial sessions that have more explanation and exercises to familiarize themselves with the methods learnt during the lectures. Students will develop the ability of implementing dynamic programming and reinforcement learning algorithms.	✓	✓	✓	✓	10 hours/sem

4. Assessment Tasks/Activities (ATs)

(ATs are designed to assess how well the students achieve the CILOs.)

Assessment Tasks/Activities	CILO No.				Weighting	Remarks
	1	2	3	4		
Continuous Assessment: 70 %						
Course Project(s) Students will be assessed via course projects in demonstrating their ability in solving dynamic programming and reinforcement learning problems. Students will present their projects to their peers, and they will write feedbacks for others' presentations. Students will also write a report on their projects. 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.	✓	✓	✓	✓	30%	
Assignments Students are required to explain knowledge and apply methodologies learned from the course in solving some problems.	✓	✓	✓	✓	20%	
Midterm Students will be assessed via the examination in their ability to explain concepts and apply methodologies and technologies in solving dynamic programming problems and reinforcement learning problems learned in class.	✓	✓	✓	✓	20%	
Examination: 30% (duration: 2 hours)					30%	
					100%	

5. Assessment Rubrics

(Grading of student achievements is based on student performance in assessment tasks/activities with the following rubrics.)

Applicable to students admitted before Semester A 2022/23 and in Semester A 2024/25 & thereafter

Assessment Task	Criterion	Excellent (A+, A, A-)	Good (B+, B, B-)	Fair (C+, C, C-)	Marginal (D)	Failure (F)
1. Assignments	Based on submitted written work to evaluate the ability of explaining the knowledge of dynamic programming and reinforcement learning learned in class.	High	Significant	Moderate	Basic	Not even reaching marginal levels
2. Midterm/in-class exam(s)	Based on submitted written work and programming code to evaluate the ability of explaining the core concepts of reinforcement learning and dynamic programming, as well as the usage of different dynamic programming and reinforcement learning algorithms.	High	Significant	Moderate	Basic	Not even reaching marginal levels
3. Course project(s)	Based on submitted written work and oral presentation to evaluate students' synthesis ability and clarity, their knowledge of dynamic programming and reinforcement learning, and their ability to extend their knowledge in dynamic programming and reinforcement learning.	High	Significant	Moderate	Basic	Not even reaching marginal levels
4. Examination	Ability to solve problems of reinforcement learning and Markov decision process with fundamental methods.	High	Significant	Moderate	Basic	Not even reaching marginal levels

Applicable to students admitted from Semester A 2022/23 to Summer Term 2024

Assessment Task	Criterion	Excellent (A+, A, A-)	Good (B+, B)	Marginal (B-, C+, C)	Failure (F)
1. Assignments	Based on submitted written work to evaluate the ability of explaining the knowledge of dynamic programming and reinforcement learning learned in class.	High	Moderate	Basic	Not even reaching marginal levels
2. Midterm /in-class exam(s)	Based on submitted written work and programming code to evaluate the ability of explaining the core concepts of reinforcement learning and dynamic programming, as well as the usage of different dynamic programming and reinforcement learning algorithms.	High	Moderate	Basic	Not even reaching marginal levels
3. Course project(s)	Based on submitted written work and oral presentation to evaluate students' synthesis ability and clarity, their knowledge of dynamic programming and reinforcement learning, and their ability to extend their knowledge in dynamic programming and reinforcement learning.	High	Moderate	Basic	Not even reaching marginal levels
4. Examination	Ability to solve problems of reinforcement learning and Markov decision process with fundamental methods.	High	Moderate	Basic	Not even reaching marginal levels

Part III Other Information (more details can be provided separately in the teaching plan)

1. Keyword Syllabus

(An indication of the key topics of the course.)

Dynamic programming, Markov decision processes, temporal-difference learning, value approximation, policy approximation, Q learning, bandit problems

2. Reading List

2.1 Compulsory Readings

(Compulsory readings can include books, book chapters, or journal/magazine articles. There are also collections of e-books, e-journals available from the CityU Library.)

1.	Richard S. Sutton and Andrew G. Barto, Reinforcement Learning: An Introduction
2.	Lecture Notes and Slides

2.2 Additional Readings

(Additional references for students to learn to expand their knowledge about the subject.)

NIL

1.	Dimitris Bertsimas. Dynamic Programming and Optimal Control
2.	Csaba Czepesvári, Algorithms for Reinforcement Learning
3.	Dimitris Bertsimas. Reinforcement Learning and Optimal Control