SDSC6007: DYNAMIC PROGRAMMING AND REINFORCEMENT LEARNING

Effective Term
Semester A 2025/26

Part I Course Overview

Course Title
Dynamic Programming and Reinforcement Learning

Subject Code
SDSC - Data Science
Course Number
6007

Academic Unit
Data Science (DS)

College/School
College of Computing (CC)

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

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.

Course Intended Learning Outcomes (CILOs)

	CILOs	Weighting (if app.)	DEC-A1	DEC-A2	DEC-A3
1	Explain the basic ideas, concepts and principles of dynamic programming and reinforcement learning	10	x	x	
2	Explain the theories and solution methodologies for optimal control, dynamic programming and reinforcement learning	60	x	x	x
3	Formulate various problems to be dynamic programming/reinforcement learning problems and implement the correct solvers to solve them	15		x	
4	Model the applications of dynamic programming and reinforcement learning in real world	15	x	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 will engage in lectures that introduce key knowledge points of dynamic programming and reinforcement learning methods covered in this course	1, 2, 3, 4	29 hours/sem

2	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.	1, 2, 3, 4	10 hours/sem

Assessment Tasks / Activities (ATs)

	ATs	CILO No.	Weighting (%)	Remarks ("-" for nil entry)	Allow Use of GenAI?
1	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.	1, 2, 3, 4	30	-	Yes

2	Assignments: Students are required to explain knowledge and apply methodologies learned from the course in solving some problems.	1, 2, 3, 4	20	-	Yes
3	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.	1, 2, 3, 4	20	-	No

Continuous Assessment (%)

70

Examination (%)

30

Examination Duration (Hours)

2

Minimum Examination Passing Requirement (%)

30

Assessment Rubrics (AR)

Assessment Task

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

Criterion

Based on submitted written work to evaluate the ability of explaining the knowledge of dynamic programming and reinforcement learning learned in class.

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

Midterm/in-class exam(s) (for students admitted before Semester A 2022/23 and in Semester A 2024/25 & thereafter)

Criterion

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.

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

Course project(s) (for students admitted before Semester A 2022/23 and in Semester A 2024/25 & thereafter)

Criterion

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.

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

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
Assignments (for students admitted from Semester A 2022/23 to Summer Term 2024)

Criterion
Based on submitted written work to evaluate the ability of explaining the knowledge of dynamic programming and reinforcement learning learned in class.

Excellent
(A+, A, A-) High

Good
(B+, B) Moderate

Marginal
(B-, C+, C) Basic

Failure
(F) Not even reaching marginal levels

Assessment Task
Midterm/in-class exam(s) (for students admitted from Semester A 2022/23 to Summer Term 2024)

Criterion
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.

Excellent
(A+, A, A-) High

Good
(B+, B) Moderate

Marginal
(B-, C+, C) Basic

Failure
(F) Not even reaching marginal levels

Assessment Task
Course project(s) (for students admitted from Semester A 2022/23 to Summer Term 2024)

Criterion
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.

Excellent
(A+, A, A-) High

Good
(B+, B) Moderate

Marginal
(B-, C+, C) Basic

Failure
(F) Not even reaching marginal levels

Assessment Task
Examination (for students admitted from Semester A 2022/23 to Summer Term 2024)

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

Excellent
(A+, A, A-) High

Good
(B+, B) Moderate

Marginal
(B-, C+, C) Basic

Failure
(F) Not even reaching marginal levels

Part III Other Information

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

Reading List

Compulsory Readings

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

Additional Readings

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