The Hong Kong University of Science and Technology (Guangzhou)

Course Syllabus

[Course Title]: Introduction to Reinforcement Learning

[Course Code]: DSAA 3053

[No. of Credits]: 3

[Prerequisites]: DSAA 2012 Deep Learning

Name: Zixin ZHONG

Email: zixinzhong@hkust-gz.edu.cn

Course Description

Topics include history and motivations of reinforcement learning, tabular solution methods such as multi-arm bandits, finite Markov decision process, dynamic programming, Monte-Carlo methods and temporal difference learning, approximate solution such as on-policy approximation of action values, and off-policy approximation.

Basic knowledge on mathematics (e.g., basic of probability theory, linear algebra, calculus and optimization), programming (e.g., Python/C++/Matlab) and data science are essential and will benefit the study of this course.

Intended Learning Outcomes (ILOs)

1. Exemplify a wide variety of reinforcement learning concepts and techniques.

2. Summarize analytical methods for RL.

3. Apply the appropriate reinforcement learning tools to solve computational problems effectively.

4. Apply algorithmic design techniques to solve discrete computational problems effectively.

5. Implement and evaluate the performance of different reinforcement learning algorithms.

Assessment and Grading

This course will be assessed using criterion-referencing and grades will not be assigned using a curve. Detailed rubrics for each assignment are provided below, outlining the criteria used for evaluation.

Assessments:

Assessment Task	Contribution to Overall Course grade (%)	Due date
Mid-term test (in-class)	20%	To be determined
Written assignment	50%	To be determined
Final examination	30%	To be determined

* Assessment marks for individual assessed tasks will be released in about two weeks of the due date.

Mapping of Course ILOs to Assessment Tasks

[add to/delete table as appropriate]

Assessed Task	Mapped ILOs	Explanation
Mid-term test Final examination	[ILO1, ILO2, ILO3, ILO4, ILO5]	These tasks test students' understanding of the taxonomy of reinforcement learning algorithms (ILO 1) as well as the design of basic reinforcement learning algorithms (ILO 3), and assess students' ability to choose the proper algorithms, apply them to tack real-life problems (ILO 4), and evaluate their performance with proper metrics and by comparison with other algorithms (ILO 2, ILO5).
Written assignment	[ILO1, ILO2, ILO3, ILO4, ILO5]	These tasks assess students' ability to choose the proper algorithms, apply them to tack real-life problems (ILO1, ILO 4), and evaluate their performance with proper metrics and by comparison with other algorithms (ILO 2, ILO3, ILO5). They can also improve their collaboration skills.

Grading Rubrics

Assessed Task	Rubrics
Mid-term test Final examination	Students need to show their understanding of the taxonomy of reinforcement learning algorithms, the design of basic reinforcement learning algorithms and evaluation of algorithms.
Written assignment	Students need to (i) finish individual assignments by themselves; (ii) collaborate to decompose the project into multiple sub-projects, formulate each sub-project as a proper reinforcement learning problem, tackle the problems and evaluate the applied algorithms.

Final Grade Descriptors:

Grades	Short Description	Elaboration on subject grading description
A	Excellent Performance	Demonstrates a comprehensive grasp of reinforcement learning, expertise in problem-solving, and significant creativity in thinking. Exhibits a high capacity for scholarship and collaboration, going beyond core requirements to achieve learning goals.
B	Good Performance	Shows good knowledge and understanding of reinforcement learning, competence in problem-solving, and the ability to apply and reinforcement learning algorithms. Displays high motivation

		to learn and the ability to work effectively with others.
C	Satisfactory Performance	Possesses adequate knowledge of core reinforcement learning algorithms, competence in dealing with familiar problems, and some capacity for analysis and critical thinking. Shows persistence and effort to achieve broadly defined learning goals.
D	Marginal Pass	Has threshold knowledge of reinforcement learning, potential to achieve key professional skills, and the ability to make basic evaluation of algorithms. Benefits from the course and has the potential to develop in the discipline.
F	Fail	Demonstrates insufficient understanding of reinforcement learning and lacks the necessary problem-solving skills. Shows limited ability to think critically or analytically and exhibits minimal effort towards achieving learning goals. Does not meet the threshold requirements for professional practice or development in the discipline.

Course AI Policy

Students are allowed to use AI to help learn the knowledge but need to complete assessment tasks themselves.

Communication and Feedback

Assessment marks for individual assessed tasks will be communicated via Canvas within two weeks of submission. Feedback on assignments will include [specific details, e.g., strengths, areas for improvement]. Students who have further questions about the feedback including marks should consult the instructor/TA within five working days after the feedback is received.

Resubmission Policy

Assessment tasks are not allowed for resubmission after deadline for the fairness of all students.

Required Texts and Materials

There are no required materials. Students are encouraged to attend lectures and read some recommended materials.

Academic Integrity

Students are expected to adhere to the university's academic integrity policy. Students are expected to uphold HKUST(GZ)'s Academic Honor Code and to maintain the highest standards of academic integrity. The University has zero tolerance of academic misconduct. Please refer to Regulations for Academic Integrity and Student Conduct for the University's definition of plagiarism and ways to avoid cheating and plagiarism.

Additional Resources

Students are encouraged to familiarize themselves with following materials to obtain a good understanding of machine learning:

1. Book
    a) "Reinforcement Learning: An Introduction" by Richard S. Sutton and Andrew G. Barto
    b) "Algorithms of Reinforcement Learning" by Csaba Szepesvari
    c) "Bandit Algorithms" by Csaba Szepesvari
    d) "Reinforcement Learning: Theory and Algorithms" by Alekh Agarwal, Nan Jiang Sham, M. Kakade, and Wen Sun
2. Lecture note
    a) CS443, CS542 from UIUC by Nan Jiang
3. Webpage: Intro to RL by OpenAI

You may need help with mathematical tools with following materials:

1. `The Matrix Cookbook' [http://matrixcookbook.com] by Kaare Brandt Petersen and Michael Syskind Pedersen
2. `A First Course in Probability' by Sheldon Ross
3. `Probability and Statistics' by Morris H. DeGroot and Mark J. Schervish
4. `Probability Theory' by Achim Klenke
5. `Introduction to Linear Algebra' by Gilbert Strang
6. `Convex Optimization' by Stephen Boyd and Lieven Vandenberghe
7. `Multivariable Calculus' by James Stewart

You may get help with some important Python packages:

1. Numpy: https://numpy.org/doc/stable/reference/
2. Pandas: https://pandas.pydata.org/docs/reference/index.html
3. Scikit: https://scikit-learn.org/stable/user_guide.html

You may need help with certain topics:

1. TD Learning algorithm: Video1, Video2