The Hong Kong University of Science and Technology (Guangzhou)

UG Course Syllabus Template

Course Title: Reinforcement Learning: Principles and Methods

Course Code: AIAA 3053

No. of Credits: 3

Any pre-/co-requisites: UFUG 2104 AND (UFUG 2601 OR UFUG 2602)

Name: APOSTOLOS I. RIKOS

Email: apostolosr@hkust-gz.edu.cn

Office Hours: W4 – 311 / Meeting is anytime after email contact

TA(s): Honghao LIU (hliuco@connect.hkust-gz.edu.cn), and Yi CUI (ycui785@connect.hkust-gz.edu.cn)

Course Description

The implementation of autonomous systems requires agents to learn how to make decisions. Reinforcement learning is a powerful paradigm for achieving such a goal, and it is relevant to an enormous range of tasks, including robotics, game playing, operations research, healthcare and more. This course will provide a solid introduction to the field of reinforcement learning. Theoretical concepts will be illustrated with real-world examples and case studies from areas such as robotics, game playing, healthcare, and operations research. Students will learn about the core challenges and approaches. Through the combination of lectures, written and coding assignments, and course projects, students will be equipped with modeling and learning algorithm techniques for sequential decision-making problems. Assignments will include the basics of reinforcement learning as well as an introduction to deep reinforcement learning (an extremely promising new area that combines deep learning advancements with reinforcement learning).

Intended Learning Outcomes (ILOs)

By the end of this course, students should be able to:

1. Demonstrate fundamental understanding of reinforcement learning, such as Markov Decision Process, multi-armed bandits, Q-Learning, policy gradients, and on-policy control.
2. Address challenges in autonomous decision making by integrating scientific knowledge, technical applications, and innovative technology.
3. Identify and describe different scientific methods to critically evaluate complex, emerging decision-making problems for single agent and multi-agent scenarios.
4. Recognize the importance of exploration and exploitation to learn to make decisions under unknown environments.
5. Communicate effectively in written format and programming language to convey scientific knowledge and the application of machine learning techniques.
6. Develop a broad interest in the environment and connect the knowledge to their major study.
7. Apply the knowledge in solving engineering and technological challenges.

Course Schedule

- Week 1: Introduction to Reinforcement Learning; Multi-armed Bandits
- Week 2: Markov Decision Processes (MDPs)
- Week 3: Dynamic Programming
- Week 4: Monte Carlo Methods
- Week 5: Temporal Difference Learning
- Week 6, 7: n-Step Bootstrapping (Part I & II)
- Week 8: Planning and Learning with Tabular Methods
- Week 9: On-policy Prediction with Approximation
- Week 10: On-policy Control with Approximation
- Week 11: Policy Gradients
- Week 12: Actor–Critic Methods
- Week 13, 14: Project Presentations

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
In-class Test (closed book)	25%	During every lecture
Homework	5%	17 / 04 / 2026 *
Project Report	45%	8 / 05 / 2026 *
Project Pitch Presentation	15%	8 / 05 / 2026 *
Attendance and Participation	10%	8 / 05 / 2026 *

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

Course Project Details

- Team Size: Projects can be done individually or in teams of 2–3 students. Team projects are expected to demonstrate greater scope and depth.
- Project Scope: Students may choose to:
    - Reproduce results from a published reinforcement learning paper.
    - Extend an existing method to a new problem or dataset.
    - Propose and test a novel idea (subject to instructor approval).
- Milestones:
    - Project Proposal (Required): A 1–2 page proposal is due on 25 / 04 / 2026, outlining the problem, related work, methodology, and evaluation plan.
    - Final Report & Presentation: Due as per the assessment schedule.
    - Guidance: Students are encouraged to discuss project plans with the instructor during office hours.

Assessment(s) Grading Rubrics

In-class Test (25%) – Closed book

- Conceptual Understanding of RL principles (100%)

Homework (5%)

- Correctness (80%)
- Completeness (20%)

Project Report (45%) – Teams of 2-3 students

- Problem Definition & Literature Review (20%)
- Methodology (30%)
- Results & Analysis (30%)
- Clarity & Presentation (20%)

Project Pitch Presentation (15%)

- Clarity & Engagement (30%)
- Content & Technical Depth (40%)
- Q&A Handling (20%)

Attendance and Participation (10%)

- Regular Attendance (50%)
- Engage in Class Discussions and Contribution (50%)

Mapping of Course ILOs to Assessment Tasks

Assessed Task	**Mapped ILOs**	**Explanation**
Demonstrate thorough knowledge of the literature and a comprehensive understanding of scientific methods and techniques relevant to AI	ILO1, ILO2, ILO3, ILO5	Students conduct a structured literature review and justify method choices using evidence from prior work. They compare and synthesize findings across sources, demonstrating understanding of core scientific methods (problem formulation, experimental design, evaluation metrics, and validity considerations) and communicating the rationale clearly in writing.
Demonstrate practical skills in building AI systems	ILO2, ILO7	Students implement an AI pipeline end-to-end (data preparation, model selection, training, tuning, and evaluation) using standard tools and workflows. Performance is demonstrated via reproducible code, appropriate baselines, and documented engineering decisions (e.g., compute constraints, deployment considerations, and robustness checks).
Critically apply theories, methodologies, and knowledge to address	ILO1, ILO3, ILO4, ILO5, ILO6	Students use theoretical concepts and methodological reasoning to analyze an AI problem, select appropriate models/algorithms, and defend

fundamental questions in AI		assumptions. They evaluate alternatives, interpret results critically (including limitations and failure modes), and connect conclusions back to foundational AI questions, presenting arguments logically in written form.
Independently pursue research or innovation of significance in AI application	ILO4, ILO5, ILO6, ILO7	Students propose and execute an independent project: defining a problem, setting objectives, planning milestones, and iterating based on results. They demonstrate initiative in exploring novel ideas or non-trivial improvements (method, data, evaluation, or application), and communicate contributions, risks, and limitations with professional rigor.
Demonstrate skills in oral and written communication sufficient for a professional career	ILO5, ILO6, ILO7	Students produce professional-quality written deliverables (report/paper-style documentation) and present orally (slides/demo) to a technical audience. Assessment emphasizes clarity, structure, correct use of evidence and visuals, appropriate technical depth, and ability to respond to questions and critique constructively.

Final Grade Rubrics

A+ : [100, 96] | A : (96, 92] | A- : (92, 88] | B+ : (88, 84] | B : (84, 80] | B- : (80, 76] | C+ : (76, 72] | C : (72, 68] | C- : (68, 64] | D : (64, 60] | F : (60, 0].

Final Grade Descriptors

Grades	Short Description	Elaboration on subject grading description
A	Excellent Performance	Demonstrates outstanding mastery of AI concepts and methods, with rigorous reasoning and accurate use of literature. Produces high-quality implementations and/or analyses with strong experimental design, appropriate baselines, and well-justified conclusions. Work is original or insightful, clearly communicated, and largely free of errors.
B	Good Performance	Demonstrates solid understanding of key concepts and competent application of methods. Implementation and evaluation are generally correct with minor gaps (e.g., limited ablations, weaker justification, or small technical mistakes). Communication is clear overall, and conclusions are mostly supported by evidence.
C	Satisfactory Performance	Demonstrates adequate understanding of fundamental ideas and can complete required tasks at a basic level. Solutions are partially correct with noticeable limitations in methodology, evaluation, or interpretation. Work meets minimum requirements but lacks depth, rigor, or clarity in places.
D	Marginal Pass	Demonstrates limited understanding and inconsistent application of concepts; work is incomplete or weakly justified. Technical

		execution and evaluation contain significant issues, and conclusions may be poorly supported. Meets only the minimum threshold for passing.
F	Fail	Does not demonstrate sufficient understanding of fundamental concepts or ability to apply methods. Work is substantially incomplete, incorrect, or not submitted, with major flaws in reasoning, implementation, and/or academic integrity/requirements. Does not meet the minimum learning outcomes.

Course AI Policy

Three principles for using (generative) AI in the course:

- **Always Acknowledge AI Contributions:** When AI tools are deployed in coursework, it is essential to always cite and acknowledge their contributions Transparency in the use of AI not only upholds academic integrity but also provides clarity on the sources of information and insights.
- **Always Critically Evaluate AI-Generated Information:** AI tools can be powerful resources, but their outputs must be always assessed for accuracy and relevance It is important to verify the information generated by AI against credible sources and apply independent judgment to ensure its validity and reliability in the coursework.
- **Always Prioritize Academic Integrity:** AI should always only serve as a supplementary tool in the learning and research process, but not replacing personal efforts, understanding, and analysis in the coursework.

Communication and Feedback

- Primary Course Platform: Canvas will be used for all announcements, course materials, assignment submissions, and grade releases.
- Feedback on Assessed Work: Marks and detailed feedback for assignments will be released via Canvas within two weeks of the submission deadline. Feedback will include written comments on strengths and areas for improvement.
- Questions About Feedback: Students with questions about their marks or feedback should contact the instructor within five working days after the feedback is released.
- Contact & Consultation: For questions or to schedule a meeting, please email the instructor. Office consultations are held in W4–311 and can be arranged at a mutually convenient time after email coordination.
- Response Time: The instructor and TAs aim to respond to emails within 48 hours on weekdays.

Resubmission Policy

Late submissions will incur grade penalties as follows:

- Submitted until 1 week after deadline: 30% grade reduction.
- Submitted until 10 days after deadline: 60% grade reduction.
- Submitted until 2 week after deadline: 100% grade reduction (no credit).

Required Texts and Materials

- Richard S. Sutton, and Andrew G. Barto, “Reinforcement Learning: An Introduction,” Second Edition, MIT press, 2018.
- Dimitri Bertsekas, “Reinforcement Learning and Optimal Control,” Athena Scientific, 2019.

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

–