Xi'an Jiaotong-Liverpool University
西交利物浦大學

XJTLU Entrepreneur College (Taicang) Cover Sheet

Module code and Title	**DTS307TC Reinforcement Learning**
School Title	**School of AI and Advanced Computing**
Assignment Title	**Coursework 2**
Submission Deadline	**05/Jun/2026 23:59**
Final Word Count	
If you agree to let the university use your work anonymously for teaching and learning purposes, please type **"yes"** here.	

I certify that I have read and understood the University's Policy for dealing with Plagiarism, Collusion and the Fabrication of Data (available on Learning Mall Online). With reference to this policy I certify that:

- My work does not contain any instances of plagiarism and/or collusion.
  My work does not contain any fabricated data.

By uploading my assignment onto Learning Mall Online, I formally declare that all of the above information is true to the best of my knowledge and belief.

Scoring – For Tutor Use	
Student ID	

Stage of Marking	**Marker Code**	**Learning Outcomes Achieved (F/P/M/D) (please modify as appropriate)**			**Final Score**
		A	**B**	**C**	
1st Marker – red pen					
Moderation – green pen	**IM Initials**	The original mark has been accepted by the moderator (please circle as appropriate):			Y / N
		Data entry and score calculation have been checked by another tutor (please circle):			Y
2nd Marker if needed – green pen					

For Academic Office Use			**Possible Academic Infringement (please tick as appropriate)**	
Date Received	**Days late**	**Late Penalty**	☐ **Category A**	Total Academic Infringement Penalty (A,B, C, D, E, Please modify where necessary) _____
			☐ **Category B**	
			☐ **Category C**	
			☐ **Category D**	
			☐ **Category E**	

DTS307TC Reinforcement Learning

Coursework - Individual Project Report

Due: 05/Jun/2026 23:59
Weight: 60%
Maximum score: 60 marks

Overview

The purpose of this assignment is to gain experience in Python programming and the design of reinforcement leaning algorithms. You are expected to implement an RL algorithm that solves a specific environment and provide an explanation of the algorithm's methodology. You are expected to analyse your results, including challenges and your solutions.

Learning Outcomes Assessed

B: Critically analyse real-life problem situations and expertly map them as reinforcement learning tasks

C: Mastery of Monte Carlo Methods and Temporal Difference Learning

D: Proficiency in Deep Reinforcement Learning algorithms

E: Systematically apply RL techniques to solve practical problems in various domains, such as robotics, game playing, or autonomous systems

Late policy

5% of the total marks available for the assessment shall be deducted from the assessment mark for each working day after the submission date, up to a maximum of **five** working days.

Avoid Plagiarism

- Do ***not*** submit work from other students.
- Do ***not*** share code/work with other students
- Do ***not*** use open-source code as it is or without proper reference.

Risks

- Please read the coursework instructions and requirements carefully. Not following these instructions and requirements may result in a loss of marks.
- The assignment must be submitted via Learning Mall. Only electronic submission is accepted, and no hard copy submission.
- All students must download their file and check that it is viewable after submission. Documents may become corrupted during the uploading process (e.g. due to slow internet connections). However, students are responsible for submitting a functional and correct file for assessments.
- Academic Integrity Policy is strictly followed.

Individual Project Report (60 marks)

The goal of this project is for students to apply deep reinforcement learning (DRL) concepts and techniques to solve an Atari game. You are expected to implement a DRL algorithm of your choice, train it to achieve a competitive level of performance on the chosen game, and analyze the results.

You will select one Atari game from the OpenAI Gym Atari environment (https://ale.farama.org/environments/). Then, you will implement and train a DRL agent using an algorithm of your choice.

The project report should include a clear explanation of the game chosen, research, and summarize the current state of deep reinforcement learning, focusing on its application to Atari games. Discuss the algorithms you considered and why you chose the one you implemented. Provide a detailed explanation of the chosen algorithm and your implementation. You will also need to assess the performance of the agent, explain the benchmarks you choose, and give insights into why the chosen algorithm performs well or poorly on the selected game.

Upon completion of your implementation, you are required to submit a comprehensive technical report. The report should document your engineering decisions, the theoretical grounding of your code, and a critical analysis of the agent's performance.

> *Note: All figures and plots must be clearly labeled with axes titles and legends. Raw code snippets should be kept to a minimum in the report; focus on high-level logic and pseudo-code where necessary.*

It is crucial that your report strictly follows these guidelines:

- Do **NOT** use Stable-baselines libraries or any other reinforcement learning specific libraries in your implementation, but you are allowed to use them for benchmarking.
- Do **NOT** exceed the word count limit of **3000** words, reference and appendix excluded.
- Although you are allowed to use any generative AI tools to assist your work, please keep in mind that you should be using them **responsibly**. (Good use: Improve your report after writing it and always review its output to ensure that it is correct. Bad use: Copy-pasting an entire report from AI without any effort of your own.)

Submission Requirements

Please prepare and submit the following documents:

- A cover page featuring your student ID. This page should be the first page of your report.
- A zip file containing all the source codes and your trained model, which should be named using your full name and student ID in the following format: CW2_ID_Name.zip
- One PDF file for your report. The file should be separated from the zip file, which contains your code. The files should be named in the following format: CW2_ID_Name.pdf

Note that the quality of the code, the clarity of your writing, and the format/style of your report will be taken into consideration during the evaluation. The detailed marking scheme is outlined below.

Rubric

You can find a rubric in the next two pages. Please use them to guide your writing.

Category	Criteria	Max Marks	Marks		
			0-2	3-4	5-6
Research and Background (17 Marks)	Comprehensive summary of the current state of deep reinforcement learning.	6	Superficial coverage, lacks clear understanding or misses key areas.	Good coverage of topics with some depth and clear structure, but may lack analysis or breadth.	Comprehensive, in-depth, and well-referenced summary with analysis of current trends and advancements.
	Focused discussion on DRL applications to Atari games.	5	Minimal discussion, with vague or incomplete references to Atari games and DRL methods.	Reasonably detailed coverage of methods and performance, but may lack depth in evaluation or broader insights.	Thorough, well-structured discussion with in-depth analysis of methods, performance, and implications.
	Explanation of the chosen game and its challenges.	6	Superficial explanation of the game with little to no discussion of challenges or relevance to DRL.	Provides a solid description of the game and challenges but lacks depth in connecting challenges to DRL methods.	Comprehensive and insightful explanation, with a detailed discussion of challenges and how they are addressed by DRL techniques.
Algorithm Selection and Explanation (18 Marks)	Clear explanation of algorithms considered and rationale for selecting one.	6	Superficial discussion with minimal or unclear explanations of the algorithms and no strong rationale for selection.	Provides a good explanation of algorithms and rationale but lacks depth in comparison or trade-offs.	Comprehensive explanation of algorithms with a well-justified and insightful rationale for selecting one, supported by clear comparisons and trade-offs.
	Detailed explanation of the chosen algorithm, including key concepts and techniques.	6	Provides minimal or unclear explanation, lacking sufficient detail or focus on key concepts and techniques.	Explains key concepts and techniques well but may lack depth or fail to connect the algorithm to the problem context.	Comprehensive and well-structured explanation of the algorithm, including clear details on key concepts, techniques, and their relevance to the task.
	Insightful discussion of implementation details.	6	Limited or vague discussion of implementation details, lacking specific examples or insights.	Provides a clear description of implementation with some insights but lacks depth or detailed reflection on challenges and outcomes.	Comprehensive and insightful discussion with detailed explanations of components, challenges, and thoughtful reflections on results and improvements.

Performance Assessment (16 Marks)	Clear explanation of benchmarks chosen and their relevance.	6	Minimal or unclear explanation of benchmarks, with limited connection to the problem or relevance.	Provides a good explanation of benchmarks and relevance but lacks depth in discussing metrics or limitations.	Comprehensive explanation of benchmarks, their relevance, and the evaluation metrics, with insightful reflections on limitations and trade-offs.
	Assessment of the agent's performance using appropriate metrics.	5	Minimal or unclear discussion of metrics and their application to assessing performance.	Provides a good application of metrics and some analysis of results, but lacks depth in reflections or specific insights.	Comprehensive and insightful assessment, with detailed application of metrics, in-depth analysis of performance, and thoughtful discussion of limitations and improvements.
	Insights into why the chosen algorithm performs well or poorly on the game.	5	Minimal or vague discussion of why the algorithm performs well or poorly, with little connection to game-specific contexts.	Provides a clear discussion of strengths and weaknesses but lacks depth in evidence or recommendations for improvement.	Comprehensive and insightful analysis of the algorithm's performance, with detailed strengths, weaknesses, and well-thought-out recommendations.
Clarity and Quality of Report (9 Marks)	Clear structure, proper grammar, and professional formatting.	5	Poor structure with frequent grammar issues or inconsistent formatting	Good structure and language use but with occasional grammar or formatting inconsistencies.	Exceptionally clear structure, polished grammar, and professional formatting, creating a cohesive and visually appealing document.
			0-1	**2**	
	Effective use of visual aids (e.g., graphs, tables, or diagrams) to present data.	2	Visual aids are either irrelevant, unclear, or poorly formatted, hindering understanding.	Visual aids are highly effective, well-designed, and contribute significantly to the clarity and understanding of the data presented.	
	Proper use of citation	2	Citations are incorrect, incomplete, or inconsistent, with missing or misformatted references.	Citations are accurate, properly formatted, and consistent throughout the document, with all referenced materials properly credited.	