GHENT UNIVERSITY

Course Specifications

Valid as from the academic year 2024-2025

Reinforcement Learning (E061360)

Course size *(nominal values; actual values may depend on programme)*

Credits 6.0 **Study time 180 h**

Course offerings and teaching methods in academic year 2024-2025

A (semester 1)	English	Gent	group work	5.0h
			lecture	30.0h

Lecturers in academic year 2024-2025

Simoens, Pieter	TW05	lecturer-in-charge

Offered in the following programmes in 2024-2025	**crdts**	**offering**
Bridging Programme Master of Science in Bioinformatics(main subject Engineering)	6	A
Master of Science in Bioinformatics(main subject Engineering)	6	A
Master of Science in Computer Science Engineering	6	A
Master of Science in Computer Science Engineering	6	A

Teaching languages

English

Keywords

machine learning, reinforcement learning, deep neural networks, multi-agent systems

Position of the course

Reinforcement learning is the third major category in machine learning, alongside supervised and unsupervised learning. This technique is used for sequential decision problems, where an agent takes actions that maximize its expected cumulative reward. Combined with deep neural networks, called deep reinforcement learning, this technique is used for control problems in many domains.
The objective of this course unit is to apply both the basic principles and modern paradigms for reinforcement learning, so that the student is armed to easily understand and frame new academic literature in this domain independently. Both single-agent and multi-agent situations are discussed. In a number of assignments, students gain hands-on experience

Contents

- Modelling sequential decision making as Markov Decision Process
- Concepts, learning and planning for tabular reinforcement learning: value iteration, Q-learning, SARSA, policy improvement theorem
- Deep Value-Based RL: Deep Q Networks (DQN), Double DQN, prioritized expeirence replay
- Policy-based RL: REINFORCE, Asynchronous Advantage Actor Critic (A3C), Proximal Policy Optimization
- Model-based RL: uncertainty models, world models, model-predictive control, deep end-to-end planning and learning
- multi-agent RL: decentralized partially observable MDP, self-play, cooperative behavior

Initial competences

- Python programming skills
- Hands-on experience in machine learning, in particular deep learning

Final competences

1 Application of reinforcement learning on tractable environments
2 Acquire knowledge and insights into the state-of-the-art in the domain of reinforcement learning, and being able to situate novel techniques in the subbranches of RL
3 Being able to critically assess latest scientific advances in the domain of RL

Conditions for credit contract

Access to this course unit via a credit contract is determined after successful competences assessment

Conditions for exam contract

This course unit cannot be taken via an exam contract

Teaching methods

Group work, Lecture

Study material

Type: Handbook

Name: Book
Indicative price: Free or paid by faculty
Optional: no
Additional information: Als handboek gebruiken we A. Plaat, Deep Reinforcement Learning, Springer Nature, 2022. Dit handboek is gratis online beschikbaar. Geschatte kostprijs van de boekversie: 60 EUR.

References

[1] A. Plaat, Deep Reinforcement Learning, Springer Nature, 2022
[2] R. Sutton, A. Barto, Reinforcement Learning: An Introduction, 2nd edition, MIT Press, 2018

Course content-related study coaching

Lecturers and teaching assistants are available after each lecture, or after making an appointment

Assessment moments

end-of-term and continuous assessment

Examination methods in case of periodic assessment during the first examination period

Written assessment

Examination methods in case of periodic assessment during the second examination period

Oral assessment, Written assessment

Examination methods in case of permanent assessment

Participation, Peer and/or self assessment, Assignment

Possibilities of retake in case of permanent assessment

examination during the second examination period is possible in modified form

Extra information on the examination methods

- periodic assessment, first period: written exam with theory and application
- periodic assessment, second period: written exam with theory and application. Oral defense of revised project.
- non-periodic assesment: one-or-more implementation tasks in the domain of RL. These assesment can be group assesments or individual tasks

Calculation of the examination mark

First period:
Final score is calculated as weighted average. Non-periodic assessment is weighted by 40% in the final score, the periodic assesment by 60%. Participation to the assignments of the non-periodic assessment is a mandatory requirement for obtaining a credit. Students must obtain at least 9/20 for the non-periodic assesment and at least 9/20 for the periodic assesment. If none of the aforementioned conditions is met, students cannot get a higher grade than 8/20.

Second period
The final score is calculated using the same weights as the first period. Students who obtained 10/20 for the periodic assessment do not have to retake the periodic assessment. The score of the first period on that part is transferred.

The score for the non-periodic assessment is calculated as a correction of the score in the first period. The score for the non-periodic assessment is calculated as 0.7 x score_first_period + 0.3 x score_second_period.