



COURSE SYLLABUS

Reinforcement Learning, 7.5 credits

Förstärkningsinläring, 7,5 högskolepoäng

Course Code: TFSS25	Education Cycle: Second-cycle level
Confirmed by: Dean Aug 14, 2024	Disciplinary domain: Technology
Revised by: Director of Education Nov 28, 2024	Subject group: DT1
Valid From: Aug 1, 2025	Specialised in: A1F
Version: 2	Main field of study: Computer Science

Intended Learning Outcomes (ILO)

After a successful course, the student shall:

Knowledge and understanding

- display knowledge of central theories, algorithms, and performance criteria in Reinforcement Learning (RL),
- display an understanding of possibilities and limitations in RL,
- display knowledge of common tools and frameworks used in RL,
- display an understanding of seminal contributions in the RL research literature,

Skills and abilities

- demonstrate skills in formulating problems as a Markov Decision Process (MDP),
- demonstrate skills in choosing appropriate RL algorithms for specific MDPs,
- demonstrate skills in developing and implementing RL solutions to various problems using common tools and frameworks,

Judgement and approach

- demonstrate sound judgement in formulating problems as MDPs, selecting appropriate RL algorithms for solving specific MDPs, and evaluating solutions across multiple performance criteria.

Contents

The quest to fully realize the potential of Artificial Intelligence (AI), requires autonomous systems that can learn to make good decisions by interacting with their environment. Reinforcement learning is a paradigm that meets these requirements, and can be applied to various tasks, including game-playing, healthcare, economics, and robotics. This course gives a solid introduction to reinforcement learning with its core approaches and challenges, and is structured around several lectures, assignments, and a project.

The course includes the following elements:

- Markov Decision Processes (MDPs)

- Model-based and model-free prediction and control
- On-policy and off-policy methods
- Monte Carlo, Temporal Difference, Policy-Gradient, and Actor-Critic methods
- The exploration versus exploitation trade-off, including regret
- The bias versus variance trade-off, including stability
- Function approximation, including Deep Reinforcement Learning
- Imitation Learning and Reinforcement Learning with multi-agent interactions

Type of instruction

Lectures, exercises, and seminars.

The teaching is conducted in English.

Prerequisites

Passed courses at least 90 credits within the major subject Computer Engineering, Computer Science, Electrical Engineering (with relevant courses in Computer Engineering) or equivalent, or passed courses at least 150 credits from the Computer Science and Engineering programme, and taken courses in Artificial Intelligence, 7,5 credits, Machine Learning, 7,5 credits and Deep Learning, 7,5 credits or equivalent. Proof of English proficiency is required.

Examination and grades

The course is graded 5,4,3 or Fail.

Registration of examination:

Name of the Test	Value	Grading
Assignment ¹	5 credits	5/4/3/U
Project	2.5 credits	U/G

¹ Determines the final grade of the course, which is issued only when all course units have been passed.

Course literature

The course literature is determined 8 weeks before the start of the course.

Title: Reinforcement Learning, 2nd Edition

Author: Richard S. Sutton and Andrew G. Barto

Publisher: Bradford Books, 2018

ISBN: 9780262039246

Title: Grokking Deep Reinforcement Learning

Author: Miguel Morales

Publisher: Manning, 2020

ISBN: 9781617295454