REINFORCEMENT LEARNING IN DIGITAL FINANCE

Course syllabus

MSCA Digital Finance
Academic year: 2024-2025
Course code: MSCA_DF_24_19

Course coordinator: Wouter van Heeswijk

DIGITAL

UNIVERSITY OF TWENTE.

Funded by the European Union

Course manual Reinforcement Learning in Digital Finance

Content

1. Meet the teaching team

Wouter van Heeswijk
Assistant Professor in Operations Research & Financial Engineering
University of Twente
w.j.a.vanheeswijk@utwente.nl

Jörg Osterrieder
Associate Professor in Artificial Intelligence and Finance
University of Twente
Professor in Finance and Artificial Intelligence
Bern University of Applied Sciences
joerg.osterrieder@utwente.nl

Anne Zander
Assistant Professor in Stochastic Optimization
University of Twente
a.b.zander@utwente.nl

Martijn Mes
Full Professor in Transportation and Logistics Management
University of Twente
m.r.k.mes@utwente.nl

Branka Hadji MIsheva
Professor in Applied Data Science and Finance
Bern University of Applied Sciences
branka.hadjimisheva@bfh.ch

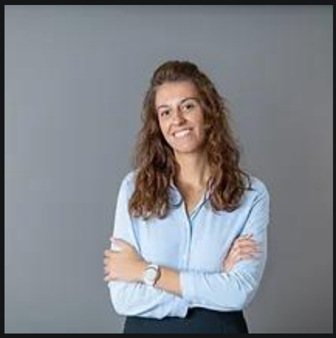

Stefano Penazzi
Senior Data Scientist
Cardo AI
stefano.penazzi@cardoai.com

Warren Powell
Professor Emeritus in Stochastic Optimization
Princeton University
Chief Analytics Officer
Optimal Dynamics
https://faculty.princeton.edu/people/warren-buckler-powell

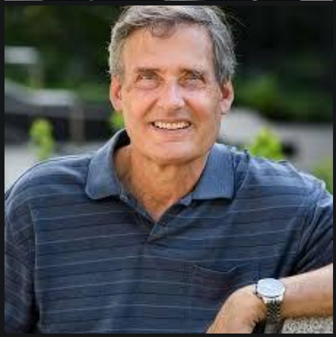

Christian Spethmann
Data Scientist
Swedbank Group
christian.spethmann@swedbank.ee

Adrian Costea
Professor in Data Mining and Econometrics
Bucharest University of Economic Studies
Regional Director
National Bank of Romania
adrian.costea@csie.ase.ro

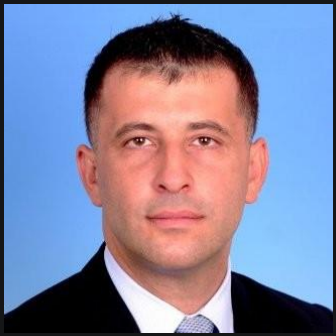

2. General course description

The course ‘Reinforcement Learning in Digital Finance’ introduces doctoral candidates to the main concepts and techniques of reinforcement learning. The focus lies on applications in digital finance, i.e., optimizing investment portfolios or managing capital in finance markets, requiring strategies that maximize returns while navigating uncertainties. Solving such complex financial problems is crucial for professionals in digital finance and financial engineering. Reinforcement learning has emerged as a powerful solution method for these dynamic resource allocation challenges, making it a key area of study for those interested in modern financial technologies.

The primary objective of this course is to introduce reinforcement learning, offering a generic introduction of core techniques with a focus on their applications in digital finance. By the end of the course, candidates will have a global understanding of reinforcement learning techniques and their application in digital finance. They will be equipped to formally describe allocation problems and select and apply appropriate reinforcement learning algorithms within this domain.

To achieve the learning objectives, the course combines lectures, a group project in which a problem environment and reinforcement learning algorithm are coded and evaluated, and a short paper describing the project. The lectures will cover the theoretical foundations of reinforcement learning, various types of algorithms, modeling techniques, and applications in finance. While the focus will be on the basics of reinforcement learning, the course will also introduce more advanced techniques, such as deep neural networks and multi-agent systems. Demonstration models and industry cases will be provided in a tutorial setting, allowing students to experiment and apply these techniques to real-world financial scenarios.

3. Prior knowledge

Doctoral candidates in this course are expected to have a basic understanding of stochastic models and (partially observable) Markov Decision Processes. Furthermore, a certain degree of familiarity with programming (the course will use Python), dynamic programming, calculus and statistics is expected. Candidates not meeting all prior knowledge requirements can still successfully complete the course but are expected to put more effort into self-study.

4. Learning objectives

This course addresses 7 learning objectives. By the end of this course, the doctoral candidate will be able to:

1. Explain the core concepts of reinforcement learning, the computational challenges it tackles and the potential applications;
2. Formulate dynamic financial resource allocation problems as (partially observable) Markov Decision Process models;
3. Model and apply basic value- and policy function algorithms;
4. Explain the relevant tunable parameters that determine learning performance, including features and learning-, exploration- and discount rates;
5. Apply neural networks in the context of (deep) reinforcement learning, as well as other contemporary developments in the domain;

6. Develop algorithmic pipelines embedding reinforcement learning techniques in the context of digital finance, starting with retrieving raw financial data, extracting appropriate features and concluding with comprehensible outputs geared towards managerial audiences;
7. Explain and present design choices and business implications to both team members and external stakeholders.

5. Study materials

Course materials will be made available via GitHub. The lecture slides form the basis of the course, the other materials serve to solidify your understanding of the topics. As a textbook, Sutton & Barto (2018, freely available online) is used.

- Lecture slides and materials provided during the lecture
- Reinforcement Learning: An Introduction (Sutton & Barto, 2018)
- Selected papers, to be provided via Github
-
- Permanent link to the Github course repository:
- https://github.com/MSCA-DN-Digital-Finance/Courses/tree/main/Cohort%201%20(2024.01-2027.12)/MSCA_DF_19%20Reinforcement%20Learning%20in%20Digital%20Finance%20%20(4EC%2C%20UTW)

6. Course setup

6.1 Activity overview

The activity overview below is indicative and may be subject to changes. Please refer to the course's GitHub page for the most up-to-date timetable.

Date	Time	Instruction mode	Topic	Teacher(s)
03-02-2025	10:00-10:15	Lecture	Welcome and introduction to MSCA Digital network	Jörg Osterrieder
	10:15-11:30	Lecture	Course introduction and general introduction to reinforcement learning	Wouter van Heeswijk, Martijn Mes
	11:45-13:00	Lecture	Markov decision processes and basics of temporal difference learning	Wouter van Heeswijk
	14:00-15:00	Lecture	Reinforcement learning in digital finance	Jörg Osterrieder
	15:15-16:00	Project	Q-learning in taxicab environment Group formation, topic selection and problem formulation	Wouter van Heeswijk, Jörg Osterrieder
	16:15-17:00	Lecture	Convergence proofs for Q-learning	Anne Zander

	17:00-18:00	Tutorial	Q-learning in taxicab environment Group formation, topic selection and problem formulation	Wouter van Heeswijk
04-02-2025	10:00-12:00	Lecture	Policy-based reinforcement learning	Wouter van Heeswijk
	13:00-15:00	Lecture	Deep reinforcement learning in finance	Jörg Osterrieder
	15:00-17:00	Project	Project: description and coding	Wouter van Heeswijk, Jörg Osterrieder
	17:00-18:00	Keynote lecture	Parameterized policies in the finance industry	Warren Powell (online)
05-02-2025	10:00-11:30	Lecture	Explainable AI in reinforcement learning	Branka Hadji Misheva, Wouter van Heeswijk
	11:45-13:00	Tutorial	Project: explainable components	Branka Hadji Misheva, Wouter van Heeswijk
	14:00-16:00	Project	Project: description and coding EB Meeting	
	16:00-17:00	Lecture	Guest lecture: Combining fuzzy clustering and artificial neural networks for financial performance predictions	Adrian Costea
	17:00-18:00	Lecture	Guest lecture: Beyond Automation: Leveraging Generative AI at Swedbank - Christian Spethmann	Christian Spethmann
06-02-2025	10:00-11:30	Lecture	Industry perspective on tabular transfer learning	Stefano Penazzi
	12:00-13:00	Presentation session	Discussion panel: prospects and barriers of Reinforcement learning in digital finance	Wouter van Heeswijk, Jörg Osterrieder, Stefano Penazzi
07-02-2025	10:00-11:30	Tutorial	Project: presentation design	Wouter van Heeswijk, Jörg Osterrieder

	11:30-13:00	Presentation session	Project progress presentation and discussions	Wouter van Heeswijk, Jörg Osterrieder
		Meeting (online)	Monthly progress meetings on project	
04-06-2025		Presentation session (online)	Final group presentations	

6.2 Assessment

Candidates are assessed based on the group project in a pass/fail setting, with the evaluation encompassing the quality of the code, the final presentation and the paper. The project can be done in groups of 2-3 people and requires substantial coding contributions from each individual member. In case of insufficient project evaluation, a single repair opportunity is provided.

6.3 Project due dates

The due dates for the project are as follows:

Due dates	Assignment	Content
03-02-2025	Project topic and group formation	Choose a project topic and form a group
30-05-2025	Codebase - Functioning RL algorithm Experimental results - Key conclusions Final presentation slides Short paper	20-minute presentation +10 minutes Q&A [hand in presentation slides] Codebase Short paper

All project files need to be handed in via GitHub before 30-05-2025.

In case project quality is insufficient upon final submission, repair opportunities will be discussed individually. High-quality projects may receive support beyond the course deadlines to transform the group work into an academic paper.

6.4 Lectures and tutorials format

The lectures and tutorials will be held in a hybrid setting; candidates are encouraged to physically attend the training week. Project progress meetings and final presentations will be held online.

1. **Lectures and tutorials**
   Lectures and tutorials will be organized in a hybrid setting during the training week. Supplementary guest lectures may be offered online.
2. **Project progress meetings**
   Project progress meetings will be held online monthly. You will meet with the teachers to present your progress and discuss open questions.
3. **Office hours**
   Office hours for organizational, project or technical questions may be scheduled on demand and will be held online.

4. **Final presentation**
   The final project presentations will be held online.

6.5 AI and group contribution rules

For all group assignments and presentations, it holds that you should hand in/present your own and original work. You must add an “author contribution & use of AI statement” to the group assignment.

The author contribution statement should include who did what (tasks) and what was the relative contribution of each group member to the overall contribution (percentage). Also, all group members should explicitly agree on the final version of the assignment.

Example author contribution statement:

** Name group member 1: Wrote the introduction of the report, produced the mathematical model of Module 1, downloaded and cleaned the data, produced output statistics and wrote answers 1.1 and 2.3. She debugged the Python code to make the mathematical program work. She read the final version of the report and made final edits. [20%]*

** Name group member 2: … [30%]*

Example of AI statement:

** We declare that no content produced by AI technology has been presented as our own work (both in reporting and coding)*

** We declare that we used ChatGPT 4.0 to improve writing at the sentence level and to better express transitions between paragraphs.*

** We declare that ChatGPT 4.0 has been used to generate initial code snippets and to generate docstrings for functions.*

Note: An extra oral assessment may be part of each assignment as a verification of the authenticity and contribution. Such an oral assessment could also be randomly assigned to a group.