

Reinforcement Learning – Opening Remarks

Christopher Mutschler



COVID-19

What we wanted to do:

- Originally, we planned to do everything in presence (or at least something hybrid)

What we are doing:

- Lecture material → online lessons (all available in the week before Q&A)
- Q&A Zoom meetings aligned with the original time plan
- Online Zoom-Exercises

Code of Conduct

For efficient zoom meetings:

- Please **mute your microphone** unless you are speaking
- **Turn on your video** to increase presence

If you see technical issues:

- Please tell me immediately in chat or by raising your hand 😊

Important:

Please do not record the meeting!


Class Logistics

- We (will start with) use studon for main communication
- If you have any questions, you can also write to




Christopher

christopher.mutschler@iis.fraunhofer.de



Lukas

lukas.schmidt@iis.fraunhofer.de




Sebastian

sebastian.rietsch@iis.fraunhofer.de


- If we will ever use a password somewhere, it will be: **FAU_RL_2021**

Syllabus


- What pre-requisites do we expect?
 - Analysis/Calculus
 - Multivariate Statistics
 - Machine Learning
 - Python
- How can you get them?
 - Attending basic lectures 😊
 - Or dive into one of them:



*Kevin Murphy:
Machine learning; a
probabilistic
perspective.*



*Christopher Bishop:
Pattern Recognition
and Machine Learning*



*Ian Goodfellow and Yoshua
Bengio and Aaron Courville:
Deep Learning*

- You will find RL literature in Lesson 1.01

Syllabus

Lecture (consultation hour: Thursday 8:30-10:00) https://fau.zoom.us/j/68176752775?pwd=UDdPjdMbXZsSjc2eFdoVzkrTEFnQT09		Exercises (Mondays: 10:15 – 11:45) https://fau.zoom.us/j/61413002667?pwd=VEE1TkdxZy9jTFRGYXZoNEh1S1pyUT09	
15.04.	Introduction to RL (Q&A: 22.04.)	19.04.	Environments & MDPs
22.04.	Dynamic Programming (Q&A: 29.04.)	26.04.	DP: Policy & Value Iteration
29.04.	Model-Free Prediction (Q&A: 06.05.)	03.05.	Intro to OpenAI Gym + TD-Learning
06.05.	Model-Free Control (Q&A: 20.05.)	10.05.	SARSA and Q-Learning in Gym
13.05.	Holiday	17.05.	Intro to PyTorch
20.05.	Value Function Approximation (Q&A: 27.05.)	24.05.	DQNs
27.05.	Policy-based RL #1 (Q&A: 10.06.)	31.05.	
03.06.	Holiday	07.06.	Policy Gradients: Vanilla Policy Gradient, Advantage Actor-Critic
10.06.	Policy-based RL #2 (Q&A: 17.06.)	14.06.	
17.06.	Model-based RL (Q&A: 24.06.)	21.06.	
24.06.	Offline RL (Q&A: 01.07.)	28.06.	Imitation Learning: Behavioral Cloning
01.07.	Dependable RL (Q&A: 08.07.)	05.07.	
08.07.	Exploration Strategies (Q&A: 15.07.)	12.07.	
15.07.	? + Q&A		

Syllabus

Tips & Tricks: How to succeed

- Play with the jupyter notebooks as I suggest in the videos
- Keep in mind: look lesson in time then you can also benefit from the Q&A meetings!
- Attend the exercises and implement them
- Q&A sessions are open to everybody! (even though you do not an exam) just like under normal conditions.

Exam

- Currently, the plan is to have **oral** exams.
- Scheduling and logistics of the exams still TBD.