KTH ROYAL INSTITUTE OF TECHNOLOGY

FDD3359: Reinforcement Learning PhD Level Course

Danica Kragic, Alexander Kravchenko, Ali Ghadirzadeh, Hang Yin, Alexis Linard, and Chris Pek

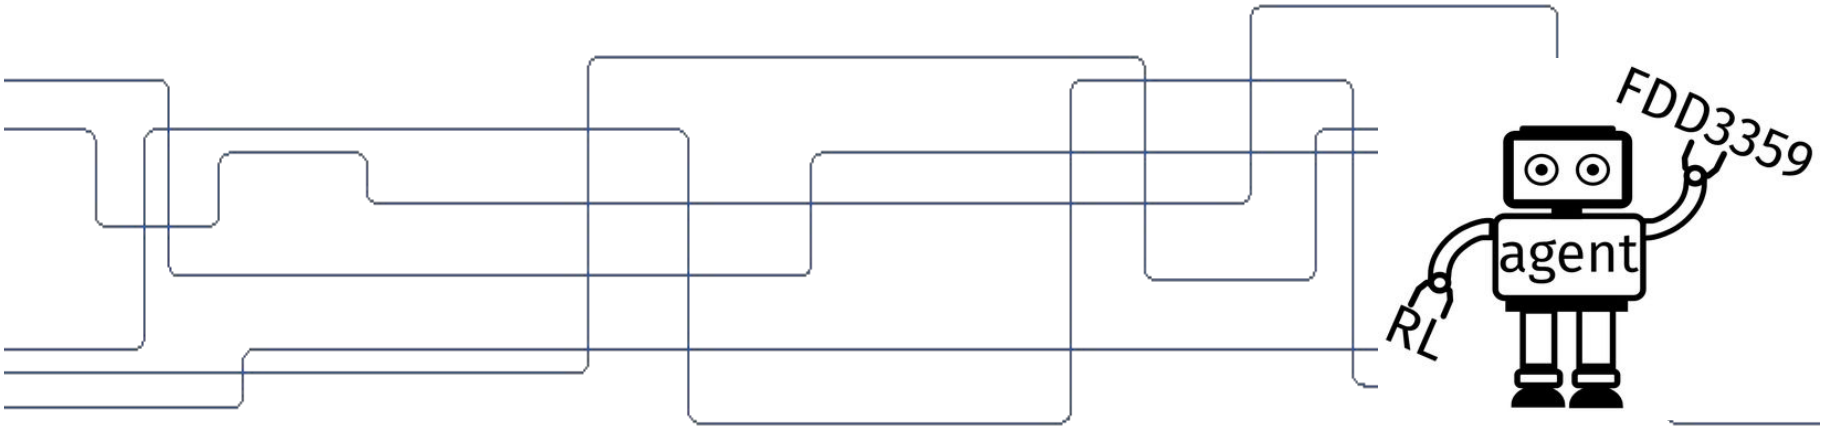

Welcome to this year’s course round

Prof. Danica Kragic Jensfelt

Alexander Kravchenko

Ali Ghadirzadeh

Hang Yin

Alexis Linard

Chris Pek

Participants in this year's course round

Participants in this year's course round

How familiar are you with RL?

26 Antworten

Participants in this year's course round

Your expectations (excerpts):

- Survey state-of-the-art and work on a more practical RL problem
- Revise existing knowledge, deeper insight into research, some more insight into implementation details as I have mostly used out of the box implementations (haven't written one myself).
- The get a deeper understanding of the theoretical background of RL.
- Acquire a toolbox of RL techniques to strengthen the therapy process provided by the robot in our lab.
- help me to understand and apply RL approaches in my research area.
- An introduction to the field of RL suitable to beginners of RL, and progression to the point of being capable to apply RL in practical scenarios outside of the course.

Course objectives

After this course, you will be able to:

1. describe the basic concepts and recent advances of RL
2. explain RL problems and their challenges
3. apply RL concepts to an existing problem formulation
4. analyze RL performance with respect to your problem formulation
5. assess how to constrain and improve your RL
6. implement a complex RL system for your own research

Passing criteria

- This course is a 6 credits course

- To pass this course, you need to:
  - work in a group of three students on an RL project applied to your research
  - present the results to the other groups (15min)
  - moderate a discussion on your results (10min)
  - actively participate in the discussions of other groups
  - submit your presentation and code

- When you pass this course, we will automatically sign you up at doctoral studies (no Ladok registration required!)

Organization

- Lectures take place every Wednesday from 2PM to 3:30PM
- Zoom link (we will record the lectures): https://kth-se.zoom.us/j/62401028768?pwd=VmdYN013TjZEQjVEeTlmT21YcEFHdz09
- Material will be posted on our public social page: https://www.kth.se/social/group/fdd3359-reinforcemen/
- Group forming starts around Easter time
- Presentation day tbd, but end of June before Midsommar

Syllabus

The tentative schedule is as follows (plus/minus 1 lecture):

- Lecture 1 (Alex, 16.02.): Introduction to Reinforcement Learning I
- Lecture 2 (Alex, 23.02.): Introduction to Reinforcement Learning II
- Lecture 3 (Ali, 02.03.): Offline Reinforcement Learning
- Lecture 4 (Ali, 09.03.): Meta-learning in RL
- Lecture 5 (Hang, 23.03.): Data-efficiency for RL in control applications
- Lecture 6 (Alexis, 30.03.): Constrained RL with temporal logic constraints
- Lecture 7 (Alexis, 06.04.): Human-robot-interaction in RL
- Lecture 8 (Chris, 13.04.): Enforcing safety constraints in cont. RL systems
- Lecture 9 (Chris, 20.04.): Correcting RL policies using shielding
- Lecture 10 (Chris, 27.04.): Applications of RL for safe, robust and acceptable autonomous systems
- Lecture 11 (Alex, 04.05.): Application of RL to chemistry and biology
- Presentation Day (end of June, before Midsommar)

Your questions?