

Study Circle in Reinforcement Learning

Coordinator: Karl-Erik Årzén

Study Circle

- A study circle and not a course
- I know probably much less about RL than you do
- Active participation

Lectures and Meetings

- The University College London (UCL) course "Reinforcement Learning" by David Silver
 - 10 Video Lectures
 - Accompanying slides
 - Exercises
 - Code
- Meetings
 - Ideally once per week
 - One per lecture
 - Before each meeting you all watch the corresponding lecture and (in some cases) do some exercises
 - During the meetings we discuss the lecture and the exercise solutions

Lecture Topics

Lecture 1: [Introduction to Reinforcement Learning](#)

Lecture 2: [Markov Decision Processes](#)

Lecture 3: [Planning by Dynamic Programming](#)

Lecture 4: [Model-Free Prediction](#)

Lecture 5: [Model-Free Control](#)

Lecture 6: [Value Function Approximation](#)

Lecture 7: [Policy Gradient Methods](#)

Lecture 8: [Integrating Learning and Planning](#)

Lecture 9: [Exploration and Exploitation](#)

Lecture 10: [Case Study: RL in Classic Games](#)

Links

Richard S. Sutton and Andrew G. Barto: "[Reinforcement Learning: An Introduction](#)", The MIT Press

The RL "bible". Silver's lectures and examples follows this book quite closely. Read if there is something you you don't understand or want to know more

Szepesvari: "Algorithms for Reinforcement Learning",

More algorithm-oriented. Available on the web.

Lucian Busoniu, Robert Babuska, Bart De Schutter, Damien Ernst: "Reinforcement learning and dynamic programming using function approximators"

Written by control engineers from Delft, A lot on continuous domains. Available on the web

Dimitri Bertsekas: "Reinforcement Learning and Optimal Control"

Book in preparation. First 5 chapters on the web. Also others books by Bertsekas on Dynamic Programming are relevant

Exercises

- Multiple GitHub repositories with exercises and code
 - <https://github.com/dennybritz/reinforcement-learning>
 - Implementation of Reinforcement Learning Algorithms. Python, OpenAI Gym, Tensorflow. Exercises and Solutions to accompany Sutton's Book and David Silver's course.
 - 242 commits
 - <https://github.com/dalmia/David-Silver-Reinforcement-learning>
 - Notes for the Reinforcement Learning course by David Silver along with implementation of various algorithms.
 - 84 commits

Requirements

- Two ways to get credits
 1. Active participation at at least 7/10 course meetings
 - 5 credits
 2. One + doing an individual assignment
 - learn and present some more advanced topic, or
 - do and present a project
 - 8 credits

Meetings

- Nominally once per week
 - Fri 1 Feb 13:00
 - Mon 11 Feb 13:00
 - Mon 18 Feb 13:00
 - Fri 22 Feb 10:00
 - Fri 1 Mar 10:00
 - Fri 8 Mar 10:00
 - Fri 15 Mar 10:00
 - Fri 22 Mar 10:00
 - Fri 29 Mar 10:00
 - + additional dates later for your presentations

Until next meeting

- Go through Lecture 2 ” Markov Decision Processes”
- Work through the OpenAI Gym Tutorial from dennybritz