Introduction to
Deep Reinforcement Learning (RL)

Sofia University

Today's Menu:

- brief intro: what is RL and why is it useful?
- logistics: lecture dates and time, etc
- course prereqs
- course literature
- coding logistics
- final grade formation
- student seminar

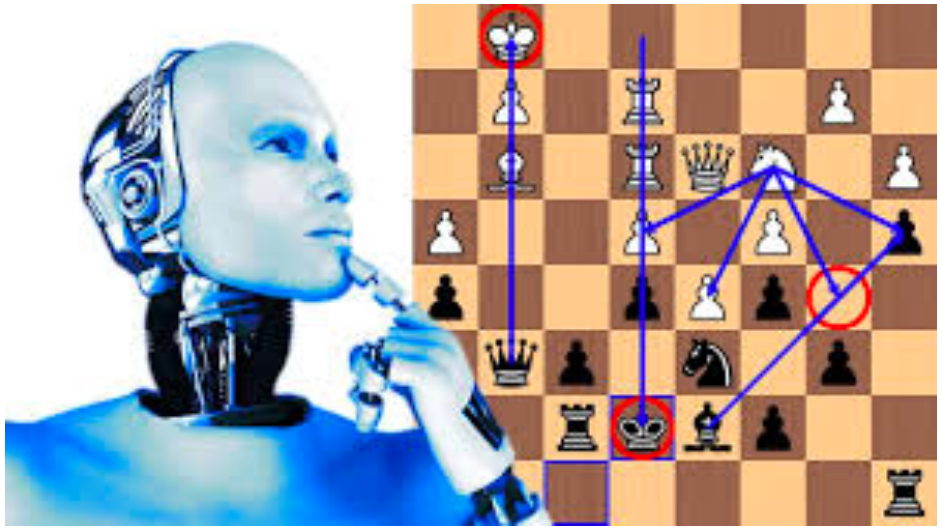

Marin Bukov

About me

Sofia University

Marin Bukov, PhD
mgbukov@phys.uni-sofia.bg

my research: quantum simulators,
non-equilibrium quantum dynamics,
applications of ML to physics

Reinforcement Learning (RL) as a branch of ML

Sofia University

→ Supervised Learning

- labelled data
- find approx. model which generalizes beyond known data

→ Unsupervised Learning

- unlabelled data
- find approx. probability distr. which generated the data

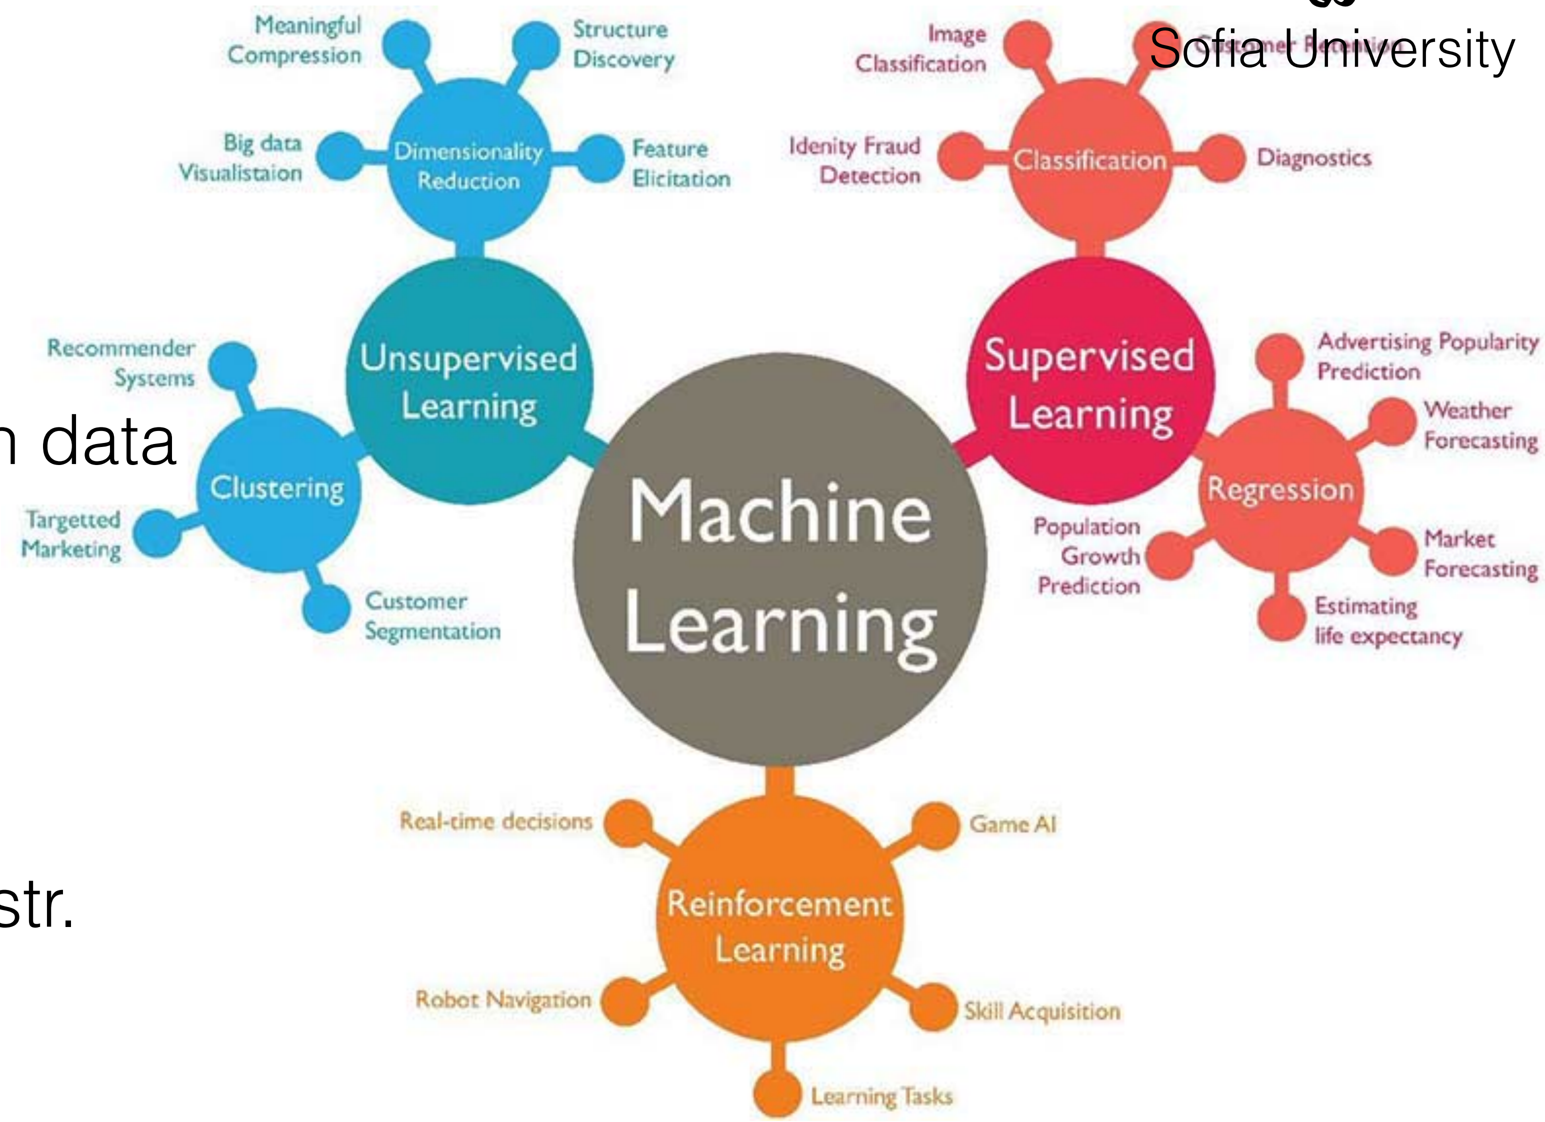

Reinforcement Learning (RL) as a branch of ML

Sofia University

Supervised Learning

- labelled data
- find approx. model which generalizes beyond known data

Unsupervised Learning

- unlabelled data
- find approx. probability distr. which generates the data

Reinforcement Learning

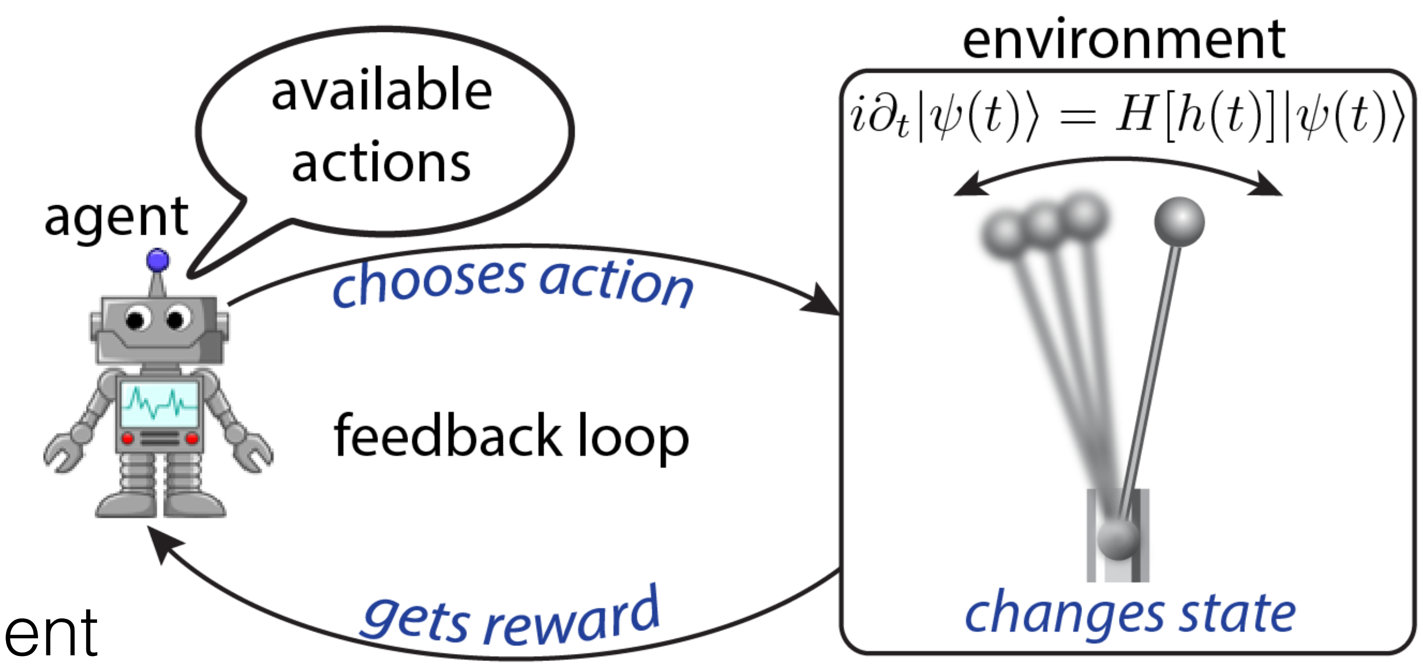

- agent learns strategy by interactions with its environment
- probability distribution which generates the learning data changes with time due to interaction with the environment

Marin Bukov

What is Reinforcement Learning?

Sofia University

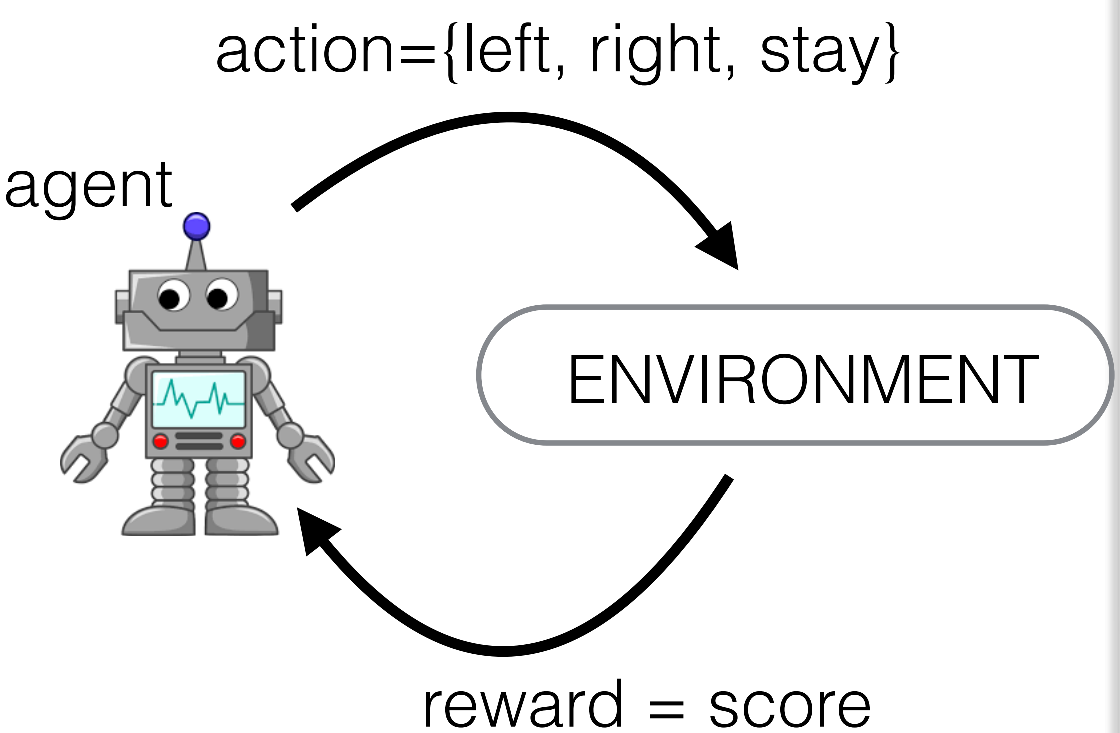

Marin Bukov

Silver et al., Nature 529 (2016) [Google DeepMind]

Reinforcement Learning
Quantum State Preparation

Sofia University

state preparation =
fidelity optimization

$$H(t) = -S^z - h_x(t)S^x$$

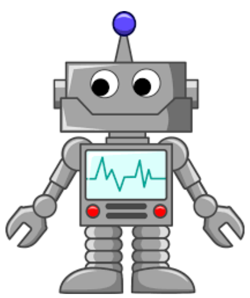

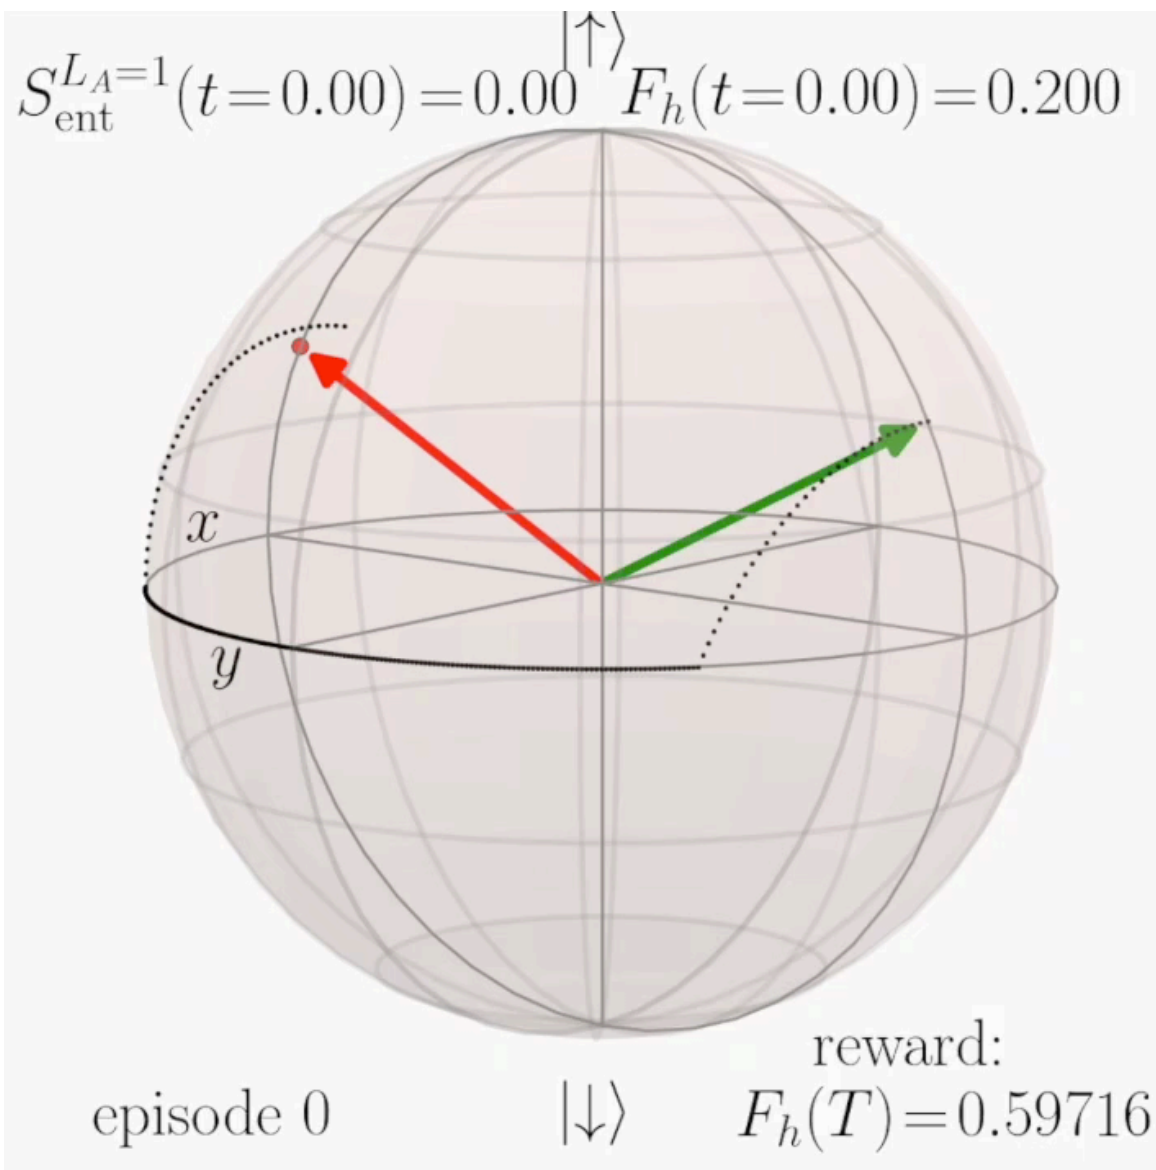

$h \in \{\pm 4\}$ bang-bang protocols

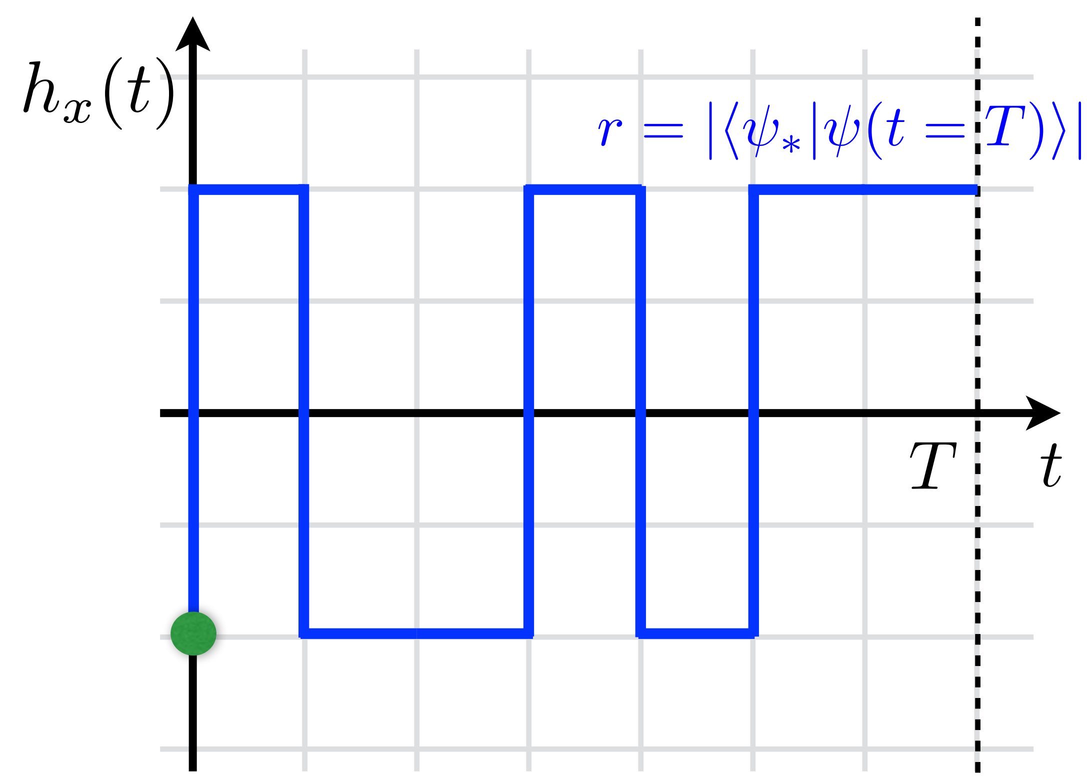

episode completed

MB et al, PRX 8 031086 (2018)
Foesel et al, PRX 8 031084 (2018)
Nui et al, NPJ Quantum Information 5 (2019)
Xiao-Ming et al, arXiv:1902.02157
Jun-Jie et al, arXiv:1901.08748

Marin Bukov

Survey Results

Introduction to
Deep Reinforcement Learning

Sofia University

Logistics:

- lecture:
    - 3-hour block, weekly
- coding session:
    - 2-hour block, weekly
- student seminar:
    - 2-hour block, weekly

Course website: **http://quantum-dynamics.phys.uni-sofia.bg/**

Google group: **rl_class_wise_20-21@googlegroups.com**

- need a volunteer for a group manager

FIX DATES AND TIMES

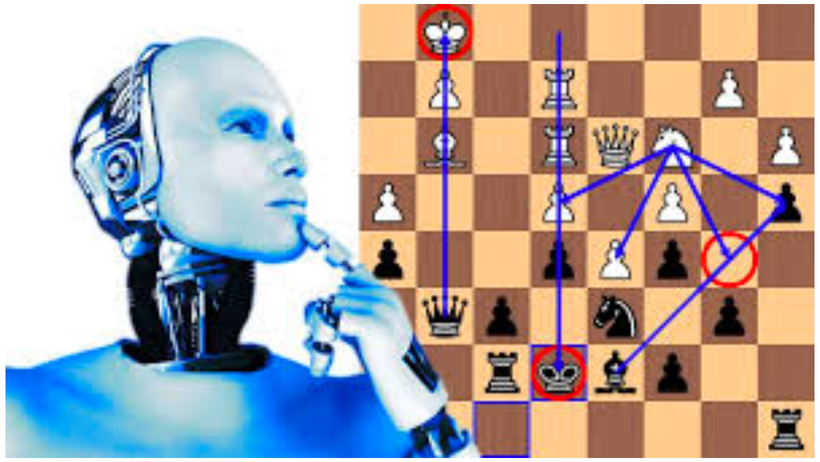

Marin Bukov

Introduction to Deep Reinforcement Learning

Sofia University

Target Audience:

- STEM (Science, Technology, Engineering, Math) students: bachelor, master, PhD

Prerequisites:

- linear algebra, analysis in many variables (calculus), probability theory
- coding: some familiarity with python (or equivalent)
- English: lectures and coding session will be exclusively in English
- *physics:* no physics background required, but it can be useful to understand some of the examples
- deep learning: no background required, but prepare to catch up (material will be provided, more later on)

Introduction to Deep Reinforcement Learning

Sofia University

Syllabus: the course has two parts

- Recap: Gaussian mixture models, Markov processes, Monte Carlo sampling, Ising models: *1 week*
- Basics of RL (no deep learning): *5 weeks*
  - concepts: states, rewards, actions, policies, value functions, etc.
  - Dynamic Programming
  - Monte-Carlo methods
  - Temporal Difference Learning
  - RL with function approximation
- Deep RL: *rest of the class*
  - policy gradient, actor-critic, deep Q-learning, etc.

Literature:

- main textbook: reading assignments (approx. one chapter per week)
    - first half: Sutton and Barto
    - second half: Berkeley video lectures by Sergey Levine

- catch-up with deep-learning: extra reading assignments
    - intro to ML: Mehta, et al
    - neural networks: Michael Nielsen's online book

- internet blog posts, seminar talks, etc.
    - be **extra careful** with any info found online: can be wrong!

Coding:

- python: numpy, spicy
- JAX (TensorFlow, PyTorch, etc.)
- Jupyter Notebooks
- Google Colab

Introduction to
Deep Reinforcement Learning

Sofia University

Final Grade: mimic scientific research —> scientific project
(can become the bedrock of your diploma thesis)

work in groups of 2-3 students (scope of project will be adjusted)

- **Project Proposal *(due in week 6): 15 % of final grade***
  - describes & motivates problem you want to study
  - describes methodology: which algorithm(s), why, literature (papers)
  - contains sub-projects, defines milestones
  - risk assessment: what can go wrong; what do you do then (contingency plan)?
  - estimate of required resources (compute hours to train model, etc.): where do you plan to run the training?

Introduction to Deep Reinforcement Learning

Sofia University

Final Grade: mimic scientific research —> scientific project

- **Project Execution *(week 6 — end of semester): 50% of final grade***
  - should be done **by the end of the semester**
  - final written project: format of Physical Review Letters
    - 4 pages, double column,
    - max 3750 words (excluding abstract and acknowledgements)
  - will provide a LaTeX template, but can use Word or whatever else.
  - contains:
    - introduction: what, why, how, key findings
    - body: numerical experiments, results, their interpretation
    - outlook: how do you advance the field, what open questions are left

Final Grade: mimic scientific research —> scientific project

- **Peer Review: *10% of final grade***
  - should be done **within one week time!**
  - critically assess and evaluate other projects (will provide criteria)
  - double-blinded:
    - author names will be hidden to reviewers
    - reviewer names will be hidden to authors
  - write a referee report (1 page)
  - each student will review two projects; instructor will review all projects

Final Grade: mimic scientific research —> scientific project

- **Project Presentation:** ***15% of final grade***
  - should be done **once peer review is over!**
  - prepare a scientific presentation of your project
  - 25 min + 5–10 min for questions
  - authors will receive the ***peer review results after the presentation***
  - presentations will be attended by all students; we will also make them open to faculty, friends, etc.

Introduction to Deep Reinforcement Learning

Sofia University

Final Grade: mimic scientific research —> scientific project

- **BONUS:** ***extra 10% to the final grade***
  - only if all of: proposal, project, review report, and presentation are done in **English**

Who are you?

Sofia University

Introduce:

- yourself
- your background
- your motivation to join the RL class
- **what kind of final project do you imagine working on?**

Google group: rl_class_wise_20-21@googlegroups.com

Student Seminar

Sofia University

GOAL:

- **learn to read, assess, scrutinize, and present scientific literature**
- **learn to go beyond textbooks**
- **figure out where concepts came to be, and how they evolved**

- student(s) present a scientific paper in RL (e.g., application to physics, CS, etc.)
- students can choose a paper of their interest (I will also make suggestions)
- everyone reads the paper; speaker has to go in-depth
- can be related to final project
- **everyone** asks question**s**
  - *weekly or biweekly (depending on interest)*
  - *3 ECTS (exam is the presentation): math students need to check*