Topics in Reinforcement Learning

Shangtong Zhang

What is RL?

- Sequential decision making problems

$$A_t \sim \pi(\cdot \mid S_t)$$

$$R_{t+1}, S_{t+1}$$

What can RL do?

AlphaGo (Silver et al. 2016)

(https://www.popularmechanics.com/technology/a19844/googles-alphago-ai-wins-first-round-against-go-champion/)

What can RL do?

AlphaStar (Vinyals et al. 2019)

(https://www.youtube.com/watch?v=nbiVbd_CEIA)

What can RL do?

OpenAI Five (OpenAI et al. 2019)

(https://www.linkedin.com/pulse/science-behind-openai-five-just-produced-one-greatest-jesus-rodriguez/)

What can RL do?

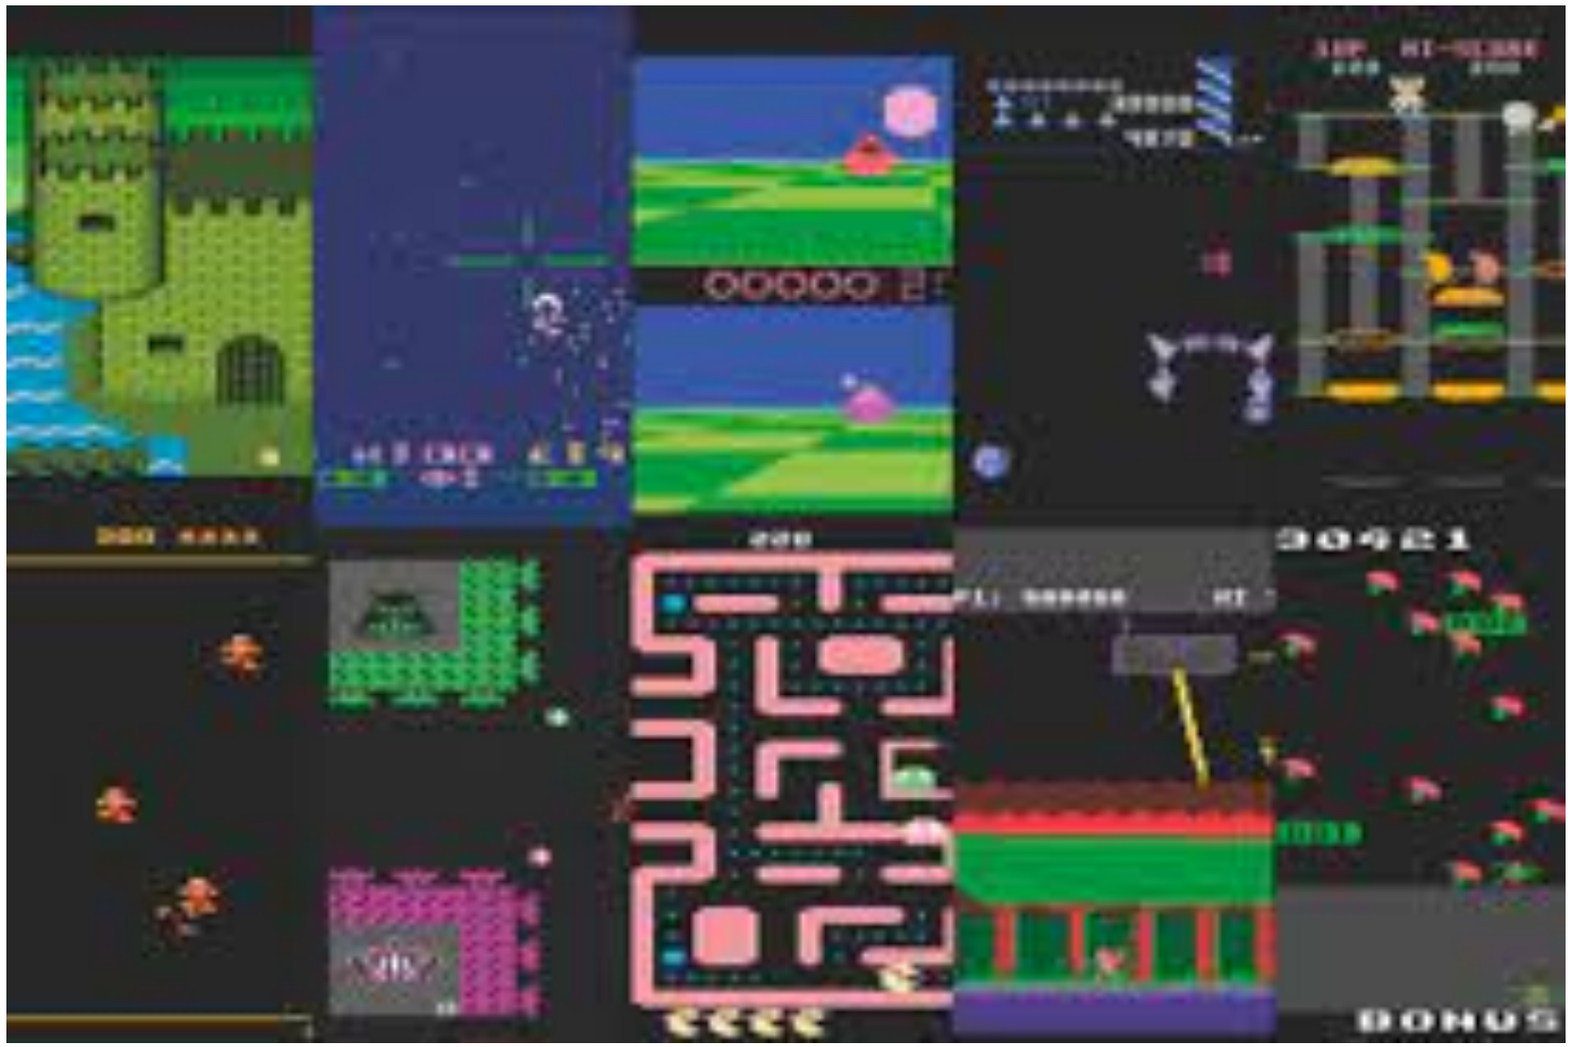

Atari games (Mnih et al. 2015)

(https://www.retrogamer.net/top_10/top-ten-atari-2600-games/)

What can RL do?

Backgommon (Tesauro, 1995)

(https://en.wikipedia.org/wiki/Backgammon)

What can RL do beyond games?

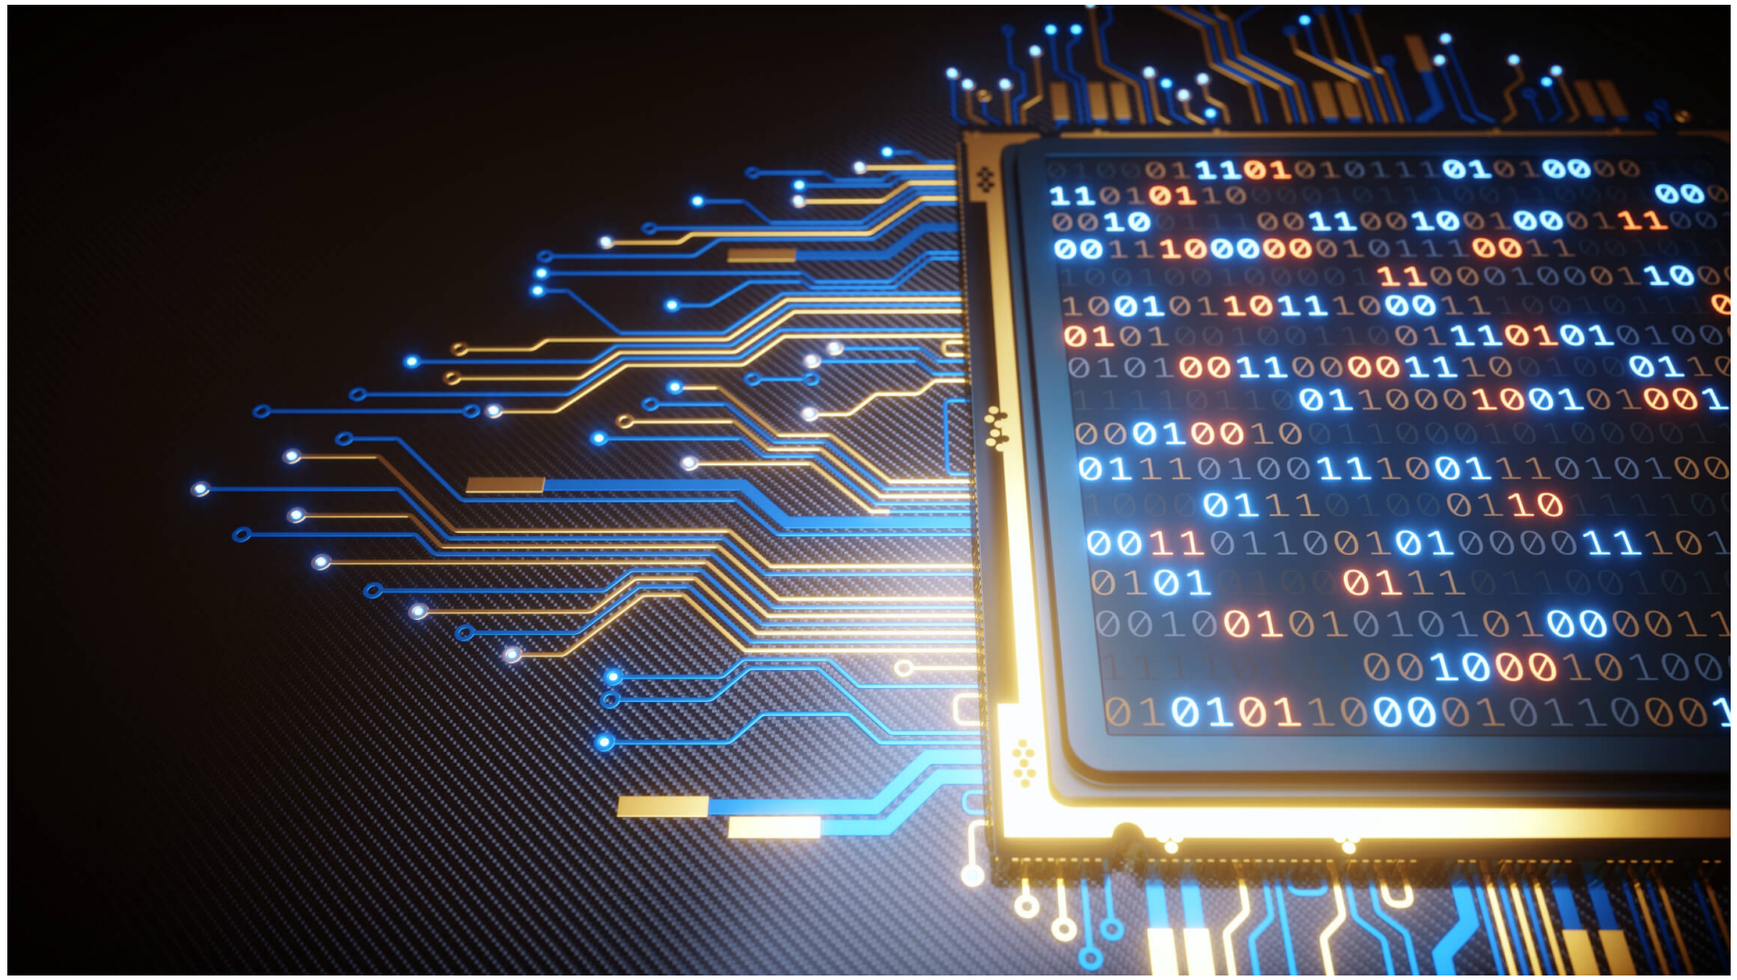

Chip design (Mirhosenini et al., 2021)

(https://www.asicnorth.com/offerings/)

What can RL do beyond games?

Navigation of stratospheric balloons (Bellemare et al., 2021)

(https://slate.com/technology/2014/12/project-loon-how-googles-internet-balloons-are-actually-working.html)

What can RL do beyond games?

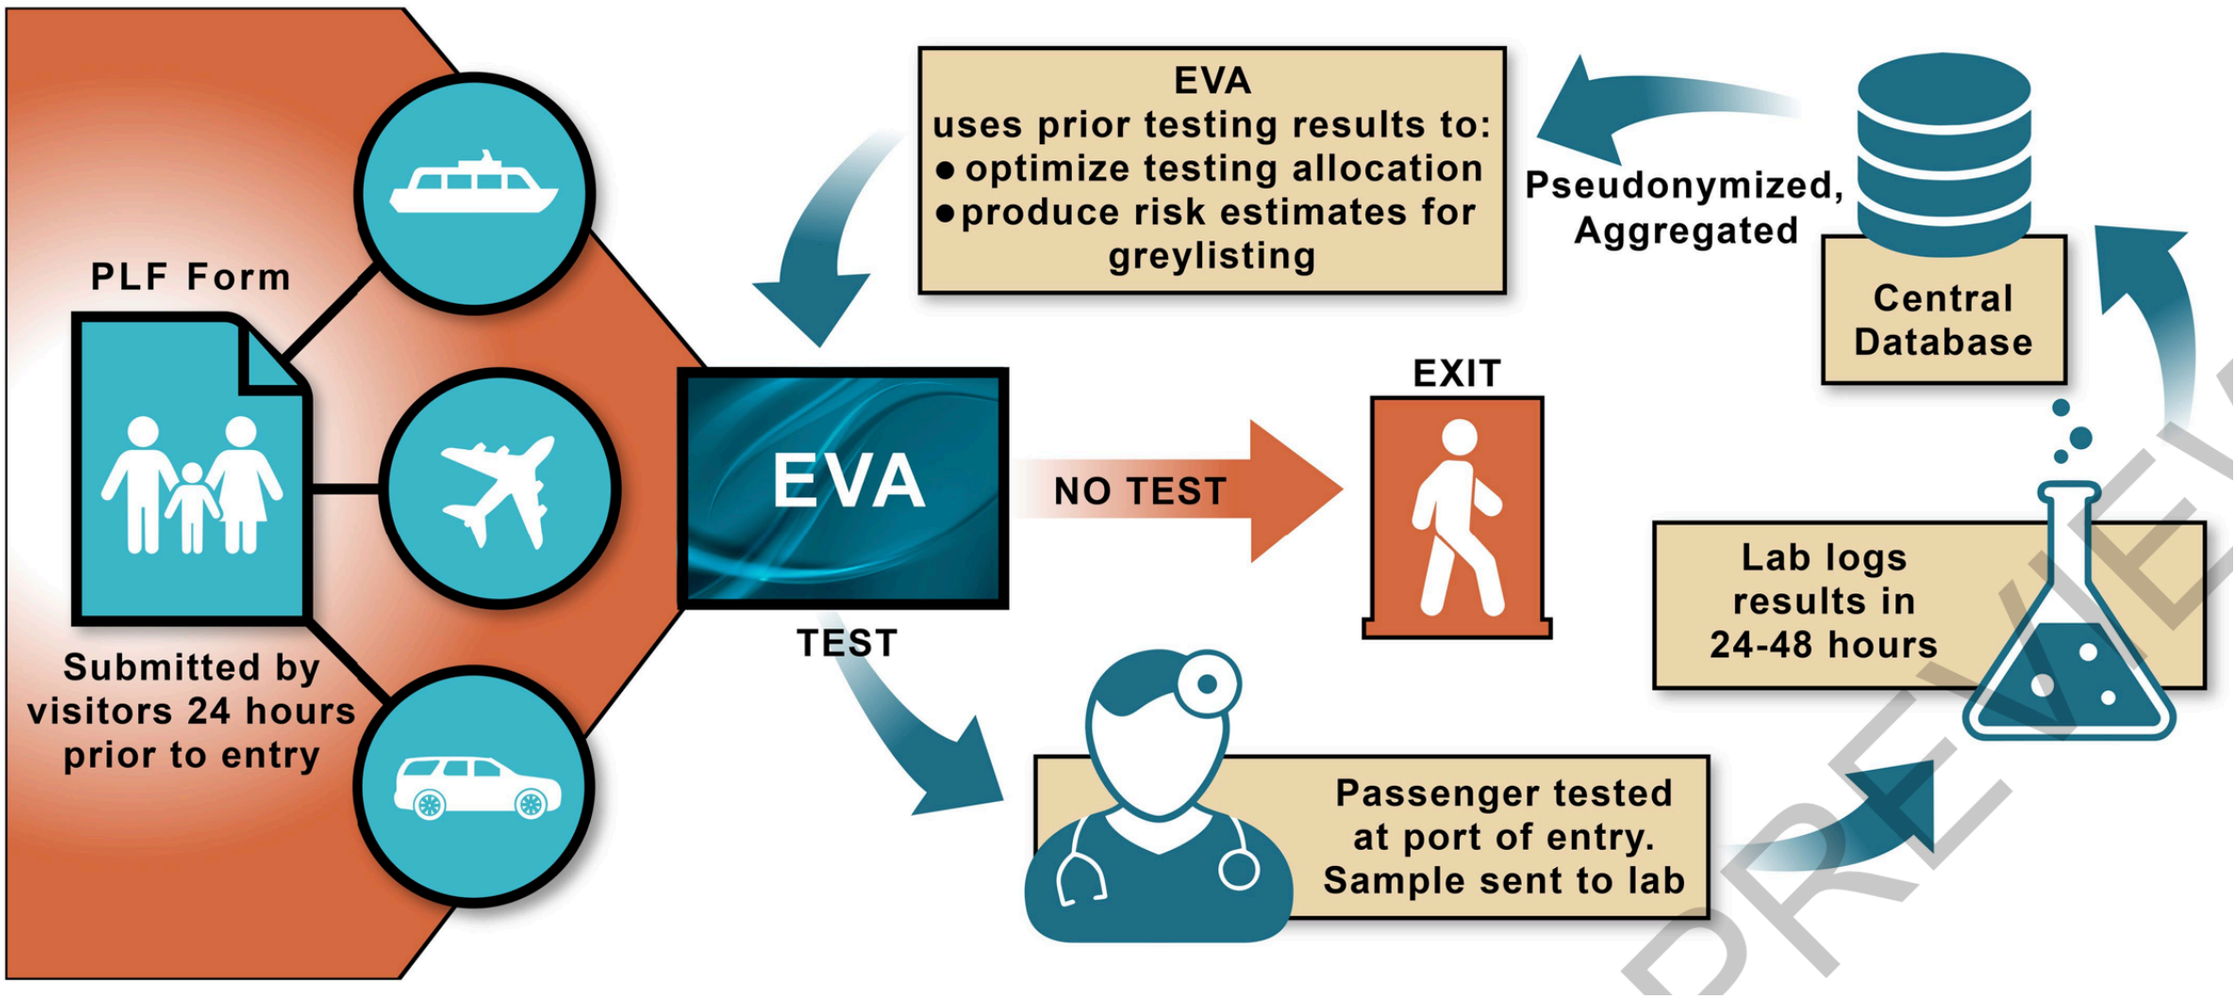

Covid-19 testing allocation (Bastani et al., 2021)

What can RL do beyond games?

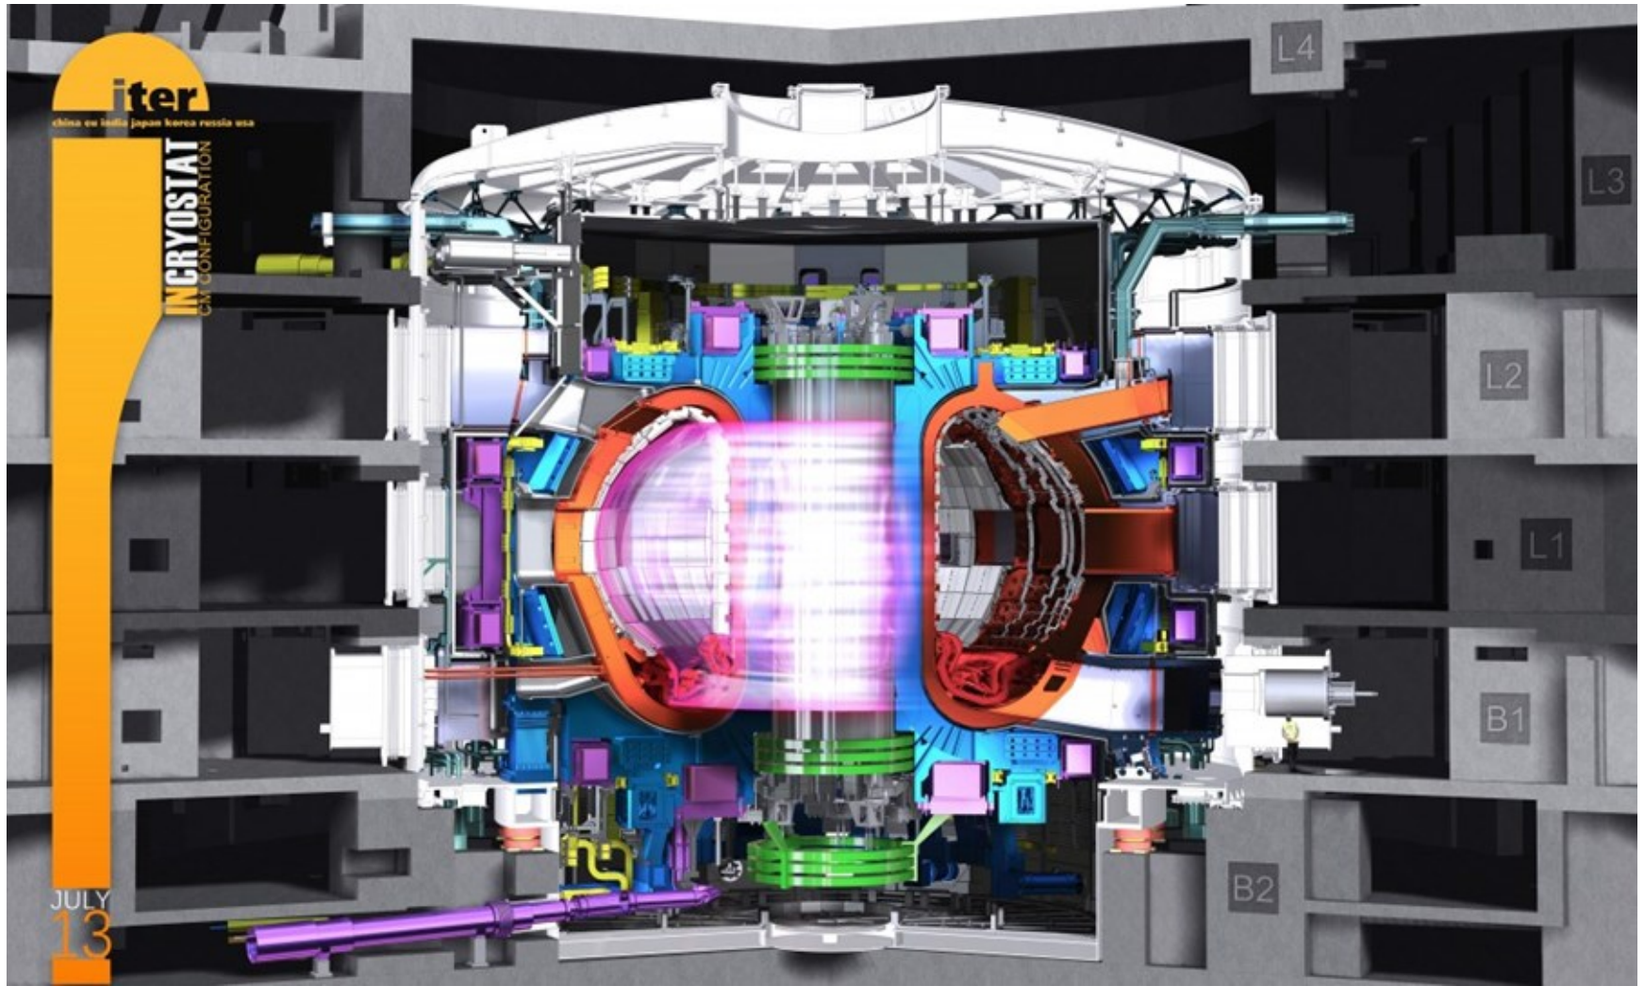

Nuclear fusion control (Degrave et al., 2022)

(https://www.iter.org/mach/Tokamak)

What are “topics” of this course?

- Focusing on theories
- Necessary background for working on RL theories
- Prerequisite:
  - Undergraduate level linear algebra
  - Undergraduate level probability

What are “topics” of this course?

- Focusing on theories
- Necessary background for working on RL theories
- Asymptotic analysis of some RL algorithms (proof)
- Finite sample analysis of some RL algorithms (proof)

What are NOT “topics” of this course?

- Deep RL
- Large scale experiments

Why do you want to take this course?

- Depth over breadth
- Learn any RL related topic quickly
- ML/RL engineers/PhDs for RL applications: DQN, A3C, DDPG, PPO, …
- PhDs: more complicated analysis, …

How to survive this course?

- Project oriented (70%)
  - Non-research project
  - Research project
- Reading assignment (30%)
  - Write a short summary and ask some questions

Non-research project

Reproduce a theoretical RL paper

- I compiled a list of papers
- Understand one
- Rewrite the proof in your own words (and notations if necessary), make it easier to read, fix “gaps”, submit a write-up
- Deliver a presentation
- You must do it solo

Non-research project

Reproduce a theoretical RL paper

- Pros:
  - Bounded workload
  - Easier
- Cons:
  - No innovation involved
    (You can of course propose some extension and transform to a research project)

Research project

Submit to some conferences (e.g., ICML)

- I compiled a list of ideas
- Understanding a paper
- Do the extension I provided (or your own)
- Write a paper and make a submission
- Solo or in a small group

Research project

Submit to some conferences (e.g., ICML)

- Cons
  - Serious commitment, more workload
    (Start now, don’t wait until the last moth)
- Pros
  - A peer-reviewed paper in a good venue (w. h. p.)
  - Risk-free:
    I provide only “good” idea
    easy transformation to the a non-research project
  - More advise from me

Which kind of project should I choose?

It’s up to YOU!

- There is no difference in terms of grading for this course
- Research project is more rewarding for your career but is harder
- Rotation with me / work with me -> research project
  (Email me after this lecture for rotation)
- You are always welcome to propose your own project
  (but talk with me first)

Research projects are limited

- It is nontrivial to find ideas that are good enough, easy enough and can fit into one term
- Take it serious

- Course website https://shangtongzhang.github.io/teaching/cs6501_fall_22/index

Q & A