


Yifan Zhu

Curriculum Vitae

459 Lagunita Drive
Stanford, CA, 94305
United States

+1 (650) 250 9557
zhuyifan@stanford.edu
fanzhuyifan.github.io



Research Interests

Reinforcement Learning

Estimating epistemic uncertainty

Estimating epistemic uncertainty, or uncertainty about the environment, is key to making good decisions in exploration. I am interested in developing principled and scalable methods to represent and estimate epistemic uncertainty, and to use them to guide exploration.

Improving sample-efficiency

Current reinforcement learning algorithms have achieved incredible success in computer games, showcasing superhuman performance. However, this achievement often comes at the cost of requiring an exceptionally large amount of data. For example, training the state-of-the-art A3C agent on Atari games requires hundreds of millions of frames. This inefficiency is a major obstacle for applying reinforcement learning to real-world problems, where data is often scarce. I am interested in developing methods to improve the sample efficiency of reinforcement learning algorithms, *by learning from more than just the reward signal*, and *by efficient exploration*.

Education

2023–Present **Ph.D. in Electrical Engineering**, *Stanford University*

○ **Advisor:** Benjamin Van Roy

2022–2023 **Master of Science in Electrical Engineering**, *Stanford University*, GPA 4.26

2018–2023 **Bachelor of Science in Mathematics**, *Stanford University*, GPA 4.16

○ Awarded the [J. E. Wallace Sterling Award for Scholastic Achievement](#).

○ Elected to [Phi Beta Kappa](#).

Awards and Honors

2019 **William Lowell Putnam Mathematical Competition (Putnam)**

○ Member of Stanford Team, Ranked 3rd Nationwide (2019)

○ Individual Rank 29.5 out of 4229 Contestants (2019)

○ Solution for A4 (2019) selected for the annual Putnam report

Publications and Preprints

Preprints

- [1] Saurabh Kumar, Henrik Marklund, Ashish Rao, Yifan Zhu, Hong Jun Jeon, Yueyang Liu, and Benjamin Van Roy. “Continual learning as computationally constrained reinforcement learning”. In: *arXiv preprint arXiv:2307.04345* (2023).

- [2] Hong Jun Jeon, Yifan Zhu, and Benjamin Van Roy. "An Information-Theoretic Framework for Supervised Learning". In: *arXiv preprint arXiv:2203.00246* (2022).

Publications

- [3] Qingxi Meng, Shubham Chandak, Yifan Zhu, and Tsachy Weissman. "Reference-free lossless compression of nanopore sequencing reads using an approximate assembly approach". In: *Scientific Reports* 13.1 (2023), p. 2082.

Experiences

Teaching

2022-2023 **SCA for *Dream It, Build It***, *Stanford Sophomore College*, Stanford

- In a three-week program for 16 Stanford rising sophomores, acted as TA in the classroom setting, mentored students on their individual projects, planned community building activities, and managed a \$2000 class budget.
- Received fantastic feedback from the students.

Internships

2021 **Cloud Software Intern**, *Intel*, Shanghai

- Optimize and profile cloud related AI workloads.
- Identify areas for optimization in architecture and hardware design.

2020-2021 **Data Scientist**, *ConvertLab*, Shanghai

- Design and build recommendation systems.
- Explore AI and data-powered marketing strategies.

Other

2023-Present **KDE developer**

Make occasional contributions to KDE projects.