

Kavosh Asadi

kavosh@alumni.brown.edu | (401) 660-7023 | US Permanent Resident | [Google Scholar](#) | [Homepage](#)

Education

- 2015-2020 Doctor of Philosophy in Computer Science, Brown University
advisor: [Michael L. Littman](#)
- 2013-2015 Masters in Computer Science, University of Alberta
advisor: [Richard S. Sutton](#)
- 2008-2013 Bachelors of Engineering, University of Tehran

Technical Interests

The alignment problem in large language models (LLMs)
Designing LLMs that can plan with long horizons
Sequential decision making (more specifically, the reinforcement-learning problem)
Applications of reinforcement learning to large-scale problems, in particular conversational agents, virtual and augmented reality, and improving generative AI using human feedback
Optimization algorithms for deep reinforcement learning, and planning in reinforcement learning
Variational approaches to generative AI such as Diffusion Models and VAEs

Industrial Experience

- 2020-Present Applied Scientist at Amazon. My work at Amazon was comprised of 3 areas:
 - I worked on advancing fundamental AI research. In particular, I started a research program to study value-function-approximation in reinforcement learning through the lens of optimization. Our efforts came into fruition by a) publishing multiple papers at NeurIPS about this topic and b) applying these advancements to internal products.
 - I focused on a few applications, most notably aligning LLMs. This effort included developing a simple but effective rejection-sampling approach, and also rethinking, researching, and implementing various ingredients of reinforcement learning from human feedback (RLHF) such as the generalized advantage estimator in the critic module. In a different project, I focused on tackling congestion control in AWS networks using imitation learning. We explored approaches that tackle the distribution-shift problem which commonly manifests itself in machine learning.
 - I focused on making it easier for students and engineers to get started in AI. This came into fruition by co-authoring the RL chapter of a fantastic book called [Dive into Deep Learning \(D2L\)](#).
- 2016 & 2017 Microsoft Research (MSR) PhD Internship
mentor: [Jason D. Williams](#)
At Microsoft I worked on training generative AI systems (before it was cool) to enable conversational agents that solve various tasks! We were among the first group to apply deep RL to the dialog domain. Our approaches were somewhat different than more recent work in that we were using the LSTM architecture (rather than Transformers), and we used sparse goal-based reward signals (rather than learning rewards from human feedback). The work was published at ACL, a leading NLP Conference, and was also deployed internally.

AI Coding Skills

Extensive experience in implementing reinforcement learning and deep learning techniques with different libraries, notably PyTorch, Tensorflow, and Keras. Proficient in leveraging the Hugging Face library for more advanced tasks such as fine-tuning in language modeling and other NLP tasks.

Advanced experience in Cloud Computing, particularly with Amazon EC2, including training large-scale models with billions of parameters on cloud infrastructure using the DeepSpeed library for GPU optimization.

Proficient in various programming languages, especially Python. Advanced experience in C, C++, and Java. Solid understanding of basic concepts and data structures in computer science.

Conference Papers

- 2024 Zuxin Liu, Jesse Zhang, **Kavosh Asadi**, Yao Liu, Ding Zhao, Shoham Sabach, Rasool Fakoore, "TAIL: Task-specific Adapters for Imitation Learning with Large Pretrained Models", International Conference on Learning Representations (ICLR)
- 2023 **Kavosh Asadi**, Shoham Sabach, Yao Liu, Omer Gottesman, Rasool Fakoore, "TD Convergence: An Optimization Perspective", Conference on Neural Information Processing Systems (NeurIPS)
- 2023 **Kavosh Asadi**, Rasool Fakoore, Shoham Sabach, "Resetting the Optimizer in Deep RL: An Empirical Study", Conference on Neural Information Processing Systems (NeurIPS)
- 2023 Omer Gottesman, **Kavosh Asadi**, Cameron Allen, Samuel Lobel, George Konidaris, Michael L. Littman, "Coarse-Grained Smoothness for Reinforcement Learning in Metric Spaces", Conference on Artificial Intelligence and Statistics (AISTATS)
- 2022 **Kavosh Asadi**, Rasool Fakoore, Omer Gottesman, Taesup Kim, Michael L. Littman, Alexander J. Smola, "Faster Deep Reinforcement Learning with Slower Online Network", Conference on Neural Information Processing Systems (NeurIPS)
- 2022 Martin Klissarov, Rasool Fakoore, Jonas Mueller, **Kavosh Asadi**, Taesup Kim, Alexander J. Smola, "Adaptive Interest for Emphatic Reinforcement Learning", Conference on Neural Information Processing Systems (NeurIPS)
- 2021 Rasool Fakoore, Jonas Mueller, **Kavosh Asadi**, Pratik Chaudhari, Alexander J. Smola, "Continuous Doubly Constrained Batch Reinforcement Learning", Conference on Neural Information Processing Systems (NeurIPS)
- 2021 **Kavosh Asadi**, Neev Parikh, Ronald E. Parr, George D. Konidaris, Michael L. Littman, "Deep Radial-basis Value Functions for Continuous Control", Conference on Artificial Intelligence (AAAI)
- 2021 Erwan Lecarpentier, David Abel, **Kavosh Asadi**, Yu Jinnai, Emmanuel Rachelson, Michael L. Littman, "Lipschitz Lifelong Reinforcement Learning", Conference on Artificial Intelligence (AAAI)
- 2019 Seungchan Kim, **Kavosh Asadi**, Michael L. Littman, George D. Konidaris, "Deepmellow: Removing the Need for a Target Network in Deep Q-learning", International Joint Conference on Artificial Intelligence (IJCAI)
- 2019 David Abel, Dilip Arumugam, **Kavosh Asadi**, Yu Jinnai, Michael L. Littman, Lawson LS. Wong, "State Abstraction as Compression in Apprenticeship Learning", Conference on Artificial Intelligence (AAAI)
- 2018 **Kavosh Asadi**, Dipendra Misra, Michael L. Littman, "Lipschitz Continuity in Model-based Reinforcement Learning", International Conference on Machine Learning (ICML)
- 2017 Jason D. Williams, **Kavosh Asadi**, Geoffrey Zweig, "Hybrid Code Networks: Practical and Efficient end-to-end Dialog Control with Supervised and Reinforcement Learning", Annual Meeting of the Association for Computational Linguistics (ACL)
- 2017 **Kavosh Asadi**, Michael L. Littman, "An Alternative Softmax Operator for Reinforcement Learning", International Conference on Machine Learning (ICML)

Book Chapter

Dive into Deep Learning. Main Authors: Aston Zhang, Zack C. Lipton, Mu Li, Alexander J. Smola. I served as a co-author for the chapter on reinforcement learning.