NEEV PARIKH

neev.v.parikh@gmail.com | 401-632-2749 | neevparikh.com | neevparikh

EXPERIENCE

ML Research Engineer
Sep 2024 – Present
METR
Berkeley, CA

- Working on AI misalignment research (model organisms, chain-of-thought monitoring, etc.)
- Worked on developing state-of-the-art evaluations for automated AI R&D; see publications.

Software Engineer II
Jul 2022 – Sep 2024
Stripe
San Francisco, CA

- Built LLM and ML-based customer insights platform, with per-customer predictions and calculated business metrics.
- Built polished data-repair tooling for self-serve, developer use, which drove self-serve repairs from **60% to 90%**, increased system adoption from **40% to 60%**, and saved **120 hours** per year.
- Built high-performance, **5min** time-to-alert, automated testing system in Java, with instrumented metrics like availability and latency.
- Built terabyte-per-hour scale, Hadoop-based data pipelines in Scala Spark to ingest financial data in a double-entry, event-based, immutable log.

Research Engineering Intern
May 2021 – Aug 2021
Common Sense Machines
Boston, MA

- Implemented large-scale, auto-regressive Seq2Seq models for working with 3D geometry from images.
- Worked with Deepspeed to explore scaling options for **500M+** param models to feasibly scale existing sequence-based models.
- Implemented a Blender-based Gym environment for RL to optimize textures on a 3D model.
- Implemented a graphics algorithm to find surface patches in a 3D wireframe (Zhang et. al., 2013)
- Dockerized AWS pipeline to create cloud-independent dev/production environment.

Research Assistant
Jun 2020 – May 2022
Brown University
Providence, RI

- Worked on original research on unsupervised, representation learning and multi-task reinforcement learning; see publications.

EDUCATION

M.Sc. in Computer Science
Brown University
Aug 2018 – May 2022 (Concurrent)
GPA: 4.0

Advised by: **Prof. George Konidaris**

B.Sc. in Computer Science
Brown University
Aug 2018 – May 2022 (Concurrent)
GPA: 3.9 (*magna cum laude*)

Advised by: **Prof. Michael Littman**

PROJECTS

Hierarchical Doom
High-throughput, distributed RL project to train asyncronous PPO-Option Critic on the VizDoom environment
neevparikh/hierarchical-doom

IP/TCP
IP/TCP system on an abstracted virtual link layer in Rust, with split horizon and poison reverse.
neevparikh/ip-tcp

Volumetric Photon Mapping
Volumetric photon mapping by extending an open-source, Rust-based path tracer, based on Bitterli et. al. (presentation).
neevparikh/rpt

Spaceport
Tiny MacOS utility written in Rust (with Obj-C bindings) to partly replace the deprecated airport utility .
neevparikh/spaceport

PUBLICATIONS

- H. Wijk et al. (2024). *RE-Bench: Evaluating frontier AI R&D capabilities of language model agents against human experts*. arXiv: 2411.15114 [cs.LG]. URL: https://arxiv.org/abs/2411.15114.
- M. Merlin, S. Parr, et al. (May 2024). “Robot Task Planning Under Local Observability”. In: *Proceedings of the 2024 IEEE Conference on Robotics and Automation*.
- C. Allen, N. Parikh, and G. Konidaris (Dec. 2021). “Learning Markov State Abstractions for Deep Reinforcement Learning”. In: *34th Neural Information Processing Systems Conference 2021*.
- K. Asadi, N. Parikh, R. Parr, G. Konidaris, and M. Littman (Sept. 2020). “Deep Radial-Basis Value Functions for Continuous Control”. In: *35th AAAI Conference on Artificial Intelligence 2021*.
- N. Parikh*, Z. Horvitz*, N. Srinvasan*, A. Shah, and G. Konidaris (Oct. 2020). “Graph Embedding Priors for Multi-task Deep Reinforcement Learning”. In: *NeurIPS 2020. KR2ML Workshop*.
- M. Merlin, N. Parikh, E. Rosen, and G. Konidaris (May 2020). “Locally Observable Markov Decision Process”. In: *International Conference on Robotics and Automation. Workshop on Perception, Action, Learning*.