Eric Wang

732-492-2378 | eric_wang@college.harvard.edu | linkedin.com/in/ericliwang/ | https://github.com/ericwang1409

Education

Harvard University Cambridge, MA
AB in Computer Science and Mathematics. GPA: 4.0/4.0 September 2022 - May 2026
Coursework: Data Structures & Algorithms, Computer Systems, Machine Learning, ML Interpretability, Probability
Activities: Harvard AI Safety, Harvard Tech for Social Good, Harvard Computer Society, Harvard Climbing

Experience

Software Engineer intern May 2025 - Present
LinkedIn Sunnyvale, CA
- Implemented internal certification tool to automate testing and deployment workflows for system infrastructure software using Java and Kafka to reduce hundreds of hours of manual testing from on-call engineers
- Designed pub/sub listeners and orchestration logic to trigger regression test pipelines and deployments on new commit events, ensuring compatibility and stability across software versions

Research Intern January 2025 - May 2025
Redwood Research Berkeley, CA
- Conducted AI safety research on early signs of collusion abilities in LLMs between agents and monitors
- Developed evaluation methodologies that identified subtle signs of strategic behavior in AI agents, highlighting potential vulnerabilities in oversight mechanisms and wrote a research paper on the results

Computer Science Teaching Assistant January 2024 – December 2024
Harvard University Cambridge, MA
- Held office hours and taught weekly recitations for 40+ students, supporting problem sets and theoretical concepts
- Teaching Assistant for CS 121 Introduction to Theoretical Computer Science -Fall 2024
- Teaching Assistant for CS 51 Abstraction and Design in Computation - Spring 2024

Machine Learning Engineer Intern June 2024 - August 2024
Analog Devices San Jose, CA
- Built an end-to-end human activity classification model using PyTorch, pre-trained with self-supervised learning techniques, and fine-tuned on open-source accelerometer data, achieving an average accuracy of 99.3%
- Architected a data pipline to download, process, and clean over 24 TB of accelerometer data, utilizing AWS S3 for storage, EC2 for cloud computing, and parallel processing to speed up data processing sevenfold
- Coordinated closely with team members and customers through biweekly meetings to ensure the model met the desired requirements such as low latency and desired memory size

Projects

Main-Memory Optimized Column-Store Database | *C, perf, valgrind* September 2024 - January 2025
- Designed and implemented a cache-conscious column-store database system from scratch in C, optimizing for main-memory access patterns to achieve a 10x speedup over traditional row-store databases
- Added multithreaded shared scans, B-tree indexing, and hash joins to speed up query processing by 5x

Sparse Autoencoders Thresholding for Safety | *Python, PyTorch* September 2024 - January 2025
- Developed and analyzed sparse autoencoders (SAEs) trained on open-source LLMs, establishing them as a competitive method for real-time model response oversight compared to residual stream activation probing
- Demonstrated SAEs to be as effective (98%) and more robust than probing SAE and residual stream activations

TaiYo! Solver | *Python, Pymunk, Deep Q-Learning, PyTorch* November 2023
- Engineered a custom game from the ground up, utilizing PyGame for engaging gameplay interfaces and PyMunk for accurate physics simulations utilizing object-oriented programming
- Implemented a Deep-Q Learning algorithm to train a model to autonomously play the game using optimal actions over the state space, outperforming over 80% of human players based on comprehensive gameplay metrics
- Won most ambitious/best idea hack at HackWellesley with a team of 3

Technical Skills

Languages & Frameworks: Python, TypeScript (React/Next.js), C/C++, PostgreSQL, PyTorch, OpenAI API
Tools and Cloud: AWS (S3, EC2), Docker, Git, Linux