

Cody Ho

+1 (408) 391 8014 | codyho@cs.stanford.edu | github.com/aesrentai | linkedin.com/in/cody-ho-006b32154/

EDUCATION

Stanford University, MS in Computer Science, AI Track, GPA: 4.0 / 4.0

Expected June 2026

- Instructor of CS40 Cloud Infrastructure and Scalable Application Deployment
- Vice President of Stanford Applied Cybersecurity

Stanford University, BS in Symbolic Systems, Computational Social Science Track

Expected June 2026

EXPERIENCE

Apple, Machine Learning Engineer | Cupertino, CA

June 2024 - June 2025

- Applied cutting edge methods for improving LLM performance including SFT, DPO, and chain of thought prompting
- Generated synthetic data for training and evaluation of LLMs on domain specific tasks
- Deployed multiple complex production apps on internal Kubernetes clusters with LLM inference backends

OpenAI, Red Team Engineer | Remote (California, USA)

October 2023 - June 2024

- Assessed the security and potential for abuse of an early access version of the Assistants API
- Developed and deployed a novel red teaming framework used to pentest LLMs using malicious API responses
- Evaluated the understanding of GPT-3.5 and GPT-4 in custom scenarios involving domain specific APIs
- Tested the security of the containerized execution environment of the Code Interpreter tool

Tatsu Lab, Research Assistant | Stanford, CA

June 2023 - June 2024

- Evaluated LLM use of APIs, discovering LLMs can produce potentially harmful outputs over 20% of the time
- Evaluated the extent to which fine tuning can improve LLM behavior when using APIs
- Examined the failure modes of LLMs and investigated potential mitigations

Infosys, Machine Learning Intern | Remote (California, USA)

June 2021 - August 2021

- Applied ML techniques to identify and flag potentially dangerous events and incidents relating to the management of microservices
- Curated a dataset of past microservice incidents used to inform future customer decisions
- Created a management platform that used WebSocket APIs to communicate with backend ML models

AWARDS

First Place, DEFCON AI Village Generative Red Teaming Contest

August 2023

- Achieved first place out of over 5000 competition participants including LLM security experts
- Elicited harmful, abusive, or potentially risky outputs from the latest LLMs using both adversarial and nonadversarial inputs

First Place, National Collegiate Cyber Defense Competition

April 2023

- Achieved first place out of over 200 teams from across the US
- Defended a computer network from a team of professional hackers using IaC hardening frameworks and automation

PROJECTS

Heads BootROM

- Added support for Heads to the System76 Lemur Pro 10 which is the firmware of the laptop I use daily
- Dockerized and improved the reproducible build system to work be over twice as fast and more reliable

Stockfish Chess Engine

- Improved the data creation and data processing pipelines used in the training of the novel NNUE architecture
- Created and evaluated modifications and improvements to the alpha-beta search algorithm, including pruning and parameter tuning

Ansible Hardening Toolkit

- Developed a series of Ansible playbooks and roles that can be used to secure a mixed Linux/Windows network
- Successfully used this framework to win first place in the NCCDC competition

SKILLS

ML Skills

Fine tuning, Prompt engineering, Retrieval Augmented Generation, Vector DBs

ML Frameworks

Hugging Face Transformers and Datasets, LangChain, PyTorch, NumPy, OpenAI APIs

Languages and Frameworks

Python, Rust, Go, Typescript, C, C++, Kotlin, Java, Bash, HCL, Flask, React, Angular

Cloud Development

AWS, Docker, Kubernetes, Terraform, Ansible, Packer, cloud-init