

Kavel Rao

Email: kavelrao@cs.washington.edu GitHub: github.com/kavelrao

Education

University of Washington Allen School , M.S. Computer Science	June 2025
<ul style="list-style-type: none">GPA: 4.0/4.0, Combined Bachelors/Masters program	
Relevant Courses: Computer Architecture, Social Reinforcement Learning, Compilers	
University of Washington Allen School , B.S. Computer Science	March 2024
<ul style="list-style-type: none"><i>Best Senior Thesis Award, Magna Cum Laude</i>, GPA: 3.97/4.0, Washington NASA Space Grant ScholarOne of five students selected for departmental exchange program at ETH Zurich, Switzerland	
Relevant Courses: Machine Learning, Database Systems, Deep Learning, Distributed Systems, Systems Programming	

Experience

Software Engineer Intern – Databricks (earned return offer)	Jun 2024 – Sep 2024
<ul style="list-style-type: none">Built and integrated login session management gRPC service to replace legacy thick-client, enabling high QPS traffic with increased reliability and supporting seamless login product experience.Designed self-encrypted token format with critical metadata to reduce database writes while ensuring security.Created MySQL table schema and indices to optimize query patterns and support future database migration.	
Technologies Used: Scala, Protobuf, MySQL, Armeria, gRPC	
Undergraduate Researcher – Allen Institute for AI & UW CS (Advisor: Yejin Choi) Feb 2022 – Present	
<ul style="list-style-type: none">Inventing AI safety tools to make LLMs more robust and safeguard user-LLM interactions through novel jailbreaking methods, synthetic data curation, and training unsafe content detection models.Co-first-author of state-of-the-art WildGuard safety moderation model detecting harmful content and model refusals, accruing 8000+ downloads in first month on HuggingFace: https://huggingface.co/allenai/wildguardDeveloped self-iterative knowledge distillation method to produce compute-efficient language models outperforming larger scale general-purpose models on quality and diversity metrics by 62%.	
Technologies Used: PyTorch, Pandas, vLLM, Numpy, HuggingFace, Transformers	
Publications: Two papers in NeurIPS 2024, one paper in AAAI 2024, one paper in EMNLP Findings 2023	
Software Engineer Intern – Stripe (earned return offer)	Jun 2023 – Sep 2023
<ul style="list-style-type: none">Identified key merchant risk metadata through historical data queries and user interviews; surfaced in risk analyst tools to improve efficiency in fraud and credit reviews (millions/month), reducing fraud incident mitigation time.Implemented flexible labeling and filtering system in merchant review interface, enabling machine learning data pipelines to improve merchant review automation.	
Technologies Used: Ruby, TypeScript, ReactJS, MongoDB, Splunk, Trino SQL	
Teaching Assistant – NLP, Programming Languages, Wireless Comm.	Mar 2022 – Present
<ul style="list-style-type: none">Designed lessons and taught 20+ students functional programming and advanced OOP in OCaml, Racket, Java.Worked with students through class time and office hours to ideate, scope, and implement innovative final projects such as multi-channel walkie talkie, radio astronomy.	
Software Engineer, Part-time – Conversica (earned return offer)	Jun 2021 – Sep 2022
<ul style="list-style-type: none">Reduced company spending by \$100,000/year building Kubernetes resource auditor with deployment pruning, cutting cluster size by 30%. Original scope was one-time cleanup, but now adopted into MLOps process.Delivered infrastructure for AI-powered chat based on BRD, scalable to 100 concurrent sessions. Implementation included DynamoDB for context and config storage, REST-based Django API for flow and business logic, Locust load testing, and model inference autoscaling with Sagemaker endpoints.Integrated and deployed 3rd-party semantic search service by building REST API wrapper for gRPC protocols. Will be used on front page of customer chat services.	
Technologies Used: Python, Django, AWS, Docker, Kubernetes, GitLab CI, Terraform, Jira	

Technical Qualifications

Languages (Proficient): Python, Java, C, SQL

Languages (Familiar): C++, Ruby, OCaml, Racket, Bash, JavaScript, TypeScript, HTML, CSS

Tools: PyTorch, Pandas, HuggingFace, Docker, Kubernetes, Django, Git, Terraform, Linux/UNIX, LaTeX