

Vinay Hiremath

+44 7404 060577
+1 248 660 0622
mail@vinayh.com
github.com/vinayh
vinayh.com/about

Education

- 2018 - 2021 **ETH Zürich and University of Zürich, MSc in Neural Systems and Computation, Zürich, Switzerland.**
- Took courses spanning statistical learning theory, probabilistic AI, deep learning, artificial and biological neuronal networks, and developmental and systems neuroscience
 - Participated in the Zürich AI Alignment reading club, and served as the first English-language editor for an engineering departmental publication (AMIV Blitz) leading to significantly more bilingual content
- 2013 - 2017 **University of Michigan, BS in Computer Science, Cognitive Science, Ann Arbor, MI, USA.**
- Dual major in Computer Science (Major GPA: 3.9/4.0) and Cognitive Science with a minor in Mathematics
 - Coursework included machine learning, algorithms, NLP, information retrieval, combinatorics, and real analysis
 - Enrolled in the LS&A Honors Program and received a major four-year scholarship from the college

Experience

- Jan. 2024–
current **Independent Grantee - ML Engineering, Lightspeed Grants.**
- Started preliminary user research into limitations of existing mechanistic interpretability tooling to identify how to better empower researchers working on reverse-engineering neural networks for safer AI systems
 - Began developing a prototype web app to make interpretability research more accessible and shareable
- May 2023–
June 2023 **Teaching Assistant, Alignment Research Engineering Accelerator (ARENA), London, UK.**
- Contributed to and taught a bootcamp on ML research engineering for AI safety to a cohort of professionals and university students, writing and updating exercises and answering student queries
 - Gathered and implemented participant feedback to improve the program's effectiveness and content, which covered transformers/LLMs, mechanistic interpretability, reinforcement learning, and distributed training
- Oct. 2022–
Apr. 2023 **Planned temporary leave.**
- Took a period of planned leave due to personal circumstances, followed by returning to a new role
- Aug. 2022–
Sep. 2022 **Teaching Assistant, Redwood Research - ML for Alignment Bootcamp (MLAB), Berkeley, CA, USA.**
- Helped to create and teach an intensive curriculum on ML engineering for AI alignment (covering transformers, mechanistic interpretability, distributed training, etc.), empowering students to later pursue AI safety careers
 - Took the initiative to rewrite the final day's content on short notice to introduce stable diffusion, newly released at the time, which students appreciated as an informative and engaging way to end the course
 - Noted code debugging was a common bottleneck for students, and opted to give a lecture on debugging for ML which improved programming productivity
- Jan. 2020–
Nov. 2021 **Instructor, TechSpark Academy, Zürich, Switzerland.**
- Designed and taught programming concepts to students aged 6-17 in after-school courses, introducing students to hands-on projects not available in school curriculum that elicited widespread interest
 - Topics included Python, AI/ML fundamentals, game design, and basic programming/logic, depending on age
- Jan. 2021–
Sep. 2021 **Researcher, ETH Zürich and University of Zürich - Institute of Neuroinformatics, Zürich, Switzerland.**
- Completed a masters thesis in biologically-plausible machine learning supervised by Prof. Benjamin Grewe, Johannes von Oswald, and João Sacramento
 - Investigated variants of the Gated Linear Networks architecture, identifying modifications that could enable scaling to larger tasks while maintaining sample/power efficiency advantages and relevance to neuroscience
- Sep. 2020–
Dec. 2020 **Teaching Assistant, ETH Zürich, Zürich, Switzerland.**
- Led discussion sections, addressed student queries, and helped write and grade assignments and exams for the *Learning in Deep Artificial and Biological Neuronal Networks* course taught by Prof. Benjamin Grewe
 - Course content explored and compared mechanisms in deep learning models and systems from neuroscience
- May 2018–
Aug. 2018 **Engineering Intern, Transcend.io, San Francisco, CA, USA.**
- Developed backend functionality as the first employee at a startup providing GDPR compliance as a service
 - Added external API integrations to automate personal data management for users of the platform's clients
- Sep. 2016–
Dec. 2017 **Research Assistant, University of Michigan EECS, Ann Arbor, MI, USA.**
- Worked on a multimodal question-answering task, implementing and evaluating character and scene recognition models on the MovieQA multimodal dataset, supervised by Prof. Rada Mihalcea and Prof. Jia Deng

- Jan. 2016– **Teaching Assistant**, *University of Michigan EECS*, Ann Arbor, MI, USA.
- Dec. 2017
- Assisted in teaching the *Introduction to Machine Learning* (EECS 445) course covering machine learning theory and implementations of topics ranging from traditional ML approaches to deep learning
 - Led weekly discussion sections, held office hours, and helped create and grade problem sets, exams, and projects; received positive feedback on learning outcomes from student surveys
- May 2016– **Engineering Intern**, *SalesforceIQ*, Palo Alto, CA, USA.
- Aug. 2016
- Improved reliability of backend systems company-wide by devising and implementing a strategy to add dynamic configuration management, health checks, and Docker containerization to services across multiple datacenters
- May 2015– **Engineering Intern**, *FarmLogs*, Ann Arbor, MI, USA.
- Aug. 2015
- Added a widely-requested feature for hyper-local nitrogen prescriptions for row farms using data science models
 - Implemented functionality in Clojure, interfacing with external APIs, satellite/soil datasets, and frontend clients
- Oct. 2014– **Research Assistant**, *University of Michigan EECS*, Ann Arbor, MI, USA.
- Dec. 2014
- Worked on a research project supervised by Dr. Richard L. Lewis and Dr. Satinder Singh Baveja, implementing domain-agnostic RL algorithms based on work by Sutton and Barto
- March 2014– **Computer Consultant**, *University of Michigan - Housing/Dining*, Ann Arbor, MI, USA.
- Aug. 2014
- Improved university-wide efficiency of cafeteria/dining point-of-sale systems by creating a backend and web interface to automate requests for interface changes for individual point-of-sale terminals
- June 2012– **Research Assistant**, *High School Honors Science Program, Michigan State University*, East Lansing, MI, USA.
- Aug. 2012
- Worked with Dr. Mitchell McGrath to measure and outline differential gene expression in sugar beet germination to improve crop output, writing Python utilities to process and analyze gene sequencing results

Activities

- Aug. 2023– current Building side projects including an interactive language model interpretability tool, a transport departures web/mobile app (React, Swift, Flask) and a life calendar web app (React, TypeScript/Node.js, Firebase).
- Feb. 2024– Facilitated an AI Safety Fundamentals discussion group introducing students and professionals to research agendas for aligning AI systems with intended goals.
- Mar. 2024
- Nov. 2021– Grantee at the Centre for Enabling Effective Altruism Learning and Research. Helped organize the SERI Existential Risk Conference 2022, participated in the AGI Safety Fundamentals course, and worked on web app side projects.
- May 2023
- Sep. 2017– Member of the Michigan Student Artificial Intelligence Lab (MSAIL) board. Helped to plan research April 2018 presentations, outreach for career opportunities, and reading groups discussing machine learning research.
- Sep. 2013– Part of a team that built a Top Five Hack at the MHacks Fall 2015 hackathon out of several hundred Dec. 2015 entries. Also attended several other hackathons including HackMIT, Hack the North, and PennApps.
- Sep. 2014– Member of the Michigan Hackers and MHacks 6 Core Team. Created slides and programming assignments Aug. 2015 for a Node.js workshop, and created a backend portal for sponsors to search through and filter university hackathon participants for recruiting purposes.
- Sep. 2014– Member of the technology team for optiMize, a student-run social innovation seed fund. Maintained May 2015 the organization's website, building a backend for handling team logistics and newsletter submissions.