

Harsh Kumar

+1-416-919-8537 | harsh@cs.toronto.edu | harsh-kumar.com | [Google Scholar](#) | [LinkedIn](#) | [GitHub](#)

PhD candidate (defending Fall 2026) working on LLM evaluation, post-training, and human–AI interaction. I study how everyday AI use shapes people, and build the evals and interventions informed by what we find.

EDUCATION

University of Toronto

Ph.D. in Computer Science (expected Fall 2026) — Advisor: Ashton Anderson

Toronto, ON

September 2021 – Present

Vellore Institute of Technology

B.Tech. in Computer Science and Engineering

Tamil Nadu, India

July 2016 – July 2020

EXPERIENCE

Graduate Research Assistant

University of Toronto

September 2021 – Present

Toronto, ON

- Study how everyday LLM use shapes human capacities (learning, creativity, and well-being) through large-scale randomized experiments; published across interdisciplinary HCI, ML, and NLP venues as well as education and mental-health venues.
- Designed and ran randomized experiments on how LLM assistance affects creative ideation and convergence (CHI 2025) and on guiding student LLM use in classrooms (CSCW 2024), identifying conditions under which AI help builds vs. substitutes human capability.
- Earlier: built contextual-bandit systems to personalize mental-health support and LLM agents for behavior-change coaching (AAAI 2024, CSCW 2025).

Research Intern (Applied Science)

Microsoft, Viva Insights

May 2025 – August 2025

Redmond, WA

- Designed LLM-based agents and evaluation approaches to help business leaders understand AI adoption, map organizational risks and opportunities, and act toward sustainable deployment (with Shamsi Iqbal).
- Collaborated with cross-functional partners to translate research findings into decision-relevant product directions; work led to a paper currently under submission.

Research Intern

Microsoft Research, Computational Social Science group

Summers 2023 & 2024

New York City, NY

- Designed controlled experiments and evaluation pipelines studying how LLMs may augment human cognition, focused on AI for education and measurable learning outcomes.
- Led one of the first large-scale randomized controlled trials of LLM math tutoring (AIED 2025); featured in Microsoft's *New Future of Work* report.

Software Engineer

J.P. Morgan Chase & Co.

January 2020 – August 2021

Mumbai, India

- Built and shipped an end-to-end production NLU system with user-feedback-driven improvement loops — model training, deployment, data pipelines, and CI/CD under real-world scalability and reliability constraints.

Research Analyst Intern

McKinsey & Company

May 2019 – January 2020

Gurugram, India

- Developed the vision system (visual SLAM; object detection and avoidance) for a semi-autonomous warehouse UAV.

PUBLICATIONS

Full list incl. workshop papers & preprints: harsh-kumar.com

- *Position: We Need Large Language Models Optimized For Our Well-Being.* **ICML 2026.** [link]
- *When AI Gives Advice: Evaluating AI and Human Responses to Online Advice-Seeking for Well-Being.* **CHI 2026.** [arXiv]
- *Investigating the Effects of LLM Use on Critical Thinking Under Time Constraints.* **CHI 2026.**
- *Human Creativity in the Age of LLMs: Randomized Experiments on Divergent and Convergent Thinking.* **CHI 2025** — **Honorable Mention (top 5%).** [doi]

- *Math Education with Large Language Models: Peril or Promise?* **AIED 2025**. [link]
- *Large Language Model Agents for Improving Engagement with Behavior Change Interventions: Application to Digital Mindfulness*. **CSCW 2025**. [doi]
- *Investigating the Role of Situational Disruptors in Engagement with Digital Mental Health Tools*. **CSCW 2025 — Honorable Mention (top 5%)**. [doi]
- *Platform-Based Adaptive Experimental Research in Education: Lessons Learned from the Digital Learning Challenge*. **LAK 2025**. [doi]
- *Guiding Students in Using LLMs in Supported Learning Environments: Effects on Interaction Dynamics, Learner Performance, Confidence, and Trust*. **CSCW 2024**. [doi]
- *Disclosures & Disclaimers: Investigating the Impact of Transparency Disclosures and Reliability Disclaimers on Learner-LLM Interactions*. **HCOMP 2024**. [link]
- *Supporting Self-Reflection at Scale with Large Language Models: Insights from Randomized Field Experiments in Classrooms*. **L@S 2024**. [doi]
- *Using Adaptive Bandit Experiments to Increase and Investigate Engagement in Digital Mental Health*. **AAAI 2024 (IAAI)**. [link]
- *Investigating the Role of Context in the Delivery of Text Messages for Supporting Psychological Wellbeing*. **CHI 2023 — Best Paper (top 1%)**. [doi]
- *A Text Messaging Intervention to Support the Mental Health of Young Adults: User Engagement and Feedback from a Field Trial*. **Internet Interventions, 2023**. [link]
- *Small Steps over Time: A Longitudinal Usability Test of an Automated Interactive Text Messaging Intervention*. **Journal of Affective Disorders, 2023**. [link]

SELECTED AWARDS & PRESS

- **XPRIZE Digital Learning Challenge** — US\$500,000 grand prize (Adaptive Experimentation Accelerator team); led development of the contextual-bandit personalization system used for large-scale adaptive classroom experiments (2023).
- **DARPA AI Tools for Adult Learning** — US\$250,000 award for QuickTA, an LLM-based learning tool combining curated content, adaptive RL algorithms, and intelligent tutoring; deployed to 1,500+ students (2023).
- **Press — research featured in:** Time, “*Are We Losing Our Minds to AI?*” (2026) • The Economist, “*Will AI Make You Stupid?*” (2025) • The New York Times, opinion on AI, coders & jobs (2025) • Science News, on our LLM-timing & critical-thinking study (2026) • Forbes, on using ChatGPT for learning (2024).

TECHNICAL SKILLS

LLM training & evaluation: supervised fine-tuning (full / LoRA), Hugging Face Transformers & TRL, vLLM inference, multi-GPU training (SLURM clusters), LLM-as-judge and human-evaluation design, benchmark construction
Experimentation: randomized experiments / A-B testing, causal inference, adaptive experiments (contextual bandits), crowdsourced studies (Prolific, MTurk), statistical analysis in R and Python
Engineering: Python, PyTorch, R, SQL, JavaScript; Flask/Django, Docker, CI/CD; production ML systems