

Smriti Singh

+1 (737) 217-3371 | smritisingh26@utexas.edu | linkedin.com/in/smriti-singh-ss2000 | smritisingh26.github.io | Mountain View, CA

ML Research Engineer specializing in agentic AI systems and LLM evaluation. Research focuses on fairness, reasoning, and reliability in large language models, with publications at NeurIPS, ICLR, NAACL, COLING, and ACL. Bridges applied science research and production ML systems deployment.

TECHNICAL SKILLS

ML and AI: Agentic AI, Multi-agent Systems, Agent Evaluation, RLHF, Observability, Production ML Systems, LLM Fine-Tuning, RAG, SFT, LoRA/QLoRA, LLM Evaluation, Data Curation, MCP, Prompt Engineering, NL-to-SQL, NLP, Transformers, LLaMA, LlamaIndex, VLMs, LLM Inference
Languages and Frameworks: Python, PyTorch, SQL, Git

Tools and Infrastructure: LlamaParse, Redis, Postman, Vercel, Google Colab, ML System Design, LLM Pipelines, Evaluation Frameworks, ChromaDB, Phoenix (Arize)

NVIDIA Stack: NeMo Agent Toolkit, NeMo Guardrails, NIM, OpenShell, Dynamo, NemoClaw

Product & Communication: Technical stakeholder communication, problem formulation, customer-facing product demos

SELECT PUBLICATIONS

- Reasoning or Rhetoric? An Empirical Analysis of Moral Reasoning Explanations in Large Language Models - Logical Reasoning in LLMs @ ICLR 2026 - Singh*, Kasat*, Chadha, Jain
- Born With a Silver Spoon? Investigating Socioeconomic Bias in Large Language Models - LLM Eval @ NeurIPS 2025 - Singh*, Keshari*, Jain, Chadha
- From Prejudice to Parity: A New Approach to Debiasing Large Language Model Word Embeddings - COLING 2025 - Rakshit*, Singh*, Keshari, Jain, Chadha
- Language Models (Mostly) Do Not Consider Emotion Triggers When Predicting Emotions - NAACL 2024 - Singh, Caragea, Li
- "Female Astronaut, Because Sandwiches Won't Make Themselves Up There": Towards Multimodal Misogyny Detection in Memes - WOAHS @ ACL 2023 - Singh*, Haridasan*, Mooney
- BREE-HD: A Transformer-Based Model to Identify Threats on Twitter - IEEE Access 2023 - Kumbale, Singh, Poornalatha, Singh
- "Hold On Honey, Men at Work": A Semi-Supervised Approach to Detecting Sexism in Sitcoms - ACL SRW 2021 - Singh*, Anand*, Chowdhury, Waseem

Google Scholar: 80 citations, h-index 6, i10-index 4 | scholar.google.co.in/citations?user=idOI-IAAAAAJ

WORK EXPERIENCE

Zacks Investment Research

Jun 2024 - Present

Lead ML Research Engineer (Apr 2025 - Present) | ML Engineer (Jun 2024 - Apr 2025) | Chicago, IL (Remote)

- Led ML system design for an agentic AI pipeline to generate initiation reports on macro-cap tickers; projected to generate \$1M+ in annual revenue.
- Led end-to-end development of a fine-tuned LLM to transform financial articles into tweets; produced 200+ officially posted tweets within 3 weeks.
- Built an automatic RLHF feedback pipeline for a finance-domain LLM, boosting model performance by 33% on domain-specific benchmarks.
- Designed a fact-checking student-teacher LLM chain for hallucination prevention, ensuring reliability of AI-generated financial content.
- Developed a robust LLM evaluation framework for complex NL-to-SQL tasks, improving accuracy by 41%.
- Stack: Agentic AI, RAG, LLM Fine-Tuning, RLHF, LlamaIndex, LlamaParse, VLMs, NL-to-SQL, Redis, Postman, LLM Evaluation

Esperanto Technologies

May 2023 - Aug 2023

Machine Learning Research Intern | Mountain View, CA

- Investigated ethical and moral capabilities of SOTA LLMs in socially nuanced situations; work contributed to research direction published at ICLR 2026.

The University of Texas at Austin

Aug 2022 - May 2024

Graduate Research Assistant | Austin, TX

- Thesis advised by Jessie Li; focused on LLM emotion perception and evaluation. Co-authored paper accepted at NAACL 2024.

Meedan

Jan 2022 - Jul 2022

NLP Researcher | San Francisco, CA

- Built an AI system identifying COVID-19 questions on social media with 91% accuracy; performed large-scale geodemographic analysis.

MITACS GRI

Jun 2021 - Sep 2021

Global Research Intern | Toronto, Canada

- Developed a neural embedding framework to analyze user anxiety patterns across digital platforms.

NOTABLE PROJECTS

Research Monitoring Agent | NeMo Agent Toolkit, NIM, ChromaDB, Phoenix

- 4-agent pipeline (Planner, Retriever, Summarizer, Evaluator) on NeMo Agent Toolkit to stay current on Gen AI research; persistent cross-session memory via ChromaDB deduplicates across runs; OTel observability via Phoenix.

Guardrailed Clinical Research Agent | NeMo Guardrails, NIM, OpenShell, NemoClaw

- Agentic POC with NeMo Guardrails input/output rails, OpenShell sandboxed tool execution, and NemoClaw policy enforcement; red-teamed against 30 adversarial probes across prompt injection, tool misuse, and PII exfiltration.

Slidelang | Next.js, Vercel, MCP [built after meeting with Andrew Ng, AI Fund]

- Deck-as-code platform where humans and AI agents co-author structured JSON specs that compile deterministically into editable, shareable slides; live at slidelang-by-smriti.vercel.app.

EDUCATION

The University of Texas at Austin

Aug 2022 - May 2024

M.S. Computer Science | Minor: Computational Linguistics | NSF Certified Portfolio in Ethical AI | GPA: 3.83/4.0

Manipal Institute of Technology

Jul 2018 - May 2022

B.Tech Information Technology | Minor: Big Data | GPA: 8.67/10.0

AWARDS AND LEADERSHIP

Honors: PyData Global Impact Scholar, Grace Hopper Scholar, EmTech MIT Scholar

Competitions: Google Hashcode Hub Winner, Smart India Hackathon College Winner, BITS Goa Research Contest Winner

Leadership: Founder - Women in AI Research Mentorship Lab (2026); Vice Chair - ACM-W Manipal; Mentor - Women in CS @ UT Austin

Speaking: Speaker at PyData Global and FruitPunch AI; AI Safety Blog on Substack; Reviewer - ICLR 2026 Workshops