

Smriti Singh

+1 (737) 217-3371 | smritisingsh26@utexas.edu | [linkedin.com/in/smriti-singh-ss2000](https://www.linkedin.com/in/smriti-singh-ss2000) | smritisingsh26.github.io | Mountain View, CA

Forward Deployed Engineer specializing in data-centric AI, partnering with frontier labs to evaluate agentic systems and produce the expert data that moves model performance. Bridges published research with production deployment, from customer discovery through shipped POC. Published at NeurIPS, ICLR, EMNLP, NAACL, COLING, and ACL.

TECHNICAL SKILLS

ML, AI and Product: Agentic AI, Multi-Agent Systems, Agent Orchestration, MCP, Customer Discovery, POC Design, Rapid Prototyping, Requirements Translation, Technical Storytelling, LLM Fine-Tuning, RAG, RLHF, SFT, LoRA/QLoRA, LLM Evaluation, Agent Evaluation, Prompt Engineering, NL-to-SQL, Transformers, LLaMA, LlamaIndex, VLMs, Production ML Systems

Languages, Frameworks and Tools: Python, PyTorch, SQL, Git, AWS, Kubernetes, Modal, Daytona, Harbor, Snowflake, Datadog, Arize Phoenix, LlamaParse, Redis, Postman, ChromaDB, NVIDIA NeMo Agent Toolkit, NeMo Guardrails, NIM, Dynamo, ML System Design, LLM Pipelines, Evaluation Frameworks

WORK EXPERIENCE

Snorkel AI

Forward Deployed Engineer | Redwood City, CA | Aug 2026 - Present

- Partnering with frontier AI labs to evaluate agentic pipelines and generate high-quality expert data that improves state-of-the-art model performance.

Zacks Investment Research

Lead ML Research Engineer (Apr 2025 - Jun 2026) | ML Engineer (Jun 2024 - Apr 2025) | Chicago, IL (Remote)

- Owned end-to-end product vision and technical execution for an agentic AI pipeline generating institutional-grade initiation reports on macro-cap tickers; projected to generate \$1M+ in annual revenue.
- Drove multi-week customer discovery with editors, analysts, and senior stakeholders to define AI success criteria; translated tribal knowledge into a structured 4-agent pipeline architecture (Planner, Retriever, Summarizer, Evaluator).
- Designed and shipped a three-phase agentic eval pipeline (fact-checking, historical depth, coherence) that reduced factual errors by 90%, plus an NL-to-SQL eval framework improving accuracy by 41%.
- Built an automatic RLHF feedback pipeline boosting model performance 33% on domain benchmarks; led fine-tuning of an LLM that produced 200+ officially posted financial summaries within 3 weeks.

Esperanto Technologies

Machine Learning Research Intern | Mountain View, CA | May 2023 - Aug 2023

- 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

Graduate Research Assistant | Austin, TX | Aug 2022 - May 2024

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

Meedan

NLP Researcher | San Francisco, CA | Jan 2022 - Jul 2022

- Built an AI system identifying COVID-19 questions on social media with 91% accuracy and geodemographic analysis.

MITACS GRI

Global Research Intern | Toronto, Canada | Jun 2021 - Sep 2021

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

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 - WOAHA @ 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: 95 citations, h-index 6, i10-index 5 | scholar.google.co.in/citations?user=idOI-IAAAAAJ

EDUCATION

The University of Texas at Austin

M.S. Computer Science | Minor: Computational Linguistics | NSF Certified Portfolio in Ethical AI | GPA: 3.83/4.0 | Aug 2022 - May 2024

Manipal Institute of Technology

B.Tech Information Technology | Minor: Big Data | GPA: 8.67/10.0 | Jul 2018 - May 2022

AWARDS AND LEADERSHIP

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

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

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

Speaking: PyData Global, FruitPunch AI; AI Safety Blog

Reviewer: ICLR 2026 Workshops, NeurIPS 2026