

ANIMESH JHA

✉ animjha@stanford.edu [in LinkedIn](#) [🔗 anime-sh](#)

EDUCATION

Stanford University

2024 – 2026

MS in Computer Science (AI Specialization)

Indian Institute of Technology, Kharagpur

2019 – 2023

Computer Science and Engineering (B.Tech) - Department Rank 3/76, Best Undergraduate Thesis

PUBLICATIONS

What to Forget in Unlearning?

[PDF]

Animesh Jha*, Arpandeeep Khatua*, Youssef Allouah, Sanmi Koyejo

MemFM Workshop, ICML 2026

MAEB: Massive Audio Embedding Benchmark

[PDF] [Code]

A Assadi, I Chung, ..., A Jha, ..., N Muennighoff, K Enevoldsen

Under Review NeurIPS 2026

Certified Unlearning for Neural Networks

[PDF] [Code]

Anastasia Koloskova*, Youssef Allouah*, Animesh Jha, Rachid Guerraoui, Sanmi Koyejo

ICML 2025

RL-Guided Data Selection for Language Model Finetuning

[PDF] [Code]

Animesh Jha*, Harshit Gupta*, Ananjan Nandi*

Reliable ML Workshop, NeurIPS 2025

Local NMPC on Global Optimised Path for Autonomous Racing

[PDF] [Code]

Animesh Jha*, Dvij Kalaria*, Parv Maheshwari* et al.

OCAR Workshop, ICRA 2021

RESEARCH EXPERIENCE

Efficient and Certified Machine Unlearning | Prof. Sanmi Koyejo | Stanford

August 2024 – Present

- Developed the first certified unlearning mechanism for general neural networks via differentially private post processing
- Using techniques from sketching, algorithmic stability, and compressed sensing to extend certified unlearning to large models

Compilation for Distributed Databases | Prof. Fredrik Kjolstad | Stanford

April 2025 – December 2025

- Represented relational algebra as tensor algebra, leveraging distributed tensor algebra to enable efficient query execution
- Designed a scheduling interface to optimize query execution and validate compliance with data movement policies

Robust Distributed Learning | Prof. Simon S. Du | UW | Undergraduate Thesis

2021 – 2023

- Designed and analyzed Byzantine Fault Tolerant algorithms for efficient distributed training with non convex loss functions
- Established lower bounds on the number of iterations for finding ϵ approximate critical points in the presence of adversaries

INDUSTRY EXPERIENCE

Cartesia AI | Intern - Serving Infrastructure | San Francisco

June 2025 – September 2025

- Designed modular task framework to decouple CPU/GPU scaling, optimizing resource utilization & reducing inference costs
- Reduced Batch Speech-To-Text latency by $> 75\%$ for long audio ($> 1\text{hr}$) via optimized chunking and in-memory caching

Rubrik, Inc. | SWE - Platform (Core Infrastructure) | Bangalore

June 2023 – August 2024

- Reduced upgrade time by 45% and saved 120 developer hours daily, through safe parallelization of orchestration steps
- Migrated 10TB/day ETL pipelines from AWS EMR to Kubernetes, boosting throughput by 35% and reducing costs by 67%

SELECTED PROJECTS

Curriculum Learning for Efficient LLM Reasoning | CS 329A (Self Improving Agents)

Autumn 2025

- Formulated minibatch sampling as Online Convex Optimization, used Mirror Descent for adaptive batches during training
- Finetuned Qwen3 via GRPO and adaptive sampling on MATH, reduced degenerate groups and achieved faster convergence

RL based Data Selection for Finetuning LLMs | CS 224R (Deep RL) | [Code]

Spring 2025

- Formulated data selection as a Sequential Decision Making problem, learnt policies via PPO to select training subsets
- Achieved accuracy close to full-dataset training while selecting only 5% of the original data, reducing training time by $2\times$

ACHIEVEMENTS

- Secured All India Rank **60** in the **Kishore Vaigyanik Protsahan Yojna** (SA, 2017) conducted by Government of India
- Awarded the **National Talent Search Examination Scholarship** (2017) by the Government of India.
- **National Finalist** (among the top 5 teams out of 25000+ participants) at the **Uber Hacktag 1.0 2021**.
- ACM ICPC Regionalist: Qualified for and placed **44th** out of 7000+ teams at the Gwalior Pune ICPC Regionals 2020.