

Ishaan Watts

Pittsburgh, PA | 412-925-7734

✉ iwatts@cs.cmu.edu

🌐 wattsishaan.github.io

📄 ishaan-watts-b7a679226

🎓 Google Scholar

EDUCATION

Carnegie Mellon University (CMU)

Master of Science in Machine Learning;

Coursework: Advance Intro to Machine Learning*, Probability & Statistics*

Pittsburgh, PA

December 2026

Indian Institute of Technology Delhi (IIT-D)

Bachelor of Technology in Engineering Physics; **Minor:** Computer Science; **CGPA:** 9/10

Coursework: Algorithms, Calculus, Digital Electronics, Computer Networks, Linear Algebra, Signals & Systems, Economics

New Delhi, India

May 2023

PUBLICATIONS

C=CONFERENCE, P=PREPRINT

- [P1] **Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities** Gemini Team, Google DeepMind (including Ishaan Watts) [Tech Report] [Paper]
- [C.4] **RTP-LX: Can LLM Evaluate Toxicity in Multilingual Scenario?** Adrian de Wynter, Ishaan Watts, et al. [AAAI '25] [Paper]
- [C.3] **PARIKSHA: A Large-Scale Investigation of Human-LLM Evaluator Agreement on Multilingual & Multi-Cultural Data** Ishaan Watts, Varun Gumma, Aditya Yadavalli, Vivek Seshadri, Swami Manohar, Sunayana Sitaram [EMNLP '24] [Paper]
- [C.2] **MAPLE: Multilingual Evaluation of Parameter Efficient Finetuning of LLMs** Divyanshu Aggarwal, Ashutosh Sathe, Ishaan Watts, and Sunayana Sitaram [ACL Findings '24] [Paper]
- [C.1] **MEGAVERSE: Benchmarking LLMs across languages, modalities, models and tasks** Sanchit Ahuja, Divyanshu Aggarwal, Varun Gumma, Ishaan Watts, Kalika Bali, and Sunayana Sitaram [NAACL '24] [Paper]

EXPERIENCE

Google DeepMind | Pre-Doctoral Researcher

July 2024 - July 2025

Projects: Gemini 2.5-Pro Contributor [P1], Multi-Model Composition, Model Stitching vs. Adapters for Continual Learning

- Explored composition of generalist & specialist LLMs for **continual domain adaptation & compositional tasks** (code × i18n)
- Improved **CALM** architecture & **beat LoRA** under varied budget constraints; presented poster at **Google Research Week '25**
- Designed a **routing-based composition** of 64 models for inference-time expert selection, beating a 2x larger monolithic model
- Engineered robust **JAX-based** training, evaluation, and bulk inference pipelines for composed models within **Gemini codebase**

Microsoft Research | NLP Research Intern

May 2023 - June 2024

Projects: Toxicity Detection Dataset [C.4], Indic LLM Evaluation Leaderboard [C.3], Multilingual Finetuning [C.2] and Benchmarking of LLMs [C.1], Education Bot [Media Coverage]

- Built culturally-nuanced **evaluation framework** in 10 Indic languages, assessing 30+ LLMs via **120k+** human/LLM judgments
- Released **38-language toxicity benchmark**, evaluating 10 S/LLMs on 8 harm type, exposing failures in detecting subtle biases
- Studied **QLoRA finetuning** on 90+ models, analyzing cross-lingual transfer and benchmarked 20 LLMs across **83 languages**
- Developed & deployed **Shiksha Copilot**, GenAI tool helping teachers create educational content, in **100+ schools** in India

Torch Investment Management | Machine Learning Engineer Intern

September 2022 - December 2022

Project: Stock Price Modeling via Deep Learning

- Optimized **LightGBM** model to predict Top-30 US S&P 500 stocks with 60% accuracy & led data modeling for Saudi market
- Innovated new-feature linking stock price with tweet sentiment by **topic filtering** (BART) & **sentiment analysis** (FinBERT)

Udaan | Data Scientist Intern

May 2022 - July 2022

Project: Holistic User-Embeddings via Graph Neural Networks

- Generated user-representation from buyer-seller interaction graph via multi-relational, multi-entity **hetero-graph autoencoder**
- Integrated embeddings into **fraud detection** pipelines, reduced Udaan dispute loss by **2.45%** and offered full-time position

SKILLS

Languages: Python, Java, SQL, R, Bash, C++, Matlab, \LaTeX **Technologies:** Git, Docker, RAG, CUDA, HPC, Vector DB, Spark

Libraries: PyTorch, TensorFlow, JAX, Numpy, Pandas, Matplotlib, Scikit-Learn, Snsrape, Streamlit, Langchain, OpenCV

HONORS & AWARDS

Narotam Sekhsaria (\$40,000), **KC Mahindra** (\$6,000) and **JN Tata Fellowship** (\$12,000) to pursue graduate studies at CMU

Peer Bonuses at Google DeepMind | **1st place** in Microsoft Global Hackathon '24 | **6th place** in Amazon ML Challenge '23

3x IIT-D Semester Merit Awards for being top 3 in department | **KVPY Fellowship** by Government of India, ranking top 2%