

Ishaan Watts

✉ wattsishaan18@gmail.com  [Google Scholar](#)  [Homepage](#)  [LinkedIn](#)  [Twitter/X](#)

EDUCATION

Carnegie Mellon University (CMU) – Master of Science Aug 2025 - Current
Masters in Machine Learning (MSML)

Indian Institute of Technology Delhi (IIT-D) – Bachelor of Technology July 2019 - May 2023
Major: Engineering Physics; Minor: Computer Science; Department Rank: 4/49 CGPA: 8.996/10

RESEARCH EXPERIENCE



Google DeepMind July 2024 - July 2025
Pre-Doctoral Researcher with [Dr. Partha Talukdar](#) Bengaluru, India
Projects: Gemini 2.5-Pro Contributor, Multiple Model Composition, Model Stitching vs. Adapters for Continual Learning




Microsoft Research May 2023 - June 2024
Research Intern with [Dr. Sunayana Sitaram](#) Bengaluru, India
Projects: Indic LLM Evaluation Leaderboard (PARIKSHA), Multilingual & Multicultural Toxicity Detection Dataset (RTP-LX), Analysis of Multilingual Finetuning of LLMs (MAPLE), Multilingual Benchmarking (MEGAVERSE), Shiksha Copilot




Griffith University May 2021 - July 2021
Research Intern with [Prof. Saiful Islam](#) Remote / Queensland, Australia
Project: Malware Detection using Machine Learning



PUBLICATIONS



C=CONFERENCE, P=PREPRINT, *=EQUAL CONTRIBUTION

[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]

[C.4] **RTP-LX: Can LLMs Evaluate Toxicity in Multilingual Scenarios?**   
Adrian de Wynter, [Ishaan Watts](#), et al. [AAAI 2025]
AAAI 2025 Artificial Intelligence for Social Impact Track

[C.3] **PARIKSHA: A Large-Scale Investigation of Human-LLM Evaluator Agreement on Multilingual and Multi-Cultural Data**   
[Ishaan Watts](#), Varun Gumma, Aditya Yadavalli, Vivek Seshadri, Swami Manohar, Sunayana Sitaram
Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing [EMNLP 2024]

[C.2] **MAPLE: Multilingual Evaluation of Parameter Efficient Finetuning of LLMs**  
Divyanshu Aggarwal*, Ashutosh Sathe*, [Ishaan Watts](#), and Sunayana Sitaram
Findings of the Association for Computational Linguistics: ACL 2024 [ACL Findings 2024]

[C.1] **MEGAVERSE: Benchmarking LLMs across languages, modalities, models and tasks**  
Sanchit Ahuja, Divyanshu Aggarwal, Varun Gumma, [Ishaan Watts](#), Ashutosh Sathe, Millicent Ochieng, Rishav Hada, Prachi Jain, Maxamed Axmed, Kalika Bali, and Sunayana Sitaram
Proceedings of the 2024 Conference of the North American Chapter of ACL [NAACL 2024]

HONORS & AWARDS

- **Narotam Sekhsaria Scholar** — awarded \$40,000 for graduate studies; top 25 out of 6,000+ applicants. 2025
- **K.C. Mahindra Scholar** — awarded \$6,000 for graduate studies; among top 100 of 2,500+ applicants. 2025
- Received **Peer Bonuses** at Google DeepMind for productive collaborations by colleagues & senior researchers. 2025
- Bagged **1st position** in Topic Challenge & **Honors** in Executive Challenge at Microsoft Global Hackathon. 2024

- Secured **6th position amongst 5,000** teams all over India in the Amazon ML Challenge 2023. 2023
- Granted **IIT-Delhi Semester Merit Award** for being amongst top 7% students in 3 out of 8 semesters. 2021-2022
- Achieved **Rank 1635 in JEE-Advanced & 99.94 percentile in JEE-MAINS** among 1.4 million candidates. 2019
- Recipient of **KVPY Fellowship** (Kishore Vaigyanik Protsahan Yojana) given by DST, Government of India. 2018

SELECTED PROJECTS

- **Model Stitching vs. Adapters for Continual Learning** July 2024 - Dec 2024
Advisors: [Dr. Prateek Jain](#), [Dr. Partha Talukdar](#) Google DeepMind
 - Explored composition of generalist & specialist LLMs to enable **continual adaptation** to new domain-specific data.
 - Improved performance over existing **CALM** technique and **beat LoRA** under different parameter & data budgets.
 - Curated **compositional tasks** (code X i18n) & presented findings as poster at Google Research Week [GRW'25].
- **Multiple Model Composition** Jan 2025 - July 2025
Advisors: [Dr. Prateek Jain](#), [Dr. Partha Talukdar](#), [Dr. Marc'Aurelio Ranzato](#) Google DeepMind
 - Designed a **routing-based composition** of generalist & expert LLMs with dynamic inference-time expert selection.
 - Beat a similar-sized monolithic model on all tasks & a **4x larger model** on most using generalist + 64 expert setup.
 - Engineered **JAX-based** training, evaluation & bulk inference pipelines for composed models in **Gemini codebase**.
- **Indic LLM Evaluation Leaderboard** Jan 2024 - June 2024
Advisors: [Dr. Sunayana Sitaram](#), [Dr. Vivek Seshadri](#), [Dr. Swami Manohar](#) Microsoft Research India
 - Collaborated with [People+AI](#) and [Karya Inc.](#) to build a fair and transparent evaluation platform for LLMs. 🏆
 - Curated **culturally-nuanced** prompt through native-speakers in **10 Indian languages** and evaluated 30+ LLMs.
 - Built leaderboards from **90k human & 30k LLM-based** evaluations via pairwise comparisons & direct assessment.
 - Studied Human-LLM agreement and potential biases in evaluations across languages and models. [EMNLP'24]
- **Multicultural and Multilingual Toxicity Dataset Curation** Dec 2023 - April 2024
Advisor: [Dr. Adrian de Wynter](#) Microsoft Redmond
 - Released culturally-nuanced toxic multilingual benchmark in **38 languages** to facilitate safer deployment of LLMs.
 - Evaluated **10 S/LLMs** for toxicity detection across **8 harm categories** and different language groups. [AAAI'25]
 - Observed LLMs fail to detect subtle biases (microaggressions) & performance decreases with language's availability.
- **Multilingual Finetuning and Benchmarking** Oct 2023 - Feb 2024
Advisors: [Dr. Sunayana Sitaram](#), [Dr. Kailka Bali](#) Microsoft Research India
 - Studied **QLoRA** finetuning technique across different rank-quantisation configurations, base models & datasets.
 - Finetuned **90+ models** and analysed cross-lingual transfer, catastrophic forgetting and effects of **PEFT**. [ACL'24]
 - Benchmarked multilingual capabilities & contamination of 20 LLM across 23 dataset & 83 languages. [NAACL'24]
- **Shiksha Copilot Web App** June 2023 - Nov 2023
Advisors: [Dr. Akshay Nambi](#), [Dr. Tanuja Ganu](#) Microsoft Research India
 - Engineered **Shiksha Copilot** web app to aid teachers in creating engaging content for students using Gen-AI.
 - Built an intricate and user-friendly frontend via streamlit, structured telemetry and analysed collected data.
 - Collaborated with [Shikshana Foundation](#) and deployed Copilot in **10+ schools** across Karnataka to teachers.

📰 **Media Coverage:** [Satya Nadella talk](#), [Times of India](#), [The Indian Express](#), [Analytics India Magazine](#), [Microsoft](#)
- **Holistic User-Embeddings via GNNs** May 2022 - July 2022
Mentors: [Pranjal Singh](#), [Mohit Kumar](#) Udaan
 - Developed framework to generate user representation from buyer-seller interaction graph for better segmentation.
 - Built multi-relational & multi-entity graph and modeled **Hetero-Graph AutoEncoder** with novel a loss function.
 - Decreased Udaan dispute loss by **2.45%** by incorporating embeddings in existing fraud model – **full-time offer**.

COURSEWORK

*=ONGOING

- **CMU:** Advanced Introduction to Machine Learning* (10-715), Probability & Mathematical Statistics* (36-705), Deep Learning Systems: Algorithms and Implementation* (10-714)
- **IIT-D:** Data Structures and Algorithms (COL106), Digital Electronics (ELL201), Machine Learning (COL341), Social Computing (COL865), Computer Networks (COL334), Analysis and Design of Algorithms (COL351), Linear Algebra and Differential Equations (MTL101), Calculus (MTL100), Probability and Stochastic Processes (MTL106), Signals and Systems (ELL205)
- **Online:** Stanford CS229-Machine Learning, DeepLearning.AI – Deep Learning Specialization (5 courses).

VOLUNTEERING

- Mentored six freshmen for their academic wellbeing through **BSW Student Mentorship Program** at IIT Delhi.
- Prepared Care-Packages comprising poems & articles under **NSS Mental Health Project** during the pandemic.
- Worked in the social media promotion department to spread awareness for the **non-profit 'No Help Too Big'**.

CO-CURRICULAR ACTIVITIES

- Organised & participated in India's **First Model Builder Summit** at Bangalore Convention Centre with People+AI.
- Gave an invited **Talk** at **National Research Council Canada** (NRCC) on "Inclusivity in Evaluations" in Sept '24.
- Presented paper posters at ACL 2024 in Bangkok & EMNLP 2024 in Miami and regular **Reviewer for ACL-ARR**.
- Attended **Winter School on Deep Learning** organised by Indian Statistical Institute (ISI), Kolkata in Spring 2023.
- Completed **A1 level in Spanish** as non-credited course of Humanities & Social Sciences Department, IIT Delhi.

OTHER INDUSTRIAL EXPERIENCE

Torch Investment Management

Machine Learning Engineer Intern with [Amit Sharma](#)

Project: Stock Price Modelling

Sept 2022 - Dec 2022

Remote / Noida, India

Udaan

Data Scientist Intern with [Pranjal Singh](#)

Project: Holistic User-Embeddings via GNNs




May 2022 - July 2022

Remote / Bengaluru, India

TECHNICAL SKILLS

- **Languages:** Python, Java, SQL, Bash, C, C++, ~~ML~~ \LaTeX , JAX • **Tools:** Git, Docker, ROOT, RAG, Autodesk Inventor
- **Libraries:** PyTorch, TensorFlow, Numpy, Pandas, Matplotlib, Scikit-Learn, Streamlit, Langchain, Hugging Face

REFERENCES

- Dr. Partha Talukdar Senior Staff Research Scientist, Google DeepMind 
- Dr. Sunayana Sitaram Principal Researcher, Microsoft Research 
- Dr. Adrian de Wynter Principal Applied Scientist, Microsoft Redmond 
- Prof. Abhishek Iyer Assistant Professor, IIT Delhi 