

AYUSHMAN GUPTA

✉ hey@ayushman.me  [LinkedIn](#)  [GitHub](#) ayushman.me

Education

University of Toronto

Sep. 2025 – May 2029 (expected)

Honours Bachelor of Science in Computer Science, Mathematics, & Philosophy

Toronto, Ontario

- *Relevant Coursework:* Analysis I, Abstract Linear Algebra, Probability & Statistics I, Data Structures & Analysis
- *Awards:* Lester B. Pearson International Scholarship 2025

Research Experience

IISER Kolkata

Summer 2026 – Present

Research Collaborator

Kolkata, India

- Mechanistically investigating memorization and training data attribution in LLMs.
- Exploring applications to law, healthcare, and other fields.

Independent Research

Mar. 2024 – Nov. 2024

Independent Researcher

Kolkata, India

- Led research projects and created benchmarks for evaluating fact-checking abilities of 5 LLMs across 6 diverse languages: Spanish, Italian, Portuguese, Turkish, Tamil, and Arabic.
- Employed zero-shot, chain-of-thought, and cross lingual prompting to study how fact-checking performance varies with prompting technique.
- Performed exhaustive statistical analyses such as correlation analyses and a two-way ANOVA to investigate the effects of the prompting technique and the LLM on the observed accuracy scores.
- Papers accepted at NLP4PI at EMNLP, Safe GenAI at NeurIPS, AbjadNLP at COLING

IISER Kolkata

Aug. 2023 - Aug. 2025

Research Collaborator

Kolkata, India

- Led two research projects on the generation and evaluation of code-mixed sentences across 5 low-resource language pairs, mentored by Dr. Kripabandhu Ghosh.
- Created SOTA gold-standard-agnostic metric for automatic evaluation of code-mixed sentences, eliminating the need for human annotators in the code-mixed evaluation process
- Built novel code-mixed generation algorithm informed by real-world code-mixing patterns from social media data; created novel prompting method utilizing well-defined linguistic rules for accurate English-to-CM translation.
- Benchmarked the English-to-CM and CM-to-English translation abilities of 3 LLMs across 5 language-pairs using 7 different kinds of few-shot prompts; facilitated creation of and released benchmark dataset.

Dev Experience

Flagpost AI

Summer 2026 – Present

Software Engineer Intern

Toronto, Ontario

- Opened and shipped 30+ PRs implementing and improving features and fixing bugs across platform's AI workflows and pipelines, infrastructure, backend, and frontend.

Claisen

Sep. 2025 – Oct. 2025

Software Engineer

Boston, Massachusetts

- Built AI medical consultation tool with phone calling functionality with Twilio and secure storage of medical records.

Projects

Paper Walkthroughs | PyTorch

Summer 2026 – Present

- Explainers and reimplementations of interesting papers on my website, e.g. [Deep Differentiable Logic Gate Networks](#) and [Mechanistic Training Data Attribution](#).

Melfi AI | Next.js, Node, Python, PyTorch

March 2026

- Built agentic AI therapist in 36 hours; placed top 10 at GenAI Genesis 2026 (Canada's largest AI hackathon).
- Fine-tuned Qwen3-8B using Reinforcement Learning From Human Feedback via Direct Preference Optimization
- Implemented phone calling functionality using Vapi.ai and Twilio API, and memory from previous sessions using RAG.
- Used self-hosted model server (served model via vllm, exposed over cloudflare tunnel).
- Built 3D manifold/mesh representing emotional state across conversations using Three.js

Technical Skills

Research: Custom model architectures, mechanistic interpretability, physics-informed ML, generative modelling, LLM fine-tuning and inference, RLHF.

Libraries: PyTorch, JAX, HuggingFace Transformers, TransformerLens, NumPy, Pandas

Languages: Python (5 years), Java (8 years), CPP (2 years), Typescript and Javascript (6 years)

Technologies/Frameworks: Next JS, Supabase, Postgres, Node, Bun, Prisma, Linux