Jibran Iqbal Shah

+1 647-676-3059 | jibraniqbal.shah@mail.utoronto.ca | Toronto, Ontario | jibby1729.github.io

Education

University of Toronto — Expected May 2027
Master of Engineering, Computer Engineering (Co-op) — *Toronto, ON*

University of Toronto — July 2025
*Honours Bachelor of Science, Mathematics Specialist (**GPA: 3.97/4.00**)* — *Toronto, ON*

Awards and Honors

Putnam Top 500 (2023) – Top 50 in Canada on the Putnam Competition, Top 500 worldwide.

Samuel Beatty In-Course Scholarship (2025) – Awarded to the top 3 UofT Mathematics students yearly.

Work Experience

Block (Square) — May 2026 – Dec 2026
Machine Learning Intern — *Toronto, ON*

- Proprietary AI research on speech-to-speech models.

Robot Vision and Learning Lab — May 2025 – Aug 2025
Machine Learning Research Intern (UTEA) — arxiv.org/abs/2602.16035 (Publication) — *Toronto, ON*

- Trained a Deep Learning based trajectory forecaster (HAICU) using PyTorch on 200K+ trajectories, reducing prediction errors by 20-40% overall, for an autonomous driving project with the **Toyota Research Institute**.
- Designed and implemented a novel CasADi-based Stochastic Model Predictive Control (SMPC) planner, achieving 98.5% collision avoidance guarantee across 2000+ simulation scenarios.
- Optimized gradient-based IPOPT solvers using mathematical optimization techniques, reducing average trajectory planning latency from 2s to 150ms (92% improvement), enabling real-time planning at 5Hz for dynamic scenarios.

Fields Institute — May 2024 – Aug 2024
Fields Summer Research Intern (FUSRP) — arxiv.org/abs/2511.22905 (Publication) — *Toronto, ON*

- Used Martingale Theory to prove novel Central Limit Theorems for Random Multiplicative functions with applications to the discrete logarithm problem in Cryptography.

Projects

TargetTTS — Apr 2026
VoiceFilter Target-Speaker Speech-to-Text — Github Repository

- Created a target-speaker extraction based ASR model that reduced WER from 92.2% to 41.9% on noisy multi-speaker full-overlap mixtures using ECAPA-TDNN embeddings and Whisper ASR.
- Reduced non-target leakage by 60% and CER by 61% on the same benchmark by training a speaker-conditioned mask extractor on synthetic LibriSpeech mixtures.

RL for GPT2 — Nov 2025
Implemented "RL Finetunes Small Subnetworks in LLMs" from scratch — Github Repository

- Post-trained GPT2-medium (400M params) using Reinforcement Learning to generate positive sentiment completions using PPO and GRPO, with a sentiment classifier as the reward model.
- Re-implemented the paper "Reinforcement Learning Finetunes Small Subnetworks in Large Language Models" (arXiv:2505.11711), analyzing sparse subnetworks formed, and presented it in an RL seminar to 20+ people.

Skills and Frameworks

Languages: Python, C++, JavaScript, TypeScript, HTML/CSS, SQL, Bash

Frameworks/Libraries: PyTorch, transformers, trl, CasADi, pandas, FastAPI, wandb

Tools/Technologies: AWS/S3, Git, Docker, PostgreSQL, GitHub Actions, DigitalOcean, Linux/Unix, Slurm