Li Jian Zhang

Toronto, ON | *647-858-6883* | *lijian.zhang@mail.utoronto.ca* | *linkedin.com/in/lijian-zhang* | *github.com/lzhang4672*

Education

University of Toronto (St. George) **BSc in Computer Science (Comp Sci Specialist & Stats Major)** | cGPA : 3.76/4.0 *2022-2026*

Research Experience

Research Assistant | *University of Toronto — Prof. Xujie Si* *June 2025 – Present*

- Interpretable ML and neurosymbolic reasoning (rule extraction, logical constraints, verifiable models).
- Implemented scalable post-hoc rule-extraction for transformer models; improved faithfulness and runtime.
- Built neuro-symbolic generation and logical-decoding methods.
- Extended rule-based interpretability to LLMs and GNN models.
- Contributed heavily in major experiments / pipelines, optimizations, manuscripts, and weekly experiment reports.

Student ML Researcher Internship | *Flybits* *February 2025 – December 2025*

- Co-developed a privacy-preserving, perspective-aware neuro-symbolic dataset and pipeline.
- Led image/text generation modules and rule-driven alignment; contributed methods and evaluation to manuscript.

AI Lab Automation Developer | *Acceleration Consortium — SDL6: Human Organ Mimicry* *April 2025 – Present*

- Built Bayesian optimization and experiment-planning loops for autonomous wet-lab workflows.
- Mapped SOPs to executable steps and integrated robotic control with clear, testable APIs.

Research Assistant | *University of Toronto — Prof. Feng Ji* *Sept 2024 – March 2025*

- Developed approach to address missing labels in psychology by combining clustering techniques and synthetic data generation using large language models.
- Applied advanced evaluation methods and leveraged proxy models to synthesize training examples, improving downstream classification on psychological datasets such as GoEmotions and sarcasm detection tasks.

Research Assistant | *University of Toronto — Prof. Irene Yi* *June 2024 - November 2024*

- Built a rule-guided LLM pipeline for SEC filings: custom parser, section prioritization, and grounded quoting.
- Owned data, prompts, and evaluation; shipped a reproducible codebase with clear docs.
- Reduced manual data annotations drastically while maintaining high accuracy and F1 scores.

Manuscripts / Preprints

NEUROLOGIC: Post-hoc FOL rule extraction for transformer encoders Under revision | Preprint | 2025

VisionLogic: Rule-based explanations for vision models Under review | Anonymized preprint | 2025.

NSGGM: Neuro-symbolic molecule generation with SMT constraints Under review | Anonymized preprint | 2025

Presentations

Rethinking Deep Learning Autodiff for Rust & C | *Poster Link* *December 2024*

- Exploratory work on the viability of *Enzyme AD*, an experimental LLVM-based automatic differentiation tool, for training deep neural networks, specifically the GPT-2 architecture.

Skills

Languages : Python, C/C++
ML Libraries : PyTorch, HuggingFace Transformers/Diffusers, NumPy, Pandas, PyTorch Geometric
Research Areas : Trustworthy and Interpretable ML; Neurosymbolic AI and Knowledge Representation; Representation Learning; Graph ML; Optimization and Probabilistic Modeling, AI Safety, AI Alignment.
Coursework : CSC413 (Neural Nets and Deep Learning) (A+), CSC311 (Intro Machine Learning) (A+), CSC420 (Computer Vision) (A), CSC486 (KR&R) (In Progress), CSC384 (Intro AI) (A), STA302 (Data Analysis) (A), MAT237 (Advanced Calculus) (A-)
Systems : Linux, CUDA, multi-GPU training, SLURM (Alliance/Compute Canada), Docker, Git, LaTeX

Awards

Entrance Scholarship | *Trinity College* *September 2022*

Dean's List Scholar *2022, 2023, 2024, 2025*