

# John Zhang

☎ 519-616-4119 ✉ [johnzhanguoft@gmail.com](mailto:johnzhanguoft@gmail.com) [in linkedin.com/in/john-z-zhang](https://www.linkedin.com/in/john-z-zhang) [github.com/john-zhang-uoft](https://github.com/john-zhang-uoft)

## Education

---

### University of Toronto

Expected May 2025

*Bachelor of Applied Science in Engineering Science - Machine Intelligence*

Toronto, Ontario

- **Coursework:** Advanced Algorithms, Deep Learning, Artificial Intelligence, Operating Systems, Software Engineering, Performance Optimization, Data Structures, Decision Support Systems

## Experience

---

### Intel (Altera)

May 2023 – Aug 2024

*Software/Machine Learning Engineer*

San Jose, California & Toronto, Ontario

- Led R&D into a novel GenAI feature that improves hardware design flows by leveraging LLM design generation.
- Built a data processing and augmentation pipeline to create a hardware system dataset with millions of tokens.
- Outperformed SOTA prompt-engineered foundation models by 300% on test set by leveraging performance efficient fine tuning (QLoRA) and RAG to fine-tune an open source LLM with Transformers library on Azure ML Studio.
- Designed a custom language for LLM outputs, increasing total inference speed by 100% from baseline.
- Developed and integrated C++ and Tcl code to support the migration to a modern database for Quartus tools, using Object Oriented Programming patterns for improved maintainability.
- Architected a testing framework for a user-facing debug tool and used it to validate the tool's reliability on thousands of inputs from customer designs.
- Continuously resolved customer requests and bugs by implementing software features and fixes, including creating a Quartus Python module for merging customer debug files.

### The Matter Lab

Sept 2021 – Apr 2023

*Part-Time Undergraduate Machine Learning Researcher*

Toronto, Ontario

- Developed a pipeline to process 40,000 datapoints with xtb, Open Babel, and PyTorch, building a molecule dataset.
- Trained and tuned graph convolutional neural networks to predict molecule solubility with Pytorch Geometric while experimenting with various architectures and physical chemistry features.
- Demonstrated that a electronic surface mesh approach is valid on solubility task, beating some architectures such as XGBoost and GC by 10% in RMSE loss on ESOL solubility task.

### Rocscience Inc.

May 2021 – Aug 2021 & May 2022 – Aug 2022

*Software Engineering Intern*

Toronto, Ontario

- Developed a program in C++ that fixes mesh holes by transforming meshes into voxel grids with OpenVDB.
- Improved triangle element feature query speed in C# by more than 10x by redesigning algorithms to leverage Bounding Volume Hierarchies.
- Collaborated on and presented a cross-platform model viewer in Typescript using Vue.js and Electron.
- Developed an augmented reality interactive 3D geotechnical model viewer app using Unreal Engine 4 and deployed it on the Microsoft HoloLens.

## Projects

---

### Winner of National Undergraduate Big Data Challenge | *STEM Fellowship*

- Co-wrote [scientific research paper](#) analyzing communities and the spread of COVID-19 misinformation on Twitter.
- Processed and analyzed 300,000 Tweets from Twitter's graph analytics API by developing a Python data pipeline and applying PCA, t-SNE and regression, winning the Outstanding Science Communication Award.

### Finalist of Hack the North | *Hack the North*

- Created a [Chrome extension](#) that automatically takes lecture notes and allows users to select and copy text seen in videos with Google Cloud Vision's API and Flask, and was selected as a finalist in a hackathon of over 400 teams.

## Technical Skills

---

**Languages:** C++, Python, C, C#, Typescript, Javascript, SQL, Bash, Tcl, Verilog, MATLAB

**Technologies:** Transformers, Pytorch, JAX, React, Vue.js, Node.js, Flask, HTML, CSS

**Tools:** Git, Perforce, Azure ML, Linux, Docker, Jira