Ted Zhang

Cell Phone: (647) 643-3466 E-mail: th3zhang@uwaterloo.ca LinkedIn: https://www.linkedin.com/in/ted-h-zhang/

Education

BCS, Computer Science, *University of Waterloo* Apr 2026

Cumulative GPA: 93.5/100

Waterloo, Ontario

- Spring 2022 **First in Class Engineering Scholarship**, Fall 2021, Spring 2022, Winter 2023 **Dean's Honours List**
- Relevant Courses: Algorithms, Data Structures, OOP, Compilers, Linear Algebra, Probability, Statistics, Combinatorics

Work Experience

Software Engineering Intern, *Kinaxis Inc.* May 2024 – Aug 2024

Toronto, Ontario

- Built a support chatbot for advanced log querying, log analysis, and root cause analysis on a large-scale distributed system.
- Engineered 2 **REST API**s using **Express.js**, **OpenAI embeddings**, **GPT-4**, and a **PostgreSQL** database hosted on **Azure**, with a **React.js** and **CSS** frontend. Crafted an **ETL** pipeline that ingests **350,000** logs a day using **Apache Airflow** and **Datadog API**.
- Deployed **3** applications with **Azure Kubernetes Service** using **Docker** containers, **Kubernetes** manifests, and **Helm** charts.
- Created **2 FastAPI** endpoints and corresponding unit tests to fetch table metadata using **Databricks API** and **SQL Warehouse**.

Machine Learning Developer Intern, *Kinaxis Inc.* May 2023 – Aug 2023

Toronto, Ontario

- Enhanced runtime by **45%**, reduced code size by **40%**, and improved readability by refactoring data utilities using **PySpark**.
- Designed and implemented a **Python** module that detects erroneous, anomalous, and missing **time series** data before it impacts forecasting accuracy using **sklearn**, **Ruptures, Kalman Filters**, and an **Augmented Dickey-Fuller** test.
- Analyzed and aggregated data from a large-scale **Apache Hive** data warehouse using **PySpark**, **HQL**, and **Azure Databricks**, reducing processing time by **53%**.

Software Engineering Intern, *BlackBerry Limited* Sept 2022 – Dec 2022

Waterloo, Ontario

- Developed an unsupervised **NLP** model with an **0.87 f1-score** and **92% accuracy** for log anomaly detection using **sklearn**, **hashing vectorizers**, **scalers**, **LSTM** autoencoders, Google's **BERT transformer**, and **isolation forest**.
- Tested and integrated the new model into the existing codebase, pushing over **2,000 lines** of code into **production**.
- Optimized the machine learning data pipeline, reducing redundant API calls by **55%** and memory usage by **52%** with **multiprocessing**, **memory tracing**, and **profiling**.
- Implemented a **CI/CD** pipeline using **Git**, **GitLab CI/CD**, **GitLab Runner**, **Docker**, and **shell scripts**.

Machine Learning Developer Intern, *Advanced Micro Devices Inc. (AMD)* Jan 2022 – Apr 2022

Markham, Ontario

- Spearheaded development and training of a **computer vision** model trained on **700,000 images** with **94% validation accuracy** using **TensorFlow** and **Keras**, transfer learning using **InceptionV3**, and image preprocessing using **OpenCV.**
- Containerized an inference optimization library along with various other pretrained models using **Docker** and **shell scripts**.

Relevant Projects

ML^2: Machine Learning Money Lines - ml-squared.ca/

- Aggregated over **500,000** lines of NBA player data using **pandas** and developed a tree-based **regression** model using **XGBoost**, boasting a mean absolute error of **4.07** with **58%** of predictions within 3 points of actual performance.
- Developed a **REST API** using **Django**, hosted on **AWS EC2** using **NGINX** and **Gunicorn**, with a **React.js** and **CSS** frontend.

Skills

- Languages: Python, C++, JavaScript, HTML, CSS, SCSS, Java, SQL, HQL
- Frameworks: Express.js, Node.js, TensorFlow, PySpark, sklearn, Django, Docker, Kubernetes, XGBoost, React.js, FastAPI
- Cloud & Other: AWS EC2, Gitlab CI/CD, Git, Azure Container Registry, AKS, PostgreSQL, Hive, Airflow, Time Series Analysis