

Profile

I'm a Toronto-based software developer with experience bringing data intensive applications to production on the AWS and GCP clouds. My specialty is data engineering and MLOps.

Skills & Knowledge

- I'm an AWS certified Machine Learning Specialist and I have in-depth experience running ML workloads and data-intensive distributed systems on AWS and GCP (App Engine, Cloud Run, Vertex AI).
- I have deep experience designing clean, well-tested, and performant distributed systems using Python, Bash, and Node.js.
- I have significant experience in applied natural language processing using Gensim, NLTK, and spaCy.

Employment History

CBC/Radio-Canada, Toronto, ON (Remote)

Senior Machine Learning Developer

Oct 2022 – Present

Machine Learning Developer

April 2021 – Oct 2022

- Contributed to REST and GraphQL APIs spanning multiple teams. I led the migration of a major API version, reducing total lines of code by 25% in a large codebase serving ~80 requests/second.
- I designed a Cloud Function architecture for end-to-end model lifecycle and scheduling. This system currently serves 3 audience-facing ML models with a total load of over 1 million requests/day.
- Implemented a minhash model for related videos that increased relative click-thru by 17% compared to an existing session-based model. This system now serves ~300k recommendations per day.
- Converted an existing session-based recommendation model from Tensorflow 1 to 2, enabling my teammates to extend this model in numerous A/B tests.
- Mentored and sponsored teammates through 1-1s, code reviews, facilitating cross-team initiatives, providing feedback on design documents, and setting technical direction.

CaseWare International, Toronto, ON

Server Developer – Machine Learning

May 2016 – March 2021

- Developed a named entity recognition model that doubled the successful extraction rate for lease parties compared to a baseline model. I deployed this model, with unit tests, using AWS SageMaker.
- Independently retrieved, cleaned, and structured a large dataset of XML financial statements. I containerized this ETL pipeline and deployed it to in-house servers.
- Trained and evaluated models for predicting future amendments to financial statements. Achieved a 10% reduction in model error by implementing multiple feature selection algorithms.

Academic Background

Master of Science, Applied Computing

Sept 2015 – Jan 2017

University of Toronto, Toronto, ON

Honours Bachelor of Arts (with high distinction), Computer Science

Sept 2011 – April 2015

Wilfrid Laurier University, Waterloo, ON

- Received an Undergraduate Student Research Award to work on a team conducting novel graph theory research. Our results were published in Discrete Applied Mathematics (2016).

Honours and Awards

MScAC Program Award – 2021 Alumni of the Year

Nov 2021

- Presented to alumni who have made significant contributions to the development of the program.

MITACS Accelerate Fellowship

May 2016 – Dec 2016

Addictive Mobility Scholarship in Applied Computing

Dec 2015