

# Bokai (Jason) Lai

(647) 328-9168 • [bokai.lai@rotman.utoronto.ca](mailto:bokai.lai@rotman.utoronto.ca) • [linkedin.com/in/jason-lai](https://www.linkedin.com/in/jason-lai) • Portfolio: <https://jasonlbk.com/>

---

Business-minded Data Scientist with experience in **machine learning (supervised/unsupervised)**, **statistical analysis**, **analytics engineering**, and **applied AI**. Skilled in Python, SQL, PySpark, and BI tools, with a strong ability to turn complex data into scalable models, actionable insights, and business-facing analytics products

## TECHNICAL SKILLS

---

**Programming Language:** Python, SQL (Expert; Prev Head TA for Database Course), R, JavaScript

**Machine Learning:** Scikit-learn, PyTorch, TensorFlow, XGBoost, LightGBM, Time series forecasting, A/B testing, Causal Inference

**Data / Engineering / BI:** PySpark, Pandas, NumPy, ETL pipelines, AWS, Docker, FastAPI, MLflow, Git, Tableau, MS Excel

**Applied AI / LLM:** Transformers, RAG, LangChain, LangGraph, FAISS, BM25, Vector Database, LLM-as-Judge

## EDUCATION

---

**Rotman School of Management, University of Toronto**, Toronto, Ontario

**Master of Management Analytics Candidate**, 2025 – Expected July 2026

- 1<sup>st</sup> Place - Calian Case Competition (Data Analysis); 2<sup>nd</sup> Place - Scotiabank Case Competition (Credit Risk)

**University of British Columbia**, Vancouver, BC

**Bachelor of Science in Computer Science and Statistics**, 2020 – 2024

- Dean's List (2021–2023); VP of Statistics Club: led finance and sponsorship data governance for 10,000+ members

## PROFESSIONAL EXPERIENCE

---

**Scotiabank (AML)**, Toronto, Ontario

**Data Scientist Intern**, January 2026 – Present

- Built **Isolation Forest**, **Autoencoder**, and semi-supervised models in **Python (PyTorch framework)** across **6+** transaction typologies and million-scale samples. Improved **alert precision** by **18%** and reduced **false positive rate** by **23%** through threshold recalibration and feature reweighting, significantly lowering manual review costs
- **Engineered 50+** behavioral, time-series, and transaction-level **features**, including rolling-window statistics, robust z-score normalization, transaction velocity shifts, counterparty behavior, and customer activity patterns to support scalable model training and analytical monitoring
- Designed **model validation** and **monitoring** workflows using time-aware cross-validation, leakage prevention checks, feature stability analysis, and business-aligned performance metrics, improving composite model performance by **12%**
- Prototyped a governed **multi-agent** AML investigation assistant using **LangGraph**, hybrid **RAG**, audit logging, **compliance guardrails**, and human-in-the-loop review to support case triage, evidence assembly, and investigator decision-making

**Siemens (Digital Industries)**, Beijing, China

**Data Scientist Intern**, October 2024 – May 2025

- Applied **NLP** preprocessing, including tokenization, stop-word removal, and text normalization, on customer text data; used **SMOTE-Tomek** sampling to address class **imbalance** and improve model generalization for downstream classification tasks
- Implemented and tuned a hybrid model combining **CNN**, **Bi-LSTM**, and **LightGBM** in TensorFlow for client intent triage, tuned hyperparameter via K-Fold CV; **Docker**-containerized and deployed on **AWS** for reproducible low latency inference, resulting in accuracy improving from **0.76** to **0.91** and **15%** F1 Score gain, and enabling robust business decision support
- Dockerized model inference and **ETL** workflows on **AWS**, enabling reproducible deployment, lower-latency inference, and more reliable analytics delivery across business teams
- Developed a production **RAG** chatbot with persistent context and a localized DeepSeek 32B model, improving document comprehension by **37%** and increasing chatbot accuracy and user satisfaction by **20%**, supported by feedback loops and monitoring to sustain quality

**Pou Sheng International**, Dalian, Liaoning, China

**Data Analyst Intern**, July 2024 – October 2024

*Leading sports apparel retailer in Greater China, listed on HKEX (3813.HK)*

- Analyzed performance across 8 brands' retail stores, identifying profitability drivers and underperformance root causes; recommended actions that **reduced operating costs** by **12%** and improved mid-year sales by **9.8%**
- Built profitability and customer analytics frameworks using **ARPU**, repurchase rate, and sales per sq ft, leveraging data segmentation and **KPI** tracking to support recurring business performance reviews
- Developed stakeholder-facing **visual analyses and reporting** views to communicate sales trends, customer behavior, and operational gaps to business teams

## TECHNICAL PROJECTS

---

**Coinbase Support Agent | AI Agent | RAG System | [GitHub](#)**

- Designed and deployed an agentic AI customer support system using **LangGraph**, **hybrid RAG**, **tool use**, and safety **guardrails** to resolve domain-specific queries and multi-turn support workflows
- Built a retrieval layer over **60+** Coinbase Help Center documents with semantic chunking, **FAISS vector search**, **BM25 keyword search**, and citation-aware response generation to improve grounding and reduce unsupported answers
- Developed scenario-based evaluation using **LLM-as-Judge** metrics, including **faithfulness** and **answer relevancy**, achieving **100%** task success and **97%+** grounded response quality