

Zhongmang (Marc) Cheng

Data Engineer · Applied ML · HK TTPS Visa Holder

Canada: +1 438-886-8425 | Mainland China: +86 156-5240-1103 | marccheng@outlook.com
chengmarc.com | linkedin.com/in/chengmarc | github.com/chengmarc

EDUCATION

University of Toronto

Sep. 2019 – Jun. 2023

Honours Bachelor of Science — Mathematics & Its Applications

Toronto, ON

- Specialist in Probability and Statistics; Major in Mathematics
- GPA: 3.6/4.0, High Distinction | Dean's List: 2021, 2022, 2023
- Relevant coursework: Real Analysis, Time Series Analysis, Multivariate Data Analysis, Nonlinear Optimization

PROFESSIONAL EXPERIENCE

Product Data Specialist

Jun. 2023 – May 2025

Rexel Canada

Mississauga, ON (On-Site)

- Led end-to-end implementation of a new Product Information Management (PIM) system, defining data models and core workflows in Informatica Product 360 to support enterprise-wide product operations.
- Acted as the primary data interface between SAP Hybris, ERP (IBM AS400), and external industry data sources, ensuring consistency across 200,000+ SKUs.
- Designed and executed UAT frameworks and validation rules, coordinating cross-functional testing with offshore engineering teams (India, France) to ensure production readiness.
- Designed and implemented automated ETL data quality pipelines and CI/CD scripts using AWS CLI and Bash, translating business requirements into daily validation rules with automated batch deployment and monitoring.
- Contributed to enterprise LLM initiative (RexelGPT), focusing on system performance optimization and real-world deployment constraints during beta testing.

Data Analyst (Co-op)

Jul. 2021 – Apr. 2022

Deloitte

Toronto, ON (Hybrid)

- Performed validation of enterprise AML models for a Tier-1 Canadian bank, assessing model assumptions, feature selection, and statistical robustness.
- Validated AML classification models (KNN, K-Means, XGBoost) against large-scale transactional datasets (1M–20M records, Oracle), verifying feature consistency and pipeline logic under the SR 11-7 model risk framework.
- Conducted model logic and data pipeline verification using Python, R, and SQL to identify inconsistencies between theoretical design and production behavior.
- Executed SQL queries and pipeline audits in SSMS against LoyaltyOne's loyalty program data, validating model inputs and output consistency across customer transaction segments.

FREELANCE EXPERIENCE

Front-End Developer

Jun. 2025 – Present

DoubleTrends

Toronto, ON (Remote)

- Designed and developed doubletrends.com from scratch using Tailwind CSS and Node.js; implemented component-level layout architecture to ensure full responsiveness across desktop and mobile viewports.
- Architected full production deployment on Cloudflare Pages with Stripe payment integration, enabling gated feature access and live user transactions.

PROJECTS

PaySim Fraud Detection Data Warehouse | *PySpark, Hive, Hadoop* | [GitHub](#)

Mar. 2026

- Designed and implemented an offline data warehouse for financial fraud detection on 6.3M PaySim transactions, covering the full pipeline: ingestion, transformation, and aggregation.
- Structured DWD layer with ORC-Snappy storage and dual partitioning on (`tx_type`, `tx_day`): `tx_type` as the top-level key ensures ORC predicate pushdown prunes at the directory root, avoiding cross-partition traversal for all business queries filtered by transaction type.
- Modeled ODS as an external table to decouple metadata lifecycle from HDFS data — dropping the table preserves ODS data; all downstream DWD and ADS tables are managed to enforce standard lifecycle control.

MAMBA State-Space Model | *PyTorch*, *RunPod* | [GitHub](#)

Feb. 2025 – Mar. 2025

- Implemented a MAMBA state-space model (selective SSM with dual-branch projection, depthwise convolution, SiLU gating) for sequence-to-sequence forecasting on multi-feature time series; architecture was subsequently adopted by Stoic AI — a live algorithmic trading platform — for forecasting infrastructure.
- Designed a detrending pipeline: logarithmic trend extraction via nonlinear least squares, residual forecasting, and inverse reconstruction.
- Benchmarked against five baselines (SegRNN, LSTM, Seq2Seq LSTM, Attention LSTM, Transformer) on cloud-based H100 instances via RunPod.

Large Language Model (GPT-2 Style) | *PyTorch*, *tiktoken* | [GitHub](#)

Dec. 2024

- Implemented a GPT-2-style transformer with 768-dim embeddings and 12-head self-attention.
- Trained 162M parameters on a custom corpus using AdamW optimizer; tokenization via BPE (*tiktoken*).

Image Classification | *PyTorch* | [GitHub](#)

Nov. 2024

- Implemented ConvNet on MNIST/EMNIST; ResNet-34 and ResNet-101 on CIFAR-10/100.
- Compared basic blocks (two 3×3 convolutions) vs. bottleneck blocks ($1\times 1 \rightarrow 3\times 3 \rightarrow 1\times 1$); analyzed interaction between identity shortcut connections and network depth across datasets of varying difficulty.

TECHNICAL SKILLS

Big Data & Data Warehouse: Hadoop, Hive, Spark, PySpark, HDFS, YARN**Databases:** Oracle, MySQL, IBM DB2 (AS400), SQL Server, BigQuery**Cloud & Infrastructure:** AWS, Azure, RunPod, Google Colab**Data Science & ML:** PyTorch, scikit-learn, pandas, numpy, scipy, matplotlib, seaborn, jupyter**Enterprise:** Informatica P360, SAP Hybris, Tableau, Power BI**Languages & Other:** Python, R, JavaScript, Node.js, SAS, Bash, PowerShell, Git, LaTeX, Markdown**Spoken:** English (bilingual), French (professional), Mandarin (native)**WORKING PAPERS**

Z. Cheng and N. Song. “When the Loop Closes: Architectural Limits of In-Context Isolation, Metacognitive Co-option, and the Two-Target Design Problem in Human-LLM Systems.” *arXiv preprint*, under review, Mar. 2026.

Zhongmang (Marc) Cheng. “On the Impossibility of Complete Interpretability in Neural Networks: A Measure-Theoretic Analysis.” *Zenodo preprint*, Feb. 2026. DOI: [10.5281/zenodo.18678067](https://doi.org/10.5281/zenodo.18678067)

AWARDS

University of Toronto Dean’s List: 2021, 2022, 2023**Concours Opti-math, L’Université Laval:** 5th Place Provincial (2015), 7th Place Provincial (2014)**University of Waterloo Gauss Mathematics Contest:** Distinction (2013)