Yohan Zytoon

Montréal, Québec, Canada
yohanze@icloud.com linkedin.com/in/yohanzytoon github.com/yohanzytoon yohanzytoon.com

Education

Université de Montréal **Expected Dec 2026**
B.Sc. in Computer Science *Montréal, Québec*

- **Relevant Coursework:** Machine Learning, Stochastic Processes, Probability & Statistical Inference, Optimization, Algorithm Design, Distributed Systems

Experience

Intact Financial Corporation **May 2026 – Current**
Data Scientist I — Applied AI / ML Engineering *Montréal, Québec*

- Built production **LLM agents** with **LangChain and LangGraph**, orchestrating **RAG, tool-calling, and stateful multi-step workflows** over enterprise claims data to support claims-adjustment decisions.
- Implemented a **MLflow-based evaluation framework** that automatically logs prompts, workflows, configurations, metrics, and deltas against baselines, turning agent experimentation into a reproducible, standardized workflow.
- Developed repeated-run **stability evaluation** with AI-generated diagnostics for inconsistent predictions; improved fully stable predictions from **71% to 82%** and increased overall model stability by **5%**.
- Performed deep **root-cause analysis** across POC failures, tracing prediction errors to specific prompt, retrieval, and workflow components and guiding changes that helped raise **macro-F1 from 0.32 to 0.78**.
- Built a versioned, **agent-accessible knowledge system (Wiki)** with automated documentation updates and changelog tracking, giving developers and AI agents a continuously maintained source of project context and system changes.

Desjardins Assurances générales (DAG) **Sep 2025 – Dec 2025**
Data Scientist Intern — Advanced Analytics *Montréal, Québec*

- Reduced batch processing time by **80%** by parallelizing ingestion workflows, optimizing intermediate data structures, and refactoring transformations into reusable components.
- Engineered **reproducible YAML-configured ML pipelines** and integrated feature-store outputs into Azure ML scoring artifacts, standardizing experimentation, retraining, and deployment workflows.
- Developed a **segmented correlation analysis** to identify relationships between candidate predictors and model targets across subpopulations, supporting more informed **feature selection** and model development.

Bombardier Recreational Products (BRP) **May 2025 – Aug 2025**
Business Intelligence Intern *Montréal, Québec*

- Designed **Snowflake** data models with automated reconciliation checks and built **Power BI** dashboards adopted across multiple business units.

Independent Projects

VeroLoop — Provider-Neutral AI Evaluation Platform | *Python, Pydantic, MLflow, AsyncIO* — ***GitHub*** **Current**

- Building an **open-source AI evaluation platform** that defines a structured task once and benchmarks models across **OpenAI, Anthropic, and Gemini** through a unified provider-neutral execution layer.
- Implemented **typed task contracts, provider adapters, schema compatibility, async concurrency, retries, and normalized errors** for reliable and reproducible cross-provider execution.
- Built evaluation tooling for **correctness, schema validity, repeated-run stability, latency, and cost**, with versioned reports and console, JSON, HTML, and **MLflow** outputs.

Budgeted Combinatorial Search | *PyTorch, GFlowNets, BoTorch, GPyTorch* — ***GitHub*** **2026**

- Investigated **sample-efficient exploration of a ∼8B-state combinatorial space** under a strict **1,000-query oracle budget**, implementing and comparing Gaussian Process active learning, direct GFlowNet training, and a hybrid GFlowNet–AL pipeline.
- Engineered the hybrid with **GP + UCB acquisition, diversity-aware selection, warm starts, delayed activation, and plausibility rewards**; across **5-seed experiments**, classical AL achieved **3.7× higher valid-generation rate** than direct GFlowNet under the constrained budget.

Technical Skills

Programming: Python, C/C++, SQL, Bash/Shell, R
ML / AI: PyTorch, scikit-learn, JAX, XGBoost, BoTorch, GPyTorch, HuggingFace Transformers, GFlowNets
LLMs / Agents: LangChain, LangGraph, Azure OpenAI, RAG, tool-calling, prompt engineering & evaluation
MLOps / Engineering: MLflow, Azure ML, Docker, Git/GitHub, YAML/Hydra, WandB, REST APIs
Data / Infrastructure: Databricks, Snowflake, AWS (S3, SageMaker, Bedrock), Azure, PostgreSQL/MySQL, Linux