Sherif Morsi

Montreal, QC, Canada | +1 (438) 596-6468 | shmorsi@outlook.com | linkedin.com/in/sherif-morsi-a17628295
Canadian Permanent Resident | github.com/shmorsi | shmorsi.github.io

Summary

Computer Science student building **AI for long-term investing**: **asset-liability** models for pension plans, **PyTorch** portfolio and forecasting models built from scratch, and **agentic AI** systems that act inside signed bounds, each tested out of sample.

Education

Concordia University Montreal, QC, Canada
Bachelor of Computer Science *Expected May 2028*

- **Availability:** Winter 2027 (Jan to Apr 2027) for 4, 8, 12, or 16 months, anywhere in Canada
- **Coursework: Probability**, Linear Algebra, Multivariable Calculus, Machine Learning, Artificial Intelligence, Data Structures and Algorithms, Object-Oriented Programming
- **Activities:** Member, QUARCC (Quantitative Research and Competitions Club)

Experience

AI & Machine Learning Intern Aug 2026 - Sep 2026
Nokia *Cairo, Egypt*

- Trained and evaluated Python and **PyTorch** models on wireless network telemetry to forecast 5G **energy demand**
- Developed **data and ML pipelines** for 5G energy telemetry, to support automated network management and optimization
- Trained in **agentic AI** for intent-based 5G network slicing (Amazon Bedrock); presented findings to cross-functional data science and engineering teams and led the project proposal and requirements review

Credit Analyst Job Shadow Aug 2025 - Sep 2025
Qatar National Bank (QNB) *Cairo, Egypt*

- Built and documented **Python** utilities and **SQL** queries that preprocess and validate SME financial records for the bank's credit-review process, improving data quality and reliability

Projects

Deep Portfolio Constructor (AI for Long-Term Investing) | *Python, PyTorch* | **GitHub** | **Project Page** 2026

- **Wealth study.** A long-only portfolio for **long-term wealth**: a **transformer** allocates across 30 US industries and learns how much to trade, trained in PyTorch on the **growth of wealth after costs** through a **differentiable simulator**
- **Walk-forward backtest** 1960–2026, **pre-registered**: \$1 grew to **\$1,368** vs \$929 for the market (+0.58% a year, t = 1.40, not significant); 5 classical baselines incl. risk parity and a Kelly optimizer, 3 ablations, 5 seeds, 152 tests
- **Pension Lab (asset-liability management).** A 90%-funded pension plan with 60 years of **inflation-indexed liabilities**; 60/40, **liability-driven (LDI)** and glide-path policies costed by contributions and **CVaR** over 10,000 **Monte Carlo** scenarios
- Trained a policy through a **differentiable simulator** (**pre-registered**): it halved the simulated cost (\$5.0 vs \$9.9 per \$100 promised) but only matched the glide path on 34 real 20-year periods; traced the gap to patterns the simulator repeats
- **World Shock Lab (macro shocks, code).** How Fed, Bank of Canada and OPEC surprises move **53** assets (10 bp Fed surprise: TSX −**1.10%**, 90% band −1.59 to −0.61); replicated **11 of 13** published findings; live forecasts before decisions
- **Pre-registered** test of machine learning on 168 held-out announcements: a neural net did not beat the linear model on the mean (gap −0.003, 90% CI includes zero); quantile boosting won only on 5% tail loss

Time-Series Foundation Model (Deep Learning) | *Python, PyTorch, GluonTS* | **GitHub** | **Project Page** 2026

- Pretrained a **10.1M-parameter** decoder-only transformer from scratch on **443,258** series across 77 datasets; on the frozen **GIFT-Eval** test split (97 configurations), zero-shot CRPS was **0.688**× AutoARIMA, 95% CI [0.624, 0.757]
- Wrote RMSNorm, **rotary embeddings**, SwiGLU, a KV cache and a non-crossing quantile head without a model library; **mixed-precision** training on one T4 GPU
- A **data-leakage audit** found benchmark test data in a public corpus and removed **801 of 447,229** series; 189 tests in CI

Aurvane (Agentic AI Security) | *Python, Ed25519, MCP, JavaScript* | **GitHub** | **Live Demo** 2026

- Trust layer between LLM agents and their tools: **Ed25519-signed mandates**, four checks per action, Allow/Deny approvals, narrower-only sub-agent mandates, hash-chained log; Python SDK and **MCP server**; 74 tests
- On the public **AgentDojo** benchmark (629 **prompt-injection** cases, worst-case agent, protocol fixed first), attacks that landed fell from **97.5% to 13.0%**; on my own 60 paired tasks, from 30/30 to 7/30 with 0 false alarms

Technical Skills

Machine Learning: **PyTorch**, **Deep Learning**, **Transformers**, Time-Series Forecasting, **scikit-learn**, TensorFlow, QLoRA
Investing: **Asset-Liability Management**, **LDI**, Fixed Income, **Portfolio Construction**, CVaR, Monte Carlo, Econometrics
AI Agents: **Agentic AI**, **LLM Agents**, **Tool Calling**, **MCP**, Prompt-Injection Defense, **LLM Evals**, Claude Code, Cursor
Languages: **Python**, **SQL**, TypeScript, JavaScript, C++, Java, C, R, Bash
Data / Tools: **Pandas**, **NumPy**, **FastAPI**, DuckDB, **Git**, **GitHub Actions (CI/CD)**, **pytest**, Playwright, Docker, Streamlit

Certifications

Anthropic: Building with the Claude API | **IBM:** Machine Learning with Python