ANDY AU

GitHub ◇ LinkedIn ◇ andyau2718@gmail.com ◇ +1 (617) 256 1336

Summary

Applied math student specializing in ML-driven trading strategies. Developed PPO-based portfolio allocator achieving Sharpe 2.00 vs 1.75 rolling Markowitz. Adapted PatchTST for volatility forecasting with 35% MSE improvement over GARCH/Kalman baselines. Pursuing research in continuous-time portfolio optimization under uncertainty.

Education

Boston University College of Arts and Sciences — **Boston, MA**

BA Applied Mathematics & Minor in Physics — 09/2023 - 05/2026

Notable Coursework:

- **Graduate**: Advanced Stochastic Processes (MA 783), Probability Theory (MA 779), Real Analysis (MA 711).
- Undergraduate: Methods of Quant Finance, Time-Series Forecasting, Financial Derivatives.

Research

RL in Bayesian Portfolio Control — Boston University

Reinforcement Learning for Portfolio Optimization under Drift Uncertainty — 05/2025 – Present

Advisor: Prof. Mark Kon

- Extending Wang & Zhou (2019) entropy-regularized MV framework to include Bayesian parameter learning.
- Derived belief-augmented HJB with state (X_t, m_t, v_t) under quasi-linear-Gaussian dynamics.
- Characterized optimal policy structure: Gaussian with mean affine in (X_t, m_t), variance decreasing with precision.
- Establishing convergence to classical MV as exploration weight and posterior variance vanish.

Relevant Projects

DRL Portfolio Optimization — PPO — Boston University

Portfolio Optimization with Reinforcement Learning and Markowitz Baselines — 04/2025 - 08/2025

- PPO allocator for 10 ETFs (SPY/QQQ/IWM/EFA/EEM/VNQ/TLT/IEF/GLD/USO) with 274-dim features.
- Constrained portfolio concentration via HHI bands to avoid both uniform and over-concentrated allocations.
- State-dependent exploration with annealed KL divergence; monthly refit on 90-day windows for regime adaptation.
- Test 2025H1 OOS: Sharpe 2.00 (24.6% excess, 12.3% vol) vs rolling Markowitz 1.75 (31.7%, 18.1%).
- Ranked top 0.01% among 1M Monte Carlo daily-rebalanced portfolios on same period.

Volatility Forecasting — PatchTST — HKUST

PatchTST Implementation for Realized-Volatility Forecasting on Cryptocurrency — 08/2024 - 06/2025

- 30-model ensemble achieving MSE 0.000106 (individual CI: 0.000108-0.000110) on BTC 5-day forward volatility.
- Outperformed GARCH/Kalman: 35% MSE reduction, QLIKE -6.33 vs -6.22; directional accuracy 0.60 vs 0.49.
- Trained on 2014-2022 BTC data (100-day lookback); tested on 2024-2025H1 out-of-sample period.

Experience

Machine Learning Research Intern — HKUST

Deep Learning Research — 07/2024 - 09/2024

Advisor: Prof. Yang Can & Dr. Zhang Fan

- Built attentive RNN and transformer architectures for sequential pattern recognition.
- Implemented and explored PatchTST and Temporal Fusion Transformers for time-series forecasting.
- Accelerated model convergence through learning rate scheduling, gradient clipping, and Optuna optimization.

Skills

- **Quantitative Methods:** Stochastic calculus, HJB equations, EWMA/ARIMA/OU processes, GARCH/Kalman filtering, Bayesian inference (MCMC), Monte Carlo.
- **Machine Learning:** Deep RL (PPO/DQN), transformers, gradient boosting (XGBoost), ensemble methods.
- **Programming:** Python (PyTorch, Stable-Baselines3, NumPy, Pandas), SQL; git, vectorization.
- **Languages:** English/Mandarin (native), Cantonese (professional).
- **Interests:** Rock climbing, Cooking, Gaming (Top 0.1% Challenger in League of Legends).