ANDY AU

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

Research Interests

Stochastic control with learning; probabilistic analysis of reinforcement learning; continuous-time portfolio optimization; time-series forecasting and volatility modelling; mathematical finance and algorithmic trading.

Education

Boston University, College of Arts & Sciences — Boston, MA
B.A. Applied Mathematics; Minor in Physics — Sep 2023 – May 2026 (expected)
GPA: 3.57/4.00; Dean's List

Relevant coursework: Advanced Stochastic Processes (MA 783); Probability Theory (MA 779); Real Analysis (MA 711); Math of Financial Derivatives (MA 577); Time Series and Forecasting (MA 585); Intro to Stochastic Processes (MA 583); Probability (MA 581); Quantitative Finance Methods (PY 538).

Boston University, Graduate School of Arts & Sciences — Boston, MA
M.A. Mathematics (concurrent enrollment) — Sep 2025 – May 2027 (expected)

Planned coursework: Probability Theory II (MA 780); Statistical Machine Learning (MA 751); Mathematical Deep Learning (MA 752).

Publications & Manuscripts

- Au, A. *Reinforcement Learning for Mean–Variance Portfolio Control under Bayesian Parameter Uncertainty*. Senior thesis, Department of Mathematics & Statistics, Boston University, 2025. Manuscript in preparation.

Research Experience

Senior Thesis, Boston University — May 2025 – Present
Reinforcement Learning for Portfolio Optimization under Drift Uncertainty — Boston, MA

- Extend continuous-time mean–variance portfolio selection to a setting with unknown Sharpe ratio learned via Bayesian updating and Kalman–Bucy filtering.
- Formulate a belief-augmented HJB on state (X_t, m_t, v_t); identify a bilinear wealth drift term $\sigma m_t u_t$ leading to a "quasi-LQG" structure with Gaussian optimal policies but non-classical Riccati equations.
- Current work: analyze the resulting Riccati-type system and approximations, and study convergence to classical mean–variance control as posterior variance vanishes.

Machine Learning Research Intern, HKUST — Jul 2024 – Sep 2024
Deep Learning Research in Time-Series Forecasting — Hong Kong

- Built attentive RNN and transformer architectures for sequential pattern recognition in noisy financial time series.

- Implemented and evaluated PatchTST and Temporal Fusion Transformers for forecasting tasks, with complete training and evaluation pipelines.
- Improved training stability via learning-rate scheduling, gradient clipping, and Optuna-based hyperparameter tuning.

Selected Projects

DRL Portfolio Optimization — PPO Apr 2025 – Aug 2025
Portfolio Optimization with RL and Markowitz Baselines, Boston University

- Implemented a PPO-based daily allocator for 10 liquid ETFs (SPY, QQQ, IWM, EFA, EEM, VNQ, TLT, IEF, GLD, USO) with a 274-dimensional feature stack.
- Enforced diversification via Herfindahl–Hirschman index (HHI) bands to avoid both over-concentrated and near-uniform portfolios.
- Used KL-annealed PPO updates and monthly refits on rolling 90-day windows to adapt to changing regimes.
- On the 2025H1 out-of-sample period, achieved Sharpe 2.00 (24.6% excess return, 12.3% vol) vs rolling Markowitz Sharpe 1.75 (31.7%, 18.1%); ranked in the top 0.01% of 1M Monte Carlo daily-rebalanced portfolios.
- Repository: GitHub.

Volatility Forecasting — PatchTST Aug 2024 – Jun 2025
Realized-volatility forecasting for BTC-USD, HKUST

- Implemented an ensemble of 30 PatchTST models for 5-day-ahead realized volatility using 100-day lookback windows on BTC-USD data.
- Achieved test MSE 0.000106 (95% CI: 0.000108–0.000110), outperforming GARCH and Kalman filter benchmarks by roughly 35% in MSE.
- Evaluated models with MSE, MAE, QLIKE, and directional accuracy; PatchTST attained directional accuracy 0.60 vs 0.49 for the best classical baseline.
- Built a full pipeline for data ingestion, target generation, experiment management, and visualization of predicted vs realized volatility.
- Repository: GitHub.

Technical Skills

- **Quantitative methods:** stochastic calculus, martingales, HJB equations, EWMA/ARIMA/OU processes, GARCH/Kalman filtering, Bayesian inference (MCMC), Monte Carlo simulation.
- **Machine learning:** deep reinforcement learning (PPO, DQN), transformer architectures for time series, gradient boosting (XGBoost), ensemble methods.
- **Programming:** Python (PyTorch, Stable-Baselines3, NumPy, Pandas), SQL; git, vectorization, experiment tracking.
- **Languages:** English and Mandarin (native); Cantonese (professional).

- **Interests:** rock climbing, cooking, gaming (Top 0.1% Challenger in League of Legends).