

MAKSIM SILCHENKO

mthylinao@gmail.com | Singapore | +44 77 2925 4349 | github.com/thylinao1 | [LinkedIn](#)

Portfolio Page: <https://thylinao1.github.io>

EXPERIENCE

Orbuc Research, London, UK

[[Catalyst Agent Prototype](#)] Feb 2026 - Apr 2026

Quantitative Research Intern

- Developed a time-series anomaly-detection model that classifies latent states with Hidden Markov Models to flag unstable, low-predictability conditions, validated across 5 out-of-sample windows where it held up in 4 of 5.
- Built an LLM agent that automatically classifies events from live news streams, backing each decision with retrieval-augmented generation (RAG) over an embedded knowledge base, enforcing schema-constrained output, and persisting results to a SQL (SQLite) database.
- Engineered the statistical analysis behind the firm's March 2026 research report, characterising 9 historical anomaly periods and identifying 4 statistically significant shifts in correlation structure.

The TCM Group, London, UK

[[Website + LLM Prototype](#)] [[GitHub](#)] Jan 2026 - Apr 2026

Data & Analytics Consultant Intern (Bayes Business School Capstone) · Top of the Cohort

- Built a cloud LLM scoring assistant (Vercel, OpenAI) that grades open-ended training responses 0 to 10 against rubrics, deployed with Zod schema validation and automatic retry so malformed model output is never returned.
- Designed a causal-inference framework (RCT-style wait-list control, paired hypothesis tests) to measure leadership-training outcomes, presented it to TCM's senior leadership, and saw it adopted as the standard for ongoing programme evaluation.
- Built an automated evaluation harness that benchmarks the scorer against a hand-labelled gold dataset (0.90 Spearman correlation with human raters), then halved its scoring error with a cross-validated calibration that corrects the model's systematic bias.

Oxford Comma Advisory, London, UK

Oct 2025 - Jan 2026

Data Analyst Intern

- Cut university-shortlisting time from 2 hours to 3 minutes per student by achieving a 73% consultant approval rate on top-5 recommendations through a two-tower neural network trained on sparse implicit feedback with confidence-weighted negative sampling.
- Found organic leads convert about 60% faster than paid channels and that early consultations carry a 2.3x hazard ratio, using Kaplan-Meier curves and a Cox proportional-hazards model on 800+ inquiries (non-converters right-censored).
- Raised booking rates from 18% to 27% by training an XGBoost model (0.79 AUC) to predict which leads were most likely to book, then A/B-testing tailored follow-ups against the standard process outreach.

AWARDS & ACHIEVEMENTS

- Russian National Mathematics Olympiad (MIPT): Winner**, top 0.015% nationally.
- Jane Street 'Puzzle' Competition: Winner** (Feb 2026). **Bloomberg Trading Challenge: Captain**, +32% return vs Bloomberg WLS index.
- Top of the Cohort**: 'Quantitative Methods & Analytics', 'AI & Big Data', 'Capstone Project' (Bayes Business School); 'M&A' (ESADE).
- Full Scholarships**: NUS Exchange (Jan - Jun 2025), ESADE Exchange (Aug - Dec 2024);
- Fully-funded scholarship for Agentic AI Bootcamp** by the Beijing University of Posts & Telecommunications (Jun - Jul 2026).
- Extracurricular Coursework**: Stanford CS229 (Machine Learning), CS230 (Deep Learning), EE178 (Probabilistic Systems); MIT RES.6-012 (Probability); Imperial Mathematics for ML; IBM Applied Data Science Specialization; Statistical Rethinking 2026 (R. McElreath)

EDUCATION

National University of Singapore (NUS), Singapore

Jul 2026 - 2027 (Expected)

MSc in NUS School of Computing (Business Analytics), Specialised in Statistics

Bayes Business School (City, University of London), UK

Sep 2022 - Jun 2026

BSc International Business (Hons), Specialised in AI and Quantitative Methods; First-Class Honours (Highest Distinction)

UCL Economics & Mathematics Foundation Programme, UK

Sep 2021 - Jun 2022

Final grade A* · Mathematics 87%, Highest Grade in the stream

RELEVANT PROJECTS

J.P. Morgan Quantitative Research: Credit Risk & Resource Capacity Optimisation (Applied Project)

[[Project Page](#)] [[GitHub](#)] 2025

- Implemented tabular Q-learning for a finite-horizon inventory-optimisation MDP (deterministic seed, linear epsilon decay, illegal-action handling), and built the chronological 24/24-month evaluation that caught its in-sample edge as an overfit.
- Implemented covariate-shift mitigations (importance weighting, cohort-adaptive thresholding) and a diagnostic that identified the cohort failure as a base-rate problem rather than covariate shift, recommending a 45% break-even repricing instead of a model retune.

High-Dimensional Resource Allocation: Deep Learning & Reinforcement Learning (Independent Project)

[[Project Page](#)] [[GitHub](#)] 2025

- Implemented PPO with Generalised Advantage Estimation (GAE) in TensorFlow/Keras over a 117-dimensional continuous action space, beating an equal-weight baseline by 12.6% in the deterministic environment and 5 to 9% under stochastic and contagion variants.
- Built a custom perturbation-based sensitivity harness that showed a policy under-responsiveness in the trained agent, and implemented Welford running z-score normalisation over the 1,878-dimensional continuous state.

A/B-Test Experimentation Guardrail: SRM Detection & Causal Inference (Independent Project)

[[Project Page](#)] [[GitHub](#)] 2026

- Engineered the A/B-test auditing tool as a production-shaped Python package: typed dataclasses, a custom exception hierarchy, an installable CLI, a 36-case pytest suite, and ruff-linted CI across Python 3.10 to 3.12.
- Made the data loader resilient to messy real-world exports; validating the tool against a 14M-row Criteo advertising dataset surfaced three production bugs (across the loader, the router, and the adapter), each fixed and locked down with a regression test.

SKILLS

Programming & DBs: Python (NumPy, Pandas, scikit-learn, TensorFlow, PyTorch, statsmodels, SciPy, XGBoost, PyMC), SQL, DuckDB, R, C++.

Machine Learning & AI: Supervised & Unsupervised Learning, Gradient Boosting, Neural Networks (CNNs, RNNs, LSTMs, Transformers, Two-Tower), Reinforcement Learning (Q-learning, PPO, Actor-Critic), Recommendation Systems & Embeddings, LLMs, Prompt Engineering, Agentic AI (LangChain, LlamaIndex, tool-use).

Statistical Methods: Causal Inference, A/B Testing & Experimentation, Hypothesis Testing, Survival Analysis (Cox PH, Random Survival Forest, Kaplan-Meier), Bayesian Statistics, Time-Series Analysis, Probability Calibration, Randomised Controlled Trials.

Visualisation, Cloud & MLOps: Tableau, Power BI, matplotlib, seaborn, Plotly; Google Cloud Platform, AWS, Docker; MLOps (CI/CD, GitHub Actions, model deployment, feature pipelines); Git, Jupyter, LaTeX.

Spoken Languages: English (Fluent), Russian (Native), Ukrainian (Native), Belarusian (Fluent), Spanish (Professional Working).