Adam Ru Kun Dong

MSc Data Science & AI | Quantitative Research, Probabilistic ML, Reinforcement Learning

Amsterdam, The Netherlands ark.dong@outlook.com +31 6 33688847 arkdong.com arkdong arkdong

Education

MSc **Delft University of Technology**, Data Science and Artificial Intelligence Technology — Delft, Sept 2026 – Sept 2028 (expected)

- Coursework: Probabilistic AI and Reasoning (Bayesian networks, MDPs, POMDPs, multi-agent decisions), Probabilistic Models and Inference, Elements of Statistical Learning, Machine and Deep Learning.

BSc **University of Amsterdam**, Computer Science (Informatica) — Amsterdam, Sept 2022 – Sept 2026

- Coursework: Scientific Data Analysis, Machine Learning, Algorithms and Complexity.

Research and Projects

Risk-Sensitive Multi-Agent Reinforcement Learning for Trading — Bachelor Thesis, Feb 2026 – July 2026

Studied how quadratic risk penalties change the decisions of market-making and execution agents, trained with independent PPO in JAX on AMZN limit-order-book data. Thesis

- Designed controlled penalty sweeps and a joint grid, training recurrent policies in 4,096 parallel environments with coefficients calibrated on validation days.
- Showed the strongest execution penalty cuts mean squared remaining quantity by 59% and lifts completion from 95.3% to 99.9% on 24 held-out trading days (95% bootstrap intervals, one run per setting).
- Held the counterparty fixed to rule out a changing opponent as the cause, separating real inventory control from trading withdrawal.

ros.ai - Open Workspace for AI Agents with a Fine-Tuned Local Model — Open Source, Sept 2026 – present

Built an open-source workspace where AI agents work on one shared canvas, with its own fine-tuned local model. GitHub Model

- Fine-tuned NVIDIA Nemotron 3 Nano 4B with LoRA (PyTorch, Hugging Face Transformers, PEFT) on 2,955 agent turns distilled from a larger teacher model.
- Raised held-out task success from 33 to 68 of 95, measured with an evaluation pipeline that runs the app's real agent loop on unseen domains.
- Quantized the model to GGUF and shipped it offline with llama.cpp in Mac and Windows desktop apps.

AURORA - AI Content Generation and Evaluation — ABN AMRO x AISO, Mar 2026 – June 2026

Built an editorial workflow for generating and evaluating content grounded in approved ABN AMRO sources. GitHub

- Developed a retrieval and generation pipeline with intent classification, citation checks and KPI-based evaluation (Python, Pydantic, OpenAI API).
- Exposed it through FastAPI endpoints, MCP tools and a Streamlit interface, with inspectable audit trails.

Leadership

AI Student Organization (AISO), Head of Research — Amsterdam, May 2026 – present

- Lead technical AI/ML events, mentored student projects and a quantitative-finance track.

Amsterdam Investment Club (AIC), Head of AI & ML Research — Amsterdam, July 2026 – present

- Lead the AI/ML research agenda for quantitative and agentic trading systems.

Technical Skills

Programming and Tools: Python, C++, C, TypeScript, SQL, Bash, Git, Docker, Linux, PostgreSQL

ML and Statistics: JAX, PyTorch, Hugging Face Transformers and PEFT, scikit-learn, pandas, Polars

Quantitative Methods: RL (PPO), market-making and execution models, bootstrap inference, QuantLib

LLM Systems: LoRA fine-tuning, GGUF quantization, llama.cpp, OpenAI API, MCP, RAG, FastAPI

Additional Training

Certificate in Quantitative Finance (CQF): The CQF Institute, in progress

MITx MicroMasters in Statistics and Data Science: Probability, Fundamentals of Statistics; in progress

C++ Programming for Financial Engineering: QuantNet Baruch MFE, June 2026

Languages and Interests

Languages: English, Chinese, Dutch - CEFR C1 (Advanced)

Interests: AI hackathons, open-source projects, tennis