

Zhou Bojian 周柏健

MSc AI & Business Analytics · ML for Business

✉ neofastio@gmail.com ☎ +852 9102 1799 📄 github.com/zhoubojian-stevenchow in linkedin.com/in/bojian-zhou-62391637b

SUMMARY

MSc AI & Business Analytics candidate at Lingnan University Hong Kong (graduating Aug 2026) with prior audit & finance experience at Shenzhen-based firms. Focused on applied NLP/LLMs, deep reinforcement learning, and big-data ML pipelines. NVIDIA-certified in AI Infrastructure and Generative AI LLMs. Trilingual (Cantonese, Mandarin, English — IELTS 7.5), commuting daily across the HK-Shenzhen corridor.

EDUCATION

Lingnan University, Hong Kong — *MSc in Artificial Intelligence & Business Analytics* Expected Aug 2026

Core coursework: Deep Learning, AI-based Optimization, Big Data Analytics, ML Applications, Database Management, Healthcare Analytics, Computer Vision.

SKILLS

Technical Python, SQL, PyTorch, Hugging Face Transformers, PySpark, XGBoost, scikit-learn, NLP, Deep RL (DQN, Dueling DQN, PER), LLM fine-tuning, statistics, time-series forecasting

Business Problem framing, metrics design, A/B testing, reporting & storytelling, financial analysis, auditing

Tools pandas, Git/GitHub, PostgreSQL, MongoDB, KNIME, Tableau, Selenium/Playwright, Google Colab Pro+ (A100 GPU)

Languages Cantonese (native), Mandarin (native), English (professional, IELTS 7.5)

FEATURED PROJECTS

FinBERT Fine-Tuning for Financial Sentiment CDS525 · 2026

[PyTorch](#) · [Transformers](#) · [FinBERT](#) · [Grid Search](#)

- Designed full factorial grid search: **2 losses × 4 batch sizes × 4 learning rates = 32 training runs** on the Financial PhraseBank dataset.
- Achieved **best macro F1 of 89.92%** (Weighted CE, BS=32, LR=2e-5); handled class imbalance with weighted cross-entropy and focal loss.
- Delivered unified reproducible notebook, ~2,664-word Word report, and GitHub README.
- github.com/zhoubojian-stevenchow/CDS525_Group_Project

NAS-EA + DQN for Vehicle Routing with Time Windows (VRPTW) CDS526 · 2026

[Deep RL](#) · [PyTorch](#) · [Neural Architecture Search](#) · [Evolutionary Algorithms](#)

- Combined evolutionary **Neural Architecture Search** with a **Dueling DQN** agent for constrained routing optimization.
- Implemented action masking ($Q = -\infty$ for infeasible actions); diagnosed depot-reset and time-window feasibility bugs across v6→v8.
- v7 achieved **mean distance 731 with 18% CV (vs. v6's 948 / 77% CV)** — 23% solution-quality gain, 4× stability improvement across 5 seeds.
- github.com/zhoubojian-stevenchow/CDS526_NAS_RL_VRPTW_v7

European Weather Forecasting: XGBoost vs BiLSTM + Attention CDS524 · 2025–2026

[XGBoost](#) · [BiLSTM](#) · [Attention](#) · [Feature Engineering](#)

- Designed a **2×2 factorial experiment**: {XGBoost, BiLSTM+Attention} × {standard lag features, Chinese cosmological calendar features (24 Solar Terms, Wuxing, Heavenly Stems)}.
- Non-obvious finding: Chinese calendar features **significantly improved XGBoost (p = 0.0005)** but harmed LSTM — architectures respond differently to cyclical priors.
- Delivered full report, GitHub README, and a 20-minute narrated presentation.
- github.com/zhoubojian-stevenchow/CDS524-Group-Project

Space Defender: Dueling DQN with Prioritized Experience Replay CDS524 · 2025

[Deep RL](#) · [Dueling DQN](#) · [PER](#) · [Curriculum Learning](#)

- Built end-to-end RL system iterated across **8 versions**; diagnosed reward-hacking behavior (corner-hiding) and redesigned the reward function accordingly.
- Introduced curriculum learning in v8 to progressively harden enemy spawn patterns, reaching **mean score ~1,960**.
- Documented full iteration log as a realistic record of RL debugging, not just a final metric.
- github.com/zhoubojian-stevenchow/cds524_assignment_1_Space_Defender_by_Dueling_DQN_and_PER

ADDITIONAL PROJECTS

- **CDS527 PySpark Starbucks Text Classification** — benchmarked 5 classifiers × 6 embeddings (30 configs); Tuned BERT + Logistic Regression achieved 71.6% accuracy. Co-authored Nike big-data case study.
- **Cantonese ASR Benchmark** — Whisper fine-tuning on Common Voice 17.0 Cantonese subset; built as portfolio for a China Mobile HK R&D internship application.
- **CDS529 LLM Roleplay Benchmarking** — BFI-2 personality recognition and Empath voice-fidelity evaluation across 5 LLMs × 100 fictional characters.
- **CIFAR-10 CNN Dissertation** — Undergraduate thesis on convolutional architectures for image classification.

EXPERIENCE

Auditor / Finance Associate — *Shenzhen-based firms* Prior to MSc

Audit engagements, financial analysis, and client reporting. Developed the business-framing and risk-oriented mindset now applied to ML system design and evaluation.

CERTIFICATIONS

- NVIDIA — AI Infrastructure
- NVIDIA — Generative AI & Large Language Models
- IELTS — Overall Band 7.5