

Yi Yan Ng

London, United Kingdom

✉ ng.yiyan314@gmail.com | 🌐 hectopascal.github.io/ | 📊 Hectopascal | 📄 [ngyiyan](#)

Research Profile

Research engineer focused on LLM post-training and reinforcement learning, with work spanning data selection, credit assignment, offline RL, and model evaluation. Strong systems background in multimodal training, distributed training and production ML, with 4+ years of industry experience.

Research & Open-Source Projects

Probe-Guided Curriculum Learning

Workshop manuscript

Data selection for LLM post-training | GRPO/DAPO, verl

2026

- Contributed post-training experiments and coding evaluation for a study of success-probe-guided prompt selection on OctoThinker-1B, comparing against GRPO, DAPO, and LILO baselines.
- Studied prompt selection using lightweight logistic-regression probes on model activations to identify informative prompts before generating full rollouts.
- Reported study result: probe-only selection achieved 44.7% MBPP mean@4 accuracy versus 44.0% for the DAPO baseline, using 0.7M versus 2.3M rollouts on the 1B model.

Offline Reinforcement Learning for Generative Policies

Manuscript in preparation

Offline RL | Diffusion and flow policies

2026 - Present

- Designed and conducted empirical evaluations of offline RL methods using diffusion and flow policies.
- Ran and analyzed multi-seed experiments and ablations across D4RL locomotion and AntMaze, studying gradient clipping, energy temperature, and candidate perturbations.
- Investigated optimization instability through training diagnostics and candidate-level analyses.

AgentCreditBench

[Code](#)

Credit-assignment evaluation | Exact-oracle diagnostics, CPU benchmarks

2026

- Built an RL diagnostic benchmark using exact advantage oracles to distinguish credit-assignment accuracy from policy-gradient correctness.
- Compared advantage computations across verl, TRL, and OpenRLHF, identifying differences in normalization and edge-case handling.

Multimodal VLM Training & Scaling

[Code](#) | [Technical write-up](#)

Multimodal training systems | PyTorch, LoRA, FSDP, Torch Profiler

2026

- Built a VLM training pipeline from SigLIP-2, Qwen2.5, and an MLP projector, including image-token insertion; trained on 50K LLaVA-Pretrain samples with projector alignment followed by LoRA tuning.
- Scaled to a 1.5B language backbone across 2/4/8 V100 GPUs; achieved 4.07× throughput from 2 to 8 GPUs after correcting a memory-pressure artifact in the baseline.
- Used profiler traces to identify NCCL all-gather as the dominant bottleneck; evaluated forward prefetching and documented the lack of a throughput improvement.

Industry Experience

Meta

London, United Kingdom

Senior Machine Learning Engineer, Integrity

Aug 2024 - Present

- Post-train models using supervised fine-tuning (SFT), distillation, self-distillation, and reinforcement learning (RL); distill large LLMs into smaller students to reduce label-generation inference cost.
- Built an automated LLM-labeling pipeline, benchmarked synthetic labels against human review, and trained production classifiers on synthetic labels, removing manual annotation from the production training loop.
- Shipped multi-task MMoE, GNN entity, and adversarial-embedding detection models, increasing daily detections from 2.4M to 10.8M and PR-AUC by 184% across five violation categories.
- Redesigned an adversarial-detection embedding model's feature architecture with 80+ new features, using 8× less compute and improving recall by 13%.
- Led an 11-person cross-functional workstream moving production models to daily recurrent training, removing approximately 5,280 engineering-hours of overhead while maintaining model quality.

Farfetch UK

Machine Learning Engineer

London, United Kingdom

Jan 2022 - Jul 2023

- Built large-scale ML training and serving pipelines with Airflow, Kubernetes, Spark, and Databricks for recommendation and search teams.
- Profiled GPU/CPU utilization and reduced cloud compute spend by 23% through batch sizing, instance selection, and pipeline restructuring.
- Developed time-series forecasts for seasonal merchant stock demand, informing inventory decisions with an estimated \$1.5M annual revenue impact.

Hypotenuse AI

Machine Learning Engineer | YC-backed AI product startup

Singapore

Dec 2020 - Apr 2021

- Built an end-to-end furniture-tagging ML product: collected 6M+ product listings, curated the training data, and developed a multi-task classifier on GCP (F1 0.73).
- Reduced training time by 20% through input-pipeline optimization and refactoring.

Education

New York University

MSc Computer Science

New York, USA

Apr 2024

Coursework: NLP, Deep Learning, Cloud for ML, Big Data and ML Systems, Monte Carlo Analysis.

University of New South Wales

BSc Statistics; BSc Computer Science (Class 1 Honours)

Sydney, Australia

Aug 2020

- Thesis: Image Registration for Multispectral Pattern Recognition.
- UNSW Global Vice-Chancellor’s Scholarship.
- Coursework: Probability and Stochastic Processes, Statistical Modelling, Neural Networks, Machine Learning, Computer Vision.

Technical Skills & Languages

Programming | Python, C/C++, SQL, Java, R.

Models & methods | LLM post-training, GRPO, SFT, LoRA, distillation, self-distillation, vision-language models, GNNs, multi-task learning, benchmark design.

Frameworks & systems | PyTorch, Hugging Face, verl, vLLM, DDP/FSDP, DeepSpeed, Accelerate, SLURM, NCCL, Torch Profiler, W&B, Airflow, Spark, Kubernetes.

Languages | English (native), Mandarin (native), Japanese (JLPT N1).