

NGUYEN MINH QUANG

mq.nguyen.2023@phdcs.smu.edu.sg $\diamond$ github.com/nbtpj

EDUCATION

Singapore Management University
Doctor of Philosophy in Computer Science

Singapore
January 2024 – present

University of Engineering and Technology, Vietnam National University
Bachelor of Information Technology, Honors Program
CGPA: 3.76/4.0; High Distinction (top 5%); Outstanding Thesis (9.6/10, top 10%)

Hanoi, Vietnam
June 2023

PHD RESEARCH

Singapore Management University (advisor: Prof. Hady W. Lauw)
Reinforcement learning, efficient language models and uncertainty

Singapore
January 2024 – present

- Developed ALH, which augments value-based RL with a learned hypothesis about the environment; it scored about 86 against 67 for TD3 (average normalized score over eight MuJoCo tasks, ICML 2024)
- Derived non-asymptotic plausibility and diversity bounds for diffusion sampling from a fixed starting point (NeurIPS 2026)

RESEARCH PROJECTS

Compact Shared Memory for Language Models

October 2026 – present

- Designed a shared context store using 30–86 bytes per token across 155M–963M-parameter models; at 155M, its cache is $51\times$ smaller than a comparable full-attention cache
- Reduced test perplexity by 29% on books, 35% on papers and 48% on code against the strongest of five architecture-style baselines per domain (155M parameters, 1B training tokens)
- Evaluated scaling to 963M parameters, improving test perplexity in every domain; a separate 412M serving configuration accommodated 256 sequences of 1.05M tokens within 48 GB on an A40
- Developed the training system with AI-assisted coding; in earlier architecture experiments, fit a 100M-token training update into 27.5 GiB on an H200 (about 1B parameters)
- Implemented factored second-order training in the earlier 40M model, reaching AdamW's loss in $0.63\text{--}0.81\times$ the wall-clock time over the first 150 updates; current architecture comparisons use AdamW

Self-Challenging Language Models

June 2026 – October 2026

- Built a self-play training loop in which one model writes problems and another tries to beat them, with no task dataset or judge model; the model learned to write runnable tasks (from under 15% to about 50%)
- Reduced mean task-verification latency $13.8\times$ (0.977 s to 0.071 s) in a Ray Serve/Docker sandbox by replacing per-candidate startup with warm Lean, Python, and C++ workers; tested isolation after reuse
- Built the training stack (PyTorch, Ray, vLLM, FSDP2) in which the trainer and vLLM share each GPU, with native off-policy distributed update, for models up to 27B parameters

Learning to Summarize from Raw Text Alone

February 2026 – May 2026

- Built RL post-training that learns abstractive summarization with no human labels: a summary is rewarded when the original can be rebuilt from it, and a fixed anchor cut drift about $16\times$
- Improved ROUGE-1 and BERTScore in all 15 settings (five models from 0.3B to 8B, three datasets); three LLM judges preferred it in 62–68% of decisive pairs

Robust Policies from Offline Data

January 2025 – May 2025

- For offline RL against a simulated worst-case world, showed that one diffusion model's output spread predicts one-step prediction error about as well as a five-model ensemble (rank correlation 0.76–0.89 on three datasets)

PROFESSIONAL EXPERIENCE

Singapore Management University
Research Engineer

Singapore
July 2023 – December 2023

- Worked on a lifelong-learning model for recommendation at Preferred.AI under Prof. Hady W. Lauw

PUBLICATIONS

Nguyen Minh Quang, Hady W. Lauw. “Per-Starting-Point Plausibility and Diversity Bounds for Score-Based Diffusion Models.” *Advances in Neural Information Processing Systems* (NeurIPS 2026). Accepted.

Nguyen Minh Quang, Hady W. Lauw. “Augmenting Decision with Hypothesis in Reinforcement Learning.” *Proceedings of the 41st International Conference on Machine Learning* (ICML 2024), PMLR 235:41804–41820.

Minh-Quang Nguyen*, Duy-Cat Can*, Hoang-Quynh Le. “Zero-cost Transition to Multi-Document Processing in Summarization with Multi-Channel Attention.” *ECML PKDD 2024*, Research Track, pp. 406–422. *Co-first authors.

Three workshop papers on biomedical summarization: UETrice and UETfishes at MEDIQA 2021 (BioNLP @ NAACL); UETcorn at MEDIQA-Sum 2023 (CLEF).

SKILLS

Engineering Foundations: system design, performance optimization, sandboxing and isolation

AI Agent Orchestration: Claude Code and Codex; parallel coding workflows, shared compute scheduling, automated evaluation and cross-agent code review

Programming and Systems: Python; Slurm, Docker, Lean

ML and RL: PyTorch, Hugging Face; PPO, GRPO, asynchronous RL, replay buffers, offline and continuous-control RL; verifiable rewards and LLM evaluation

Distributed Training and Inference: FSDP2, CPU offload, Megatron tensor parallelism, sequence parallelism, Ray, Ray Serve, vLLM, verl, LoRA, FP8

Languages: Vietnamese (native), English (proficient)

AWARDS AND HONORS

SMU Presidential Doctoral Fellowship in Computing (top 30%) 2024

Outstanding Student Award; Outstanding Thesis Award, UET, VNU 2023

Second place, MEDIQA-Sum 2023 clinical dialogue summarization (ImageCLEF), team 2023

Runner-up, MEDIQA 2021 multi-answer summarization (ACL-BioNLP @ NAACL), team 2021