Richael

Auckland, New Zealand | huangd6666@gmail.com
dh8116.github.io | github.com/dh8116 | x.com/RicaV42

PROFILE

A student researcher working on reinforcement learning for language models and on GPU kernels. Leading a five-person paper on verifier error in RLVR, and publishing one Triton kernel a week, each benchmarked against PyTorch with its numbers.

RESEARCH

RLVR verifier persistence — team lead, five-person paper

Aug 2026 – present. Draft due by the end of 2026, targeting COLM/TMLR

- Two published results disagree: systematic verifier errors collapse training, while natural persistent errors in real code test suites do no harm. We hold the aggregate error rate and asymmetry fixed and sweep persistence (coverage × targeting) to produce one phase map on which both results are single points.
- Mechanism under test: a persistent error corrupts every rollout of a prompt identically, so it moves the GRPO group mean, which is the baseline.
- Scoped as an explanation, not a new algorithm or a state-of-the-art claim. Five streams: past-paper analysis, code, data, modelling, analysis.

GPU SYSTEMS & OPTIMIZATION

One Triton kernel per week, benchmarked against PyTorch and written up at dh8116.github.io/blog.

- Cut fused cross-entropy to 15.90 ms against PyTorch's 24.03 ms at vocab 131,072 (1.51x), at 1.67x less peak memory, by collapsing five live copies of the logits tensor to three.
- Raised fused softmax from ~55 GB/s to ~230 GB/s, and established T4's ~240 GB/s practical ceiling by measurement.

MACHINE LEARNING & EVALUATION

- Fine-tuned Qwen3-14B with LoRA on a 1,526-row corpus spanning five generation paths in six languages, for $1.30 of compute; served it behind vLLM at 2.3s against the 5.7s model it replaced.
- Built held-out evaluations for a production dialogue model: tuned on one set of messages, judged on another. Making each persona's relationship a stake cut stock replies from 5 to 1; an observer-only prompt cut unrequested advice from 11 to 3.
- Fine-tuned MusicGen-small for bar-locked stems and built the eval harness that scores clips the way production uses them.

PROJECTS

- Soulor (soulor.app), founder, June 2026 – present: an AI app to rehearse a hard conversation against a simulated version of the other person; live in six languages with 15 users, served through a multi-provider gateway I built.
- VNportal (vnportal.vercel.app), May 2026 – present: music community platform with AI stem generation.

EDUCATION & AWARDS

Year 11, Auckland, New Zealand. Expected graduation 2028. IB Math AA HL, NCEA Level 2

- Score 10/15, AIME | Honor Roll (Top 5%), AMC 12 | High Distinction, Australian Mathematics Competition
- Honorable Mention, HiMCM; Honorable Mention, IMMC regional; co-authored the modelling papers for both