Jinming Ren

marcobisky | marcobisky.github.io | marcobisky@outlook.com | +86 17882004164

Education

University of Electronic Science and Technology of China (UESTC) Sept 2022 — June 2026
University of Glasgow, Dual Degree Program Sept 2022 — June 2026

- **Major**: Electrical & Computer Engineering BEng; GPA: 3.87/4.0, Ranking: 2/164 (Top 1.2%)
- **Relevant Coursework**: Information Theory, Stochastic Processes, Flow Matching and Diffusion Models, Reinforcement Learning in LLM, etc.

The Chinese University of Hong Kong, Shenzhen Sept 2026

- **Major**: Computer Science MPhil

Research & Projects

LLMlab: Rapid RLVR Post-Training Verification Platform Based on Formal Languages June 2026

- Motivation: To reduce the prohibitively high cost of validating LLM training algorithms, LLMlab enables efficient algorithm performance verification on controlled synthetic data.
- Designed a deterministic formal language supporting efficient RLVR algorithm benchmarking across varying levels of language difficulty.
- Implemented a complete pretrain, SFT, GRPO, KD, OPD, and SDPO pipeline with parallel ablation experiments on 2.67M and 0.15M teacher and student models, achieving 100% accuracy on difficulty levels 0–3 and over 70% accuracy on difficulty levels 4–6.
- Integrated a visualization toolkit including PCA-projected loss landscapes with weight trajectory tracking, per-layer attention heatmaps, and exposure bias measurements for in-depth model interpretability and training dynamics analysis.

LLM Post-Training: Reproducing Self-Distillation Policy Optimization (SDPO) Paper May 2026

- Successfully reproduced the SDPO algorithm using the verl RL framework on RunPod H100 GPUs, building on a thorough understanding of the algorithm.
- Trained Qwen2.5-3B on code generation tasks using LeetCode-style feedback (runtime errors, failed test cases) as the learning signal, and tracked 40 steps of SDPO training dynamics via WandB, logging mean reward and token-level KL divergence relative to the reference policy.

System-level Co-Design of RISCV Accelerators for TinyML at the Edge Sept 2025 — April 2026
Research Assistant, Prof. Yun Li, UESTC

- Engineered a standalone Neural Processing Unit (NPU) for real-time YOLOv8n edge inference on an Artix-7 FPGA, bypassing soft-core processors via a custom RISC-V instruction extension (Xnpu).
- Architected an end-to-end Python ML compiler for automated INT16 quantization and memory-aware instruction scheduling, preserving accuracy within 0.3% mAP of the PyTorch FP32 baseline.
- Designed parameterized RTL operators featuring a 3×3 systolic MAC grid and fully hardware-accelerated post-processing (DFL, NMS), achieving 288 MACs/cycle and 23.4 GOPS peak throughput.
- Integrated asynchronous camera/UDP video pipelines and AXI4 DDR3L memory multiplexing, fully verified via `Cocotb`, `Bazel`, and `Icarus Verilog`.

Design and Visualization of a Complete Single-cycle RV32I CPU Core Jan 2025 — Mar 2025

- Designed a single-core, single-cycle RISCV 32-bit CPU from scratch in `Verilog` for RTL simulation and in `Digital` Software for working principle visualization, open-sourced on Github.
- Built a complete datapath including PC, fetcher, decoder, register file, ALU, LRU-based L1 cache, etc., compatible with basic peripherals: GPIOs, IIC, UART, etc.
- Implemented a boot program in RISCV assembly, basic delay and GPIO libraries in `C`. Compiled and simulated using RISCV GNU toolchain.

Relevant Skills

Programming `Python`, `PyTorch`, `C/C++` , `Makefile`.
Language Native Chinese, Fluent English (IELTS 7).

Awards

Top Academic Scholarship of UESTC (Top 5%) Dec 2023, Dec 2024
China National Scholarship (Top 0.2%) Dec 2024
Outstanding Graduate of Sichuan Province, 2026 Oct 2025