

Yong Zhong

CURRICULUM VITAE

Department of Computer Science and Engineering
Southern University of Science and Technology, Shenzhen, China
12310224@mail.sustech.edu.cn | github.com/OptimistiCompound

RESEARCH STATEMENT

My research concerns how the behaviour of a pretrained foundation model can be controlled and adapted without training a separate model for every task. I approach this from two directions that I take to be the same question at different depths. The first is intervention at inference time, where a small auxiliary network writes into the hidden states of a frozen model to obtain fine-grained control over its output. The second is generation of the adaptation parameters themselves, where the weights of a task-specific module are produced by a hypernetwork rather than learned by gradient descent for each task in isolation. My earlier work in robot learning, on autonomous data generation and on benchmarking spatial understanding, motivated this line by making clear how poorly current models transfer the structure they have already acquired.

EDUCATION

Southern University of Science and Technology

Shenzhen, China

B.Eng. in Computer Science and Technology, expected June 2027

Aug. 2023 – present

- Cumulative GPA 3.77/4.00; weighted average 90.9/100.
- Coursework: Algorithm Design and Analysis (98), Operating Systems (95), Advanced Programming (93), Distributed and Cloud Computing (93), Probability and Statistics for Engineering (91), Computer Organization (90), Computer Networks (89), Digital Logic (92), Data Structures and Algorithm Analysis (91), Discrete Mathematics (85).

University of California, Berkeley

Berkeley, California

Summer Sessions in Computer Science, GPA 4.00/4.00

June 2025 – Aug. 2025

- CS 188 Introduction to Artificial Intelligence (A); CS 161 Computer Security (A).

RESEARCH POSITIONS

National University of Singapore

Singapore

Visiting Research Student, advised by Prof. Yang You and Kai Wang

Feb. 2026 – present

Research on controllable adaptation of language models, covering activation steering by hypernetwork and timestep-conditioned parameter generation for masked diffusion language models.

STAR Lab, Southern University of Science and Technology

Shenzhen, China

Undergraduate Researcher, advised by Prof. Feng Zheng

Apr. 2025 – present

Research on autonomous robotic data generation, in-context policy adaptation and closed-loop evaluation in simulation.

Tsinghua Shenzhen International Graduate School

Shenzhen, China

Research Intern, advised by Prof. Zhi Wang

Sept. 2025 – Jan. 2026

Early-stage integration of vision-language-action policies into a benchmark for spatial understanding in robotic manipulation.

PUBLICATIONS

Peer-Reviewed Conference Papers.

- [C1] Yongzhong Wang*, Keyu Zhu*, **Yong Zhong***, Liqiong Wang, Jinyu Yang and Feng Zheng. “RADAR: Closed-Loop Robotic Data Generation via Semantic Planning and Autonomous Causal Environment Reset.” *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Pittsburgh, 2026. arXiv:2603.11811. *Equal contribution.

Manuscripts in Preparation and Under Review.

- [M1] **Yong Zhong** et al. “Adaptive Multi-Layer Steering with Hypernetworks.” First author. In preparation.
- [M2] “MoRA: Timestep-Conditioned Adaptation for Masked Diffusion Language Models.” Co-author. In preparation.
- [M3] “GeoMIND: A Benchmark for Spatial Understanding in Robotic Manipulation.” Co-author. Under review, *Conference on Neural Information Processing Systems (NeurIPS)*, 2026.

RESEARCH PROJECTS

Adaptive Multi-Layer Steering with Hypernetworks. *National University of Singapore, 2026. First author. Manuscript [M1].*

- Constructed a hypernetwork that produces context-dependent interventions across several hidden layers of a frozen Gemma-2-2B-IT model, conditioned jointly on the live residual stream and on the concept text.
- Established that the original KV-cache formulation was the wrong locus of intervention: joint K/V editing reached 0.444 HMean against 0.973 for the reproduced hidden-state baseline, with the deficit concentrated in concept control rather than in fluency or instruction following.
- Reproduced the FLAS reference under its published protocol on AxBench Concept16k, with held-in and held-out concept splits and GPT-4o-mini judging, and used it as the standard of comparison throughout.
- Compared static against dynamic conditioning, cross-attention against adaptive layer normalisation, and single-layer against coordinated multi-layer writes. The eight-block generator reached 1.013 HMean on held-out concepts, within 0.002 of the published reference of 1.015.
- Characterised the central trade-off of the method through a steering-strength sweep: stronger intervention raises concept expression but degrades instruction following, which argues for adaptive rather than globally fixed intervention.

MoRA: Timestep-Conditioned Adaptation for Masked Diffusion Language Models. *National University of Singapore, 2026. Co-author. Manuscript [M2].*

- Tested whether one LoRA update can serve every denoising stage of a masked diffusion language model. Early Qwen-0.6B results indicated that high-uncertainty recovery and late-stage refinement require different adaptation behaviour, which motivated conditioning on the timestep.
- Compared timestep encodings, rank-level mixing and dynamic gating. Negative results from gating and from full per-step parameter generation narrowed the design toward shared LoRA bases with timestep-dependent routing.
- Analysed update norms, pairwise cosine similarity and effective rank in order to separate strength scheduling from genuine changes of direction in weight space. These experiments informed the final formulation $\Delta W(t) = B \text{diag}(\sigma(t))A$, which the project subsequently took to 78.32% on GSM8K with LLaDA-8B-Base.

RADAR: Closed-Loop Robotic Data Generation. *STAR Lab, SUSTech, 2025–2026. Co-first author. Publication [C1].*

- Held responsibility for Phase 2 of the pipeline: integration and debugging of the Instant Policy execution module in RLBench, collection of the demonstration library, and construction of the atomic, compositional and long-horizon simulation tasks.
- Conducted the full simulation evaluation reported in Table 1, comprising seven atomic and three complex or long-horizon tasks at ten rollouts each, and produced the simulation assets appearing in Figure 3. The pipeline reached up to 90% success on the long-horizon tasks.
- Traced silent reset failures to workspace geometry, object-identity and segmentation-mask mismatches, and extreme inverse-kinematics trajectories, working from point clouds, execution traces and simulator state.
- Proposed the compound reliability interpretation adopted in the paper, under which forward-execution and causal-reset errors accumulate multiplicatively over a collection cycle.

GeoMIND: Spatial Understanding in Robotic Manipulation. *Tsinghua Shenzhen International Graduate School, 2025–2026. Co-author. Manuscript [M3].*

- Integrated preliminary execution interfaces for π_0 , OpenVLA, RDT and UniVLA into a common Isaac Sim and IsaacLab evaluation loop, aligning model-specific inputs, action outputs and simulator interfaces.
- Conducted an initial π_0 fine-tuning study and documented action-space, observation, checkpoint and simulator integra-

tion failures. The large-scale benchmark evaluation was carried out by the wider team after the internship concluded.

NaiveCPU: Single-Cycle and Five-Stage Pipelined RISC-V Processor. *SUSTech, 2025. Course project.*

- Implemented both processors in Verilog and wrote waveform-based testbenches covering instruction execution, pipeline hazards and control flow.
- Localised two silent failures by signal tracing: an interface that had never been instantiated, and a register declared at an incorrect bit width. Code repository.

HONOURS AND AWARDS

National Second Prize, Contemporary Undergraduate Mathematical Contest in Modeling	2025
Provincial Second Prize, Contemporary Undergraduate Mathematical Contest in Modeling	2024
Outstanding Student Award, Southern University of Science and Technology	2025
Outstanding Student Award, Southern University of Science and Technology	2024
First Prize, Freshman Speech Competition, Southern University of Science and Technology	2023

TECHNICAL SKILLS

Research methods	hypernetworks, activation steering, hidden-state intervention, LoRA and parameter-efficient fine-tuning, masked diffusion language models, weight-space analysis, in-context imitation learning
Evaluation	AxBench Concept16k, SteerEval, GSM8K, MATH-500, Math14k, RLBench; controlled ablation, baseline reproduction, failure diagnosis
Analysis	principal component and singular value decomposition, effective rank, cosine geometry, Fisher ratios, layer-wise screening
Frameworks	PyTorch, PyTorch Lightning, Hugging Face Transformers and PEFT, Accelerate, DeepSpeed, timm, einops, Weights & Biases, Isaac Sim and IsaacLab
Programming	Python, C, C++, Rust, Java, SQL, Verilog, RISC-V assembly, $\LaTeX$
Systems	Linux, Git, GDB, Vivado, FPGA development

LANGUAGES

Mandarin Chinese (native). English (IELTS Academic 8.0).

REFERENCES

Prof. Yang You

Presidential Young Professor
Department of Computer Science
National University of Singapore
email available on request

Kai Wang

Principal Researcher
Tencent Hunyuan

email available on request

Prof. Feng Zheng

Department of Computer Science and Engineering
Southern University of Science and Technology
email available on request

Prof. Zhi Wang

Shenzhen International Graduate School
Tsinghua University
email available on request