houzy2023@mail.sustech.edu.cn
kevinhuge.github.io
linkedin

Zhiyu Hou

Education

Southern University of Science and Technology — Shenzhen, China
B.E. in Computer Science and Engineering — Sep 2023 - Jun 2027 *(expected)*

- GPA: 3.65/4.00
- Selected Courses: Algorithm Design and Analysis(98), Computer Security (97), Mathematical Logic(95), Linear Algebra(88), Artificial Intelligence(87), Deep Learning(84)

Publication

- **Budget-Aware Submodel Extraction from Vision-Language Foundation Models**
  Ziyao Wang, **Zhiyu Hou**, Chen Chen, Jingtao Li, Yuhan Liu, Zhiqiang Shen, Ang Li, Lingjuan Lyu.
  (Under review at ECCV 2026)
  Keywords: Vision-Language Models, Model Compression, Submodel Extraction, Efficient AI
  My contribution: Implemented attention/MLP/block dropping and width reduction strategies and benchmarked compressed VLMs.
- **Provable Model Provenance Set for Large Language Models** [arxiv]
  Xiaoqi Qiu, Hao Zeng, **Zhiyu Hou**, Hongxin Wei.
  (Under review at ICML 2026)
  Keywords: LLM Provenance, Statistical Guarantee, Model Auditing
  My contribution: Built MPS pipelines, supported benchmark and prompt construction.

Research Experience

University of Maryland, College Park
Remote — Jan 2026 - Present

- Contributed to a budget-aware submodel extraction framework for Vision-Language Models (VLMs), targeting maximizes task performance under the constraint.
- Contributed to design and implement a heterogeneous structural reduction pipeline, including attention, MLP, block dropping and width reduction, to optimize accuracy–efficiency trade-offs.
- Evaluated on Qwen2.5-VL and LLaVA across multiple VQA benchmarks, achieving strong accuracy under significantly reduced parameter and compute budgets.

Southern University of Science and Technology
Research Assistant — Sep 2025 - Mar 2026

- Contributed to Model Provenance Set (MPS), a provable framework for LLM provenance analysis with statistical coverage guarantees.
- Built token-level and semantic similarity evaluation pipelines, and helped construct a benchmark of 455 Hugging Face LLMs with 100k diverse prompts for large-scale provenance analysis.
- Conducted large-scale experiments in multi-candidate attribution and auditing settings, achieving a coverage rate of **96%** with an average predicted set size of **2.1** at α = **0.05**, and demonstrating superior performance over heuristic baselines.

Awards and Honors

- Second Prize, 18th "Challenge Cup" Guangdong Provincial Competition — 2025.08
- Merit Award, 3rd Entropy Cup Cryptography Security Challenge — 2025.07
- Outstanding Volunteer, SUSTech — 2024.11
- Freshman Scholarship, SUSTech — 2023.09

Skills

Languages: English (fluent), Chinese (native)
Programming: Python, Java
Frameworks & Tools: PyTorch, Transformers, Git, Linux, Vim, LaTeX, Typst