

Tianshuo Peng

Tel: +86 152 251 67693 | Email: pengtianshuo@link.cuhk.edu.hk

Homepage: pengts.github.io | GitHub | Google Scholar

EDUCATION

Chinese University of Hong Kong, PhD

Multimedia Lab

Hong Kong, China

2025-present

- Advised by Xiangyu Yue

Wuhan University, B.S.

Artificial Intelligence, School of Computer Science

Wuhan, China

2021-2025

- GPA: 3.95/4.00 Rank: 1/51
- Advised by Zuchao Li and Lefei Zhang

RESEARCH INTERESTS

- **Agentic Foundation Models:** Agentic LLMs and VLMs for long-horizon task execution, deep research, and scientific discovery.
- **Reasoning & Interaction Patterns:** Novel reasoning and interaction paradigms for enhanced logical reasoning, planning, and decision-making.
- **Multimodal Models & Representation:** Multimodal reasoning, unified multimodal models, and vision-language representation learning.

SCHOLARSHIPS AND AWARDS

- Hong Kong PhD Fellowship Scheme (HKPFS)
- CUHK Vice-Chancellor HKPFS Scholarship
- CCF (China Computer Federation) Elite Collegiate Award (101 students nationwide)
- China National Scholarship (Top 2% nationwide; two-time recipient)
- Outstanding Graduate of Wuhan University (Top 10% schoolwide)

PROJECT EXPERIENCE

Agents-A1: A 35B General Agent Foundation Model for Long-Horizon Tasks

Shanghai AI Lab

- **Core contributor;** led the development of long-horizon search capabilities and full-domain supervised fine-tuning, while contributing deeply across the complete agent-model development lifecycle, including training infra, data construction, and evaluation analysis.
- Official models exceeded **300K downloads** on Hugging Face, while community-derived variants surpassed **500K cumulative downloads**, receiving broad attention and positive feedback from the open-source community.

InternAgent 1 & 1.5: Closed-loop Multi-agent Framework for Autonomous Research

Shanghai AI Lab

- **Core contributor;** supported autonomous research tasks for vision-language alignment training of vision-language models, as well as autonomous research tasks for an LLM-agent-based self-evolving data framework.
- Designed a unified agent system for end-to-end scientific discovery across computational and empirical science domains, supporting deep research, solution optimization, and long-term memory, and providing a general and scalable technical framework for autonomous scientific discovery.

CC-mini: Open-source Project

Ultra-light Agent Harness

- Contributed to the development of cc-mini, a lightweight agent harness; the project has received over **1K stars** on GitHub.

ACADEMIC SERVICE

- **Invited Talk:** Agentic Training: Retrospective, Current Practices, and Future Opportunities — SLAI / Shanghai AI Laboratory
- **Invited Reviewer:** ICLR, NeurIPS, ACL, EMNLP

SELECTED PUBLICATIONS

Scaling the Horizon, Not the Parameters: Reaching Trillion-Parameter Performance with a 35B Agent [Tech Report 2026]

- Agents-AI Team, Shanghai AI Lab ([Tianshuo Peng](#), core contributor)

InternAgent-1.5: A Unified Agentic Framework for Long-Horizon Autonomous Scientific Discovery [Tech Report 2026]

- InternAgent Team, Shanghai AI Lab ([Tianshuo Peng](#), core contributor)

Intern-S1-Pro: Scientific Multimodal Foundation Model at Trillion Scale [Tech Report 2026]

- Intern-S1-Pro Team, Shanghai AI Lab ([Tianshuo Peng](#))

Exploring Reasoning Reward Model for Agents [ACL Findings 2026]

- Kaixuan Fan, Kaituo Feng, Manyuan Zhang, [Tianshuo Peng](#), Zhixun Li, Yilei Jiang, Shuang Chen, Peng Pei, Xunliang Cai, Xiangyu Yue

MME-Reasoning: A Comprehensive Benchmark for Logical Reasoning in MLLMs [ICML 2026]

- Jiakang Yuan*, [Tianshuo Peng](#)*, Yilei Jiang, Yiting Lu, Renrui Zhang, Kaituo Feng, Chaoyou Fu, Tao Chen, Lei Bai, Bo Zhang, Xiangyu Yue (*Equal contribution)

Video-R1: Reinforcing Video Reasoning in MLLMs [NeurIPS 2025]

- Kaituo Feng, Kaixiong Gong, Bohao Li, Zonghao Guo, Yibing Wang, [Tianshuo Peng](#), Junfei Wu, Xiaoying Zhang, Benyou Wang, Xiangyu Yue

InternAgent: When Agent Becomes the Scientist – Building Closed-Loop System from Hypothesis to Verification [Tech Report 2025]

- InternAgent Team, Shanghai AI Lab ([Tianshuo Peng](#), core contributor)

MCP-R1: Generalized Real-World Task Agent Mastering Dozens of Tools [Preprint]

- [Tianshuo Peng](#), Jiakang Yuan, Zhijie Zhong, Yilei Jiang, Bin Wang, Peng Ye, Tao Chen, Lei Bai, Xiangyu Yue, Bo Zhang

Chimera: Improving Generalist Model with Domain-Specific Experts [ICCV 2025]

- [Tianshuo Peng](#), Mingsheng Li, Jiakang Yuan, Hongbin Zhou, Renqiu Xia, Renrui Zhang, Lei Bai, Song Mao, Bin Wang, Aojun Zhou, Botian Shi, Tao Chen, Bo Zhang, Xiangyu Yue

HiddenDetect: Detecting Jailbreak Attacks against Multimodal Large Language Models via Monitoring Hidden States [ACL 2025]

- Yilei Jiang, Xinyan Gao, [Tianshuo Peng](#), Yingshui Tan, Xiaoyong Zhu, Bo Zheng, Xiangyu Yue

Multi-modal Auto-regressive Modeling via Visual Tokens [ACM MM 2024]

- [Tianshuo Peng](#), Zuchao Li, Lefei Zhang, Hai Zhao, Ping Wang, Bo Du

A Novel Energy Based Model Mechanism for Multi-Modal Aspect-Based Sentiment Analysis [AAAI 2024]

- [Tianshuo Peng](#), Zuchao Li, Ping Wang, Lefei Zhang, Hai Zhao

FSUIE: A Novel Fuzzy Span Mechanism for Universal Information Extraction [ACL 2023]

- [Tianshuo Peng](#), Zuchao Li, Lefei Zhang, Bo Du, Hai Zhao

OTHERS

• **Languages:** Mandarin (Native), English (IELTS: 7.0)

• **Skills:** Python, PyTorch, Deep Learning;

Familiar with *megatron* / *transformers* for distributed training, with experience in training 35B&397B scale models; Proficient in developing and maintaining complex projects using Claude Code and Codex.