Junteng Liu

Vicent0205 | Personal Website | Google Scholar | jliugi@connect.ust.hk

Research Interests

Recursive self-improvement (RSI), coding agents, and long-horizon AI development; verifiable task design and reinforcement learning; inference systems; logical reasoning and tool use.

1. How can we train agents to do ultra-long-horizon RSI or AI4AI development?
2. Beyond coding agents, what other domains can we explore?

Education

2024–present **Ph.D. in Computer Science**, Hong Kong University of Science and Technology
HKUST NLP Group; advisor: Prof. Junxian He

2020–2024 **B.Eng. in Automation**, Shanghai Jiao Tong University (IEEE Honor Class)

Research Experience

Apple MLR Research Scientist Intern, Cupertino, USA Apr 2026–Sep 2026
Mentors: Yizhe Zhang and Ruixiang Zhang.

- Developed and publicly released **RSI-Inference** (first author), a benchmark of **37 real-world tasks** evaluating whether coding agents can autonomously develop and maintain **vLLM and SGLang** on H100 and B200 GPUs, covering three stages:
  - **Day-0 model support (10 tasks):** Built tasks for integrating previously unsupported models, evaluated through correctness and serving-feature checks.
  - **Inference optimization (20 tasks):** Built tasks requiring source-code changes to improve serving throughput under fixed deployment settings.
  - **Bug fix (7 tasks):** Built tasks requiring agents to identify and repair performance regressions.
- Sonnet-5.5 and Opus-5.5 ranked first and second overall, achieving **1.68×** and **1.48×** geometric-mean speedups across the **27 optimization and bug-fix tasks**, but only **26.7%** and **13.3%** pass@1 on Day-0 Model Support, respectively. DeepSeek-V4.1-Flash achieved **0.0%** support pass@1 and **1.24×** speedup.

MiniMax Research Intern, Shanghai, China Feb 2025–Feb 2026

- **Logical reasoning RL:** Developed **SynLogic** (first author, NeurIPS 2025), synthesizing rule-verifiable data across 35 logical task types and training reasoning models. SynLogic-32B outperformed DeepSeek-R1-Distill-Qwen-32B on BBEH (**25.5% vs. 19.2%**). This data and RL work contributed to the **MiniMax-M1** release.
- **Search-agent RL:** Developed **WebExplorer** (first author) and conducted **RL training experiments** for long-horizon web agents, using challenging search tasks synthesized through web exploration and query evolution. WebExplorer-8B surpassed WebSailor-72B on BrowseComp-en/zh (Avg@4: **15.7% vs. 12.0%** / **32.0% vs. 30.1%**). The data was used in the **MiniMax-M2** release and formed the basis for subsequent **MiniMax-M3** search-data improvements.

Tencent WXG Research Intern, Shanghai, China Jun 2024–Sep 2024
Mentor: Zifei Shan.

- Investigated vision-encoder bottlenecks in VLM chart understanding and trained improved encoders with contrastive learning. First-author paper: **On the Perception Bottleneck of VLMs for Chart Understanding** (Findings of EMNLP 2025).

Shanghai AI Lab Research Intern, Shanghai, China Jun 2023–Dec 2023
Mentor: Prof. Yu Cheng.

- Investigated a universal truthfulness hyperplane using probes trained on 40+ datasets; found that data

diversity improves cross-dataset generalization. First-author paper: **On the Universal Truthfulness Hyperplane Inside LLMs** (EMNLP 2024).

Publications

* denotes equal contribution. Metrics: Google Scholar / GitHub, Oct 22, 2026.

[1] **Junteng Liu**, Richard He Bai, Huangjie Zheng, Yizhe Zhang, Ruixiang Zhang. RSI-Inference: A Long-Horizon Benchmark for Self-Improving Inference Systems. 2026. [Code] **GitHub Stars: 100+**

[2] MiniMax Team (including **Junteng Liu** as a **core contributor**). The MiniMax-M2 Series: Mini Activations Unleashing Max Real-World Intelligence. arXiv:2605.26494, 2026. [Paper] [Code]

[3] Junlong Li, Wenshuo Zhao, Jian Zhao, et al. (including **Junteng Liu**). The Tool Decathlon: Benchmarking Language Agents for Diverse, Realistic, and Long-Horizon Task Execution. ICLR 2026. [Paper] [Code] **GitHub Stars: 500+**
Reported in official model evaluations by OpenAI, Anthropic, Google DeepMind (Gemini), Moonshot AI (Kimi), DeepSeek, and MiniMax.

[4] Wei Liu, Ruochen Zhou, Yiyun Deng, Yuzhen Huang, **Junteng Liu**, et al. Learn to Reason Efficiently with Adaptive Length-Based Reward Shaping. ICLR 2026. [Paper] [Code]

[5] Aili Chen, Chi Zhang, **Junteng Liu**, et al. DIVE: Scaling Diversity in Agentic Task Synthesis for Generalizable Tool Use. ICML 2026. [Paper] [Code]

[6] **Junteng Liu***, Yunji Li*, Chi Zhang, et al. WebExplorer: Explore and Evolve for Training Long-Horizon Web Agents. arXiv:2509.06501, 2025. [Paper] [Code] **Citations: 80+** **GitHub Stars: 120+**

[7] MiniMax Team (including **Junteng Liu** as a **core contributor**). MiniMax-M1: Scaling Test-Time Compute Efficiently with Lightning Attention. arXiv:2506.13585, 2025. [Paper] [Code]

[8] **Junteng Liu**, Yuanxiang Fan, Zhuo Jiang, et al. SynLogic: Synthesizing Verifiable Reasoning Data at Scale for Learning Logical Reasoning and Beyond. NeurIPS 2025. [Paper] [Code] **GitHub Stars: 200+**

[9] **Junteng Liu**, Weihao Zeng, Xiwen Zhang, et al. On the Perception Bottleneck of VLMs for Chart Understanding. Findings of EMNLP 2025. [Paper] [Code]

[10] **Junteng Liu**, Shiqi Chen, Yu Cheng, Junxian He. On the Universal Truthfulness Hyperplane Inside LLMs. EMNLP 2024. [Paper] [Code]

[11] Shiqi Chen, Miao Xiong, **Junteng Liu**, et al. In-Context Sharpness as Alerts: An Inner Representation Perspective for Hallucination Mitigation. ICML 2024. [Paper] [Code]

[12] Jinghan Zhang, Shiqi Chen, **Junteng Liu**, Junxian He. Composing Parameter-Efficient Modules with Arithmetic Operation. NeurIPS 2023. [Paper] [Code] **Citations: 260+**

[13] Yuzhen Huang, Yuzhuo Bai, Zhihao Zhu, et al. (including **Junteng Liu**). C-Eval: A Multi-Level Multi-Discipline Chinese Evaluation Suite for Foundation Models. NeurIPS 2023. [Paper] [Code] **Citations: 1,000+** **GitHub Stars: 1.8k+**

Honors & Awards

Zhiyuan Honor Scholarship, Shanghai Jiao Tong University (2020–2024).