

Tian Lan (兰天)

Algorithm Expert, Alibaba Group (Token Hub · MaaS) | Ph.D., Beijing Institute of Technology

lantiangmftby@gmail.com | [Google Scholar](#) | [GitHub: gmftbyGMFTBY](#) | [Homepage](#)

Profile

I am currently an Algorithm Expert at Alibaba Group (Token Hub · MaaS), focusing on automatic evaluation and long-horizon agent technologies, driving both frontier research and real-world business applications. I joined Alibaba as a full-time employee through the 2025 Alistar (阿里星) campus recruitment program. I received my Ph.D. in Computer Science and Technology from Beijing Institute of Technology, advised by Prof. Xian-Ling Mao (毛先领); during my research internships I was primarily mentored by Yan Wang (王琰).

My work spans both academia and industry: I have published at top-tier AI venues including ICLR, NeurIPS, ICML, ACL, EMNLP, ACM MM, SIGKDD, and TOIS, with representative works such as *Copy is All You Need*, *Block-Attention*, and *Contrastive Search*, accumulating over 1,700 citations on Google Scholar.

Research Interests

Automatic evaluation of generative models, long-horizon LLM agents (agentic search, agentic memory), retrieval-augmented text generation, and efficient LLM inference/decoding.

Education

Beijing Institute of Technology (BIT)

Beijing, China

Ph.D. in Computer Science and Technology (Combined Master-Ph.D. Program)

Sep. 2019 – Jun. 2025

Advisor: Prof. Xian-Ling Mao (毛先领); DataHammer Lab. Successfully defended dissertation on Jun. 20, 2025.

Beijing Institute of Technology (BIT)

Beijing, China

B.Eng. in Computer Science and Technology

Sep. 2015 – Jun. 2019

Research & Industry Experience

Alibaba Group — Algorithm Expert (Full-time)

2025 – Present

Token Hub (ATH) – MaaS, Agent Systems. Joined via the 2025 Alistar (阿里星) campus recruitment program. Mentors: Longyue Wang, Weihua Luo.

- Lead researcher on **Marco DeepResearch**, Alibaba's flagship agentic-research effort (agentic harness, agentic memory, agentic search, self-evolving skills).
- Designed **HSCodeComp**, an expert-level hierarchical-rule-application agent benchmark with 26 tariff experts and 23 evaluated LLMs/agents; received the **ACL 2026 Best Resource Award**.
- Proposed **Table-as-Search**, a structured table-completion paradigm unifying Deep/Wide/DeepWide agentic information seeking.

Tencent TEG — Research Intern (Intern)

Apr. 2025 – Jun. 2025

AI Platform Department. Research on AI applications in gaming.

- Contributed to **Block-Attention** (ICLR 2025), a block-wise attention mechanism accelerating prefilling for retrieval-augmented LLM inference.

Shanghai Artificial Intelligence Laboratory — Research Intern (Intern)

Jul. 2023 – Jan. 2025

InternLM Post-training Group. Mentor: Wenwei Zhang.

- Built **CriticEval** (NeurIPS 2024), the first systematic benchmark for evaluating LLM critique ability across four dimensions and nine scenarios.
- Proposed a multi-agent feedback framework for training LLMs to critique (EMNLP 2025 Findings), adopted in InternLM post-training.

miHoYo — NLP Lab, Research Intern (Intern)

Feb. 2023 – Jun. 2023

Game AI and domain-specific model post-training. Mentor: Yan Wang.

Tencent AI Lab — Research Intern (Intern)

Jun. 2021 – Jan. 2023

Text generation and retrieval-augmented generation. Contributed to **Effidit**, Tencent AI Lab's AI-powered writing assistant. Mentors: Yan Wang, Deng Cai.

- Proposed **Copy-Generator** (*Copy is All You Need*, ICLR 2023), a copy-based text-generation paradigm.
- Co-developed **Contrastive Search** / **SimCTG** (NeurIPS 2022, Spotlight); the decoding method was later inte-

grated into HuggingFace transformers (v4.24.0) as a built-in decoding strategy.
IBM Research — China — Research Intern (Intern)
Dialogue systems.

2019

Note: all positions prior to Alibaba Group were research internships; Alibaba Group (2025 – present) is my first full-time role.

Selected Publications (full list on [Google Scholar](#), 1,700+ citations)

* indicates equal contribution.

Tian Lan*, Yiqian Yang*, Qianghuai Jia, Li Zhu, Hui Jiang, Hang Zhu, Weihua Luo, Longyue Wang. *HSCodeComp: A Realistic and Expert-level Agent Benchmark for Hierarchical Rule Application*. **ACL 2026. Best Resource Award.**

Tian Lan, Felix Henry, Bin Zhu, Qianghuai Jia, Junyang Ren, Qihang Pu, Haijun Li, Longyue Wang, Zhao Xu, Weihua Luo. *Table-as-Search: Agentic Information Seeking is Table Completion*. **ACL 2026, Findings.**

Yongshi Ye, Hui Jiang, Feihu Jiang, **Tian Lan**, Yichao Du, Biao Fu, Xiaodong Shi, Qianghuai Jia, Longyue Wang, Weihua Luo. *UMEM: Unified Memory Extraction and Management Framework for Generalizable Memory*. **ICML 2026.**

Rong-Cheng Tu*, Zi-Ao Ma*, **Tian Lan***, Yuehao Zhao*, Heyan Huang, Xian-Ling Mao. *Automatic Evaluation for Text-to-image Generation: Task-decomposed Framework, Distilled Training, and Meta-evaluation Benchmark*. **ACL 2025.**

Yi-Fan Lu, Xian-Ling Mao, **Tian Lan**, Tong Zhang, Yu-Shi Zhu, Heyan Huang. *SEOE: A Scalable and Reliable Semantic Evaluation Framework for Open Domain Event Detection*. **ACL 2025.**

Dongyang Ma, Yan Wang, **Tian Lan**. *Block-Attention for Efficient Prefilling*. **ICLR 2025.**

Chen Xu, **Tian Lan**, Yu Ji, Changlong Yu, Wei Wang, Jun Gao, Qunxi Dong, Kun Qian, Piji Li, Wei Bi, Bin Hu. *Decider: A Dual-System Rule-Controllable Decoding Framework for Language Generation*. **IEEE TKDE 2025.**

Shu-Hang Liu, **Tian Lan**, Yun-He Zhang, Ran Liu, Heyan Huang, Chen Xu, Zhijing Wu, Xian-Ling Mao. *Automatic Evaluating Scientific Reviews Through Meta-reviewer's Lens: A Reliable Benchmark for Peer Review Generation*. **NLPCC 2025.**

Tian Lan*, Wenwei Zhang*, Chengqi Lyu, Shuaibin Li, Chen Xu, Heyan Huang, Dahua Lin, Xian-Ling Mao, Kai Chen. *Training Language Models to Critique With Multi-agent Feedback*. **EMNLP 2025, Findings.**

Tian Lan*, Wenwei Zhang*, Chen Xu, Heyan Huang, Dahua Lin, Kai Chen, Xian-Ling Mao. *CriticEval: Evaluating Large-scale Language Model as Critic*. **NeurIPS 2024.**

Tian-Yi Che, Xian-Ling Mao, **Tian Lan**, Heyan Huang. *A Hierarchical Context Augmentation Method to Improve Retrieval-Augmented LLMs on Scientific Papers*. **SIGKDD 2024.**

Tian Lan, Deng Cai, Yan Wang, Yixuan Su, Heyan Huang, Xian-Ling Mao. *Exploring Dense Retrieval for Dialogue Response Selection*. **ACM TOIS 2024.**

Huayang Li, **Tian Lan**, Zihao Fu, Deng Cai, Lemao Liu, Nigel Collier, Taro Watanabe, Yixuan Su. *Repetition In Repetition Out: Towards Understanding Neural Text Degeneration from the Data Perspective*. **NeurIPS 2023.**

Tian Lan, Xian-Ling Mao, Wei Wei, Xiaoyan Gao, Heyan Huang. *Towards Efficient Coarse-grained Dialogue Response Selection*. **ACM TOIS 2023.**

Tian Lan*, Deng Cai*, Yan Wang, Heyan Huang, Xian-Ling Mao. *Copy is All You Need*. **ICLR 2023.**

Yixuan Su, **Tian Lan**, Yan Wang, Dani Yogatama, Lingpeng Kong, Nigel Collier. *A Contrastive Framework for Neural Text Generation*. **NeurIPS 2022, Spotlight.**

Heqi Zheng*, Xiao Zhang*, Zewen Chi*, Heyan Huang, Yan Tan, **Tian Lan**, Wei Wei, Xian-Ling Mao. *Cross-Lingual Phrase Retrieval*. **ACL 2022.**

Yixuan Su, Fangyu Liu, Zaiqiao Meng, **Tian Lan**, Lei Shu, Ehsan Shareghi, Nigel Collier. *TaCL: Improving BERT Pre-training with Token-Aware Contrastive Learning*. **NAACL 2022, Findings.**

Tian Lan, Xian-Ling Mao, Wei Wei, Xiaoyan Gao, Heyan Huang. *PONE: A Novel Automatic Evaluation Metric for Open-Domain Generative Dialogue Systems*. **ACM TOIS 2020.**

Open-Source Projects & Impact

- **Marco DeepResearch** — Alibaba's agentic DeepResearch stack (agentic harness, memory, search). [\[GitHub\]](#)
- **Contrastive Search / SimCTG** — decoding algorithm integrated into HuggingFace transformers (v4.24.0); 460+ GitHub stars. [\[GitHub\]](#) [\[HF Blog\]](#)
- **Copy-is-All-You-Need** — ICLR 2023; 180+ GitHub stars. [\[GitHub\]](#)
- **CriticEval** — open-sourced under OpenCompass/InternLM. [\[GitHub\]](#)
- **PandaGPT** — general-purpose multi-modal instruction-following model; 750+ GitHub stars. [\[Website\]](#)
- **InternLM2** — contributor to Shanghai AI Lab's open-source LLM. [\[arXiv\]](#)
- **HammerLLM** — 1.4B-parameter efficient LLM pre-training framework (DataHammer Lab). [\[GitHub\]](#)
- **Mozi-LLM (science-llm)** — retrieval-augmented LLM toolkit for scientific literature; 120+ GitHub stars. [\[GitHub\]](#)

Awards & Honors

- ACL 2026 Best Resource Award — for HSCodeComp.
- Alistar (阿里星) — Alibaba Group, 2025.
- National Scholarship (国家奖学金) — Beijing Institute of Technology, 2024.
- Tencent Qingyun Program (腾讯青云计划) — Tencent, 2025.
- Tencent Rhino-Bird Elite Intern Program (腾讯犀牛鸟计划菁英实习生) — Tencent, 2022.
- National Scholarship (国家奖学金) — Beijing Institute of Technology, 2017.

Academic Service

Peer reviewer for ICLR, AAAI, NeurIPS, ACL Rolling Review (ARR), ICML, and ACM TOIS.

Invited Talks

- *ATH-MaaS Team ACL 2026 Best Resource Paper Sharing: A New Benchmark for Expert-level Deep Search Agents in Real-world Scenarios* — Alibaba Internal Algorithm Evening Seminar (算法晚自习). Aug. 2026.
- *DeepResearch Agent* — Invited talk (Alibaba Cloud) at ACL 2026. Jul. 2026.
- *Agent Harness Practice Guide: Vertical Harness Design for Deep & Broad Search Tasks* — Alibaba Internal Algorithm Evening Seminar (算法晚自习). Apr. 2026.
- *UMEM* — Invited talk at the China Generative AI Summit 2026. Apr. 2026.
- *Agentic Memory* — Invited talk hosted by Dongfang Li, Harbin Institute of Technology, Shenzhen. Apr. 2026.
- *DeepResearch Agent* — 1st CCF Conference on Large Language Models & AI Engineering. Apr. 2025.
- *Copy is All You Need* — Invited talk hosted by Yutao Zhu, Gaoling School of Artificial Intelligence, Renmin University of China. Mar. 2024.
- *Repetition in Repetition Out* (NeurIPS 2023) — Invited talk hosted by Ruihua Song, Gaoling School of Artificial Intelligence, Renmin University of China. Oct. 2023.