Hongming Zhang

(+1) 856-883-0473 | hzhangal0330@gmail.com | hmzhang.me | panda0881 | Hongming Zhang

Bio

Research Scientist at FAIR, Meta, working on post-training for the agentic capabilities of the Muse Spark model, with a focus on tool use, self-evolving agents, and structured reasoning models. Previously Senior Researcher and Research Lead at Tencent AI Lab, and Research Scholar at the University of Pennsylvania working with Prof. Dan Roth. Ph.D. in Computer Science from HKUST, advised by Prof. Yangqiu Song, and previously earned M.Phil. and Bachelor's degrees from HKUST.

Work Experience

Meta FAIR
Bellevue, WA

Research Scientist
Jul. 2025 - Present

- Founding member of the tool-use workstream for the Muse Spark model, developing its tool-use agentic capabilities.
- Led development of the tool-use agentic RL infrastructure for mid-training and RL-based post-training.
- Developed an agent-based self-evolving data synthesis loop, contributing to Muse Spark's SOTA performance on MCP Atlas (**82.2%**) and near-SOTA performance on τ^2-Bench Telecom (**91.5%** with original prompt, reported by Artificial Analysis).

Tencent AI Lab
Bellevue, WA

Senior Researcher -> Research Lead
Nov. 2021 - Jul. 2025

- Led a team of ten on frontier research in self-evolving LLM agents from Aug. 2023 to Jul. 2025; developed and released **Cognitive Kernel**, a **general-purpose proactive** LLM agent system for real-world interaction.
- Conducted research on **self-evolving agents**, **structured reasoning models**, and agent training for real-world interaction, with publications at ICLR, ICML, NeurIPS, ACL, and other top-tier conferences.

Education

HKUST (Hong Kong University of Science and Technology)
Hong Kong, China

Ph.D. in Computer Science
Sep. 2018 - Nov. 2021

- Supervisor: Prof. Yangqiu Song
- Research topics: **commonsense reasoning**, **selectional preference**, and **coreference resolution**

University of Pennsylvania
Philadelphia, PA

Visiting Scholar
Jan. 2020 - Nov. 2021

- Supervisor: Prof. Dan Roth
- Research topics: **commonsense reasoning** and **information extraction**

HKUST (Hong Kong University of Science and Technology)
Hong Kong, China

M.Phil. in TLE (Technology, Leadership, and Entrepreneurship)
Sep. 2016 - Aug. 2018

- Supervisor: Prof. Yangqiu Song
- Major in Computer Science; research on **graph neural networks** and **information retrieval**

HKUST (Hong Kong University of Science and Technology)
Hong Kong, China

B.E. in Electronic and Computer Engineering
Sep. 2012 - Aug. 2016

- Two minor degrees in Information Technology and Entrepreneurship
- Cumulative grade average: 3.91/4.30; ranking in major: 3/74
- Academic Achievement Medal, the highest academic award for undergraduates at HKUST

Selected Recent Publications

Self-Evolving Agents

- **Cognitive Kernel: An Open-source Agent System towards Generalist Autopilots. Hongming Zhang**, Xiaoman Pan, Hongwei Wang, Kaixin Ma, Wenhao Yu, and Dong Yu. (NAACL 2025 Demo)
  Introduces a general-purpose agent system for realizing self-evolving agents in real-world tasks.
- **R-Zero: Self-Evolving Reasoning LLM from Zero Data.** Chengsong Huang, Wenhao Yu, Xiaoyang Wang, **Hongming Zhang**, Zongxia Li, Ruosen Li, Jiaxin Huang, Haitao Mi, and Dong Yu. (ICLR 2026)
  Explores a novel approach to optimizing a model's task generation capability through self-evolution.
- **WebEvolver: Enhancing Web Agent Self-Improvement with Co-evolving World Model.** Tianqing Fang, **Hongming Zhang**, Zhisong Zhang, Kaixin Ma, Wenhao Yu, Haitao Mi, and Dong Yu. (EMNLP 2025)
  Applies self-evolving algorithms to the agentic domain through a co-evolving world model.

Structured Reasoning Models

- **Parallel-R1: Towards Parallel Thinking via Reinforcement Learning.** Tong Zheng, **Hongming Zhang**, Wenhao Yu, Xiaoyang Wang, He Xing, Runpeng Dai, Rui Liu, Huiwen Bao, Chengsong Huang, Heng Huang, and Dong Yu. (ICLR 2026)
  Introduces reinforcement learning for parallel thinking in structured reasoning models.
- **Streaming Looking Ahead with Token-level Self-reward. Hongming Zhang**, Ruixin Hong, and Dong Yu. (Technical Report 2025)
  Proposes token-level self-reward for efficient look-ahead reasoning in streaming settings.
- **Scaling Test-Time Compute for Agentic Coding.** Joongwon Kim, Wannan Yang, Kelvin Niu, **Hongming Zhang**, Yun Zhu, Eryk Helenowski, Ruan Silva, Zhengxing Chen, Srinivasan Iyer, Manzil Zaheer, Daniel Fried, Hannaneh Hajishirzi, Sanjeev Arora, Gabriel Synnaeve, Ruslan Salakhutdinov, and Anirudh Goyal. (in submission to COLM 2026)
  Studies how to scale test-time compute for structured reasoning in agentic coding tasks.

For a full publication list, see Google Scholar (total citations: 8,339; h-index: 37).

Program Committees

- **Journal Editorial Board**: Semantic Web Journal special edition on Commonsense Knowledge and Reasoning; Frontiers in Big Data.
- **Journal Reviewer**: Applied Network Science; Computer Science and Language; IEEE Transactions on Neural Networks and Learning Systems; NeuralComputing.
- **Area Chair and Conference Committee**: ICLR'24-26; ICML'25; NeurIPS'23-26; COLM'24-25; AACL'23; ACL'21-24; EMNLP'21-24; CVPR 22-24; NAACL'21-24; IJCAI'21-23; AAAI'21-23; COLING'20; AACL'20; CIKM'19; AKBC'21.
- **Others**: IJCAI 2019 WebMaster.

Honors & Awards

2024	**Outstanding Paper**, Transactions of Machine Learning Research	*Bellevue, WA, U.S.A*
2023	**Outstanding Paper**, EMNLP	*Singapore*
2022	**Engineering PhD Research Excellence Finalist Award**, HKUST	*Hong Kong, China*
2020	**SENG Academic Award for Continuing PhD students**, HKUST	*Hong Kong, China*
2019	**Microsoft Research Asia Fellowship**, Microsoft Research Asia	*Beijing, China*
2019	**Tencent Rhino-Bird Scholarship**, Tencent	*Shenzhen, China*
2018	**Hong Kong PhD Fellowship**, Hong Kong	*Hong Kong, China*
2018	**Excellent Research Award**, HKUST	*Hong Kong, China*
2016	**Academic Achievement Medal**, HKUST	*Hong Kong, China*
2014	**Academic Excellence Award**, HKUST	*Hong Kong, China*