

Jiaqi Shao

PhD, HKUST | Expected Graduation: June 2027

Target: LLM Agent / RL / Agent Systems

 [GitHub](#)  [Personal Webpage](#)  [Email](#)  [Google Scholar](#)

■ Education

Hong Kong University of Science and Technology

Doctor of Philosophy (PhD) in Electronic and Computer Engineering

2023 Fall — Present

- Supervisor: Prof. Wei Zhang (HKUST) (Long-term Collaborator: Prof. Bing Luo (DKU))

The Chinese University of Hong Kong, Shenzhen

Bachelor of Engineering in Electrical and Computer Engineer

2019 — 2023

- Stream: Computer Engineering

■ Research Focus

Developing **long-horizon LLM agents**: (1) Agentic RL algorithms that stabilize multi-turn optimization under non-stationary context (FoldAct); (2) Evaluation methodology that moves beyond end-task accuracy to measure how agents gather, revise, and calibrate evidence in the loop (SeekBench, ICLR 2026); and (3) Agent systems that support sustained autonomous execution with human-in-the-loop controllability at production scale (ByteDance).

■ Experience

ByteDance | Algorithm Intern (Long-running Agent Systems)

System Design & Implementation

Jan. 2026 – Present

- Owned the end-to-end delivery of a long-running agent system for sustained execution, iterative refinement, and human-in-the-loop intervention.
- Developed a unified framework that connected sustained execution, process evaluation, and iterative improvement for long-horizon agent behavior.
- Supported **Auto**, **Interactive**, and **Human-Interrupt** modes, making the system more controllable in both autonomous and human-in-the-loop settings.
- Established recovery and handoff mechanisms for failures and intervention, improving robustness and reliability in runs.

■ Representative Research

SeekBench: Benchmarking Epistemic Competence in LLM Search Agents

First Author | Benchmark Design & Evaluation Methodology

ICLR 2026

[\[Paper\]](#)

Skills: **Benchmark Design** | **Evaluation Methodology** | **Research Leadership**

- Developed **SeekBench**, a standardized benchmark for evaluating the epistemic competence of LLM search agents and exposing failures beyond task accuracy alone.
- Defined a trajectory-level evaluation paradigm that analyzes how agents gather, revise, and calibrate evidence throughout the search process.
- Defined three core metrics *Groundedness*, *Recovery*, and *Calibration* to quantify evidence integration, adaptation to low-quality information, and evidence-aware decision-making.

FoldAct: Efficient and Stable Context Folding for Long-Horizon Search Agents

First Author | Algorithm Design & Long-Horizon Optimization

arXiv preprint, 2025

[\[Paper\]](#)

Skills: **Agentic RL** | **Long-horizon Optimization** | **Context Folding**

- Identified the core scaling bottleneck in long-horizon search-agent RL as training non-stationarity introduced by continually updated summarization policies, rather than context length alone.
- Proposed **FoldAct**, a context-folding optimization framework with a structured folding-training pipeline that stabilizes summary-policy learning while improving efficiency under long-context constraints.
- Delivered up to **5.19×** training speedup on complex search-agent tasks without sacrificing long-horizon decision quality, validating context folding as a practical scaling primitive for agent optimization.

MorphAgent: Self-Evolving Multi-Agent Collaboration Platform

Co-first Author | System Architecture & Adaptive Collaboration
[Paper]

ICML-MAS, 2025

- Designed a decentralized collaboration framework where LLM agents dynamically evolve roles without predefined structures.
- Established a self-evolving collaboration mechanism that improved task performance, transferability, and robustness across reasoning and coding benchmarks.

Duke Kunshan University

Research Mentor

Feb. 2025 – Jan. 2026

- Coordinated a student research pipeline spanning project scoping, implementation, experiments, and paper writing, while mentoring **10+ undergraduate students** on LLM-based narrative systems, memory, and human-AI collaboration.
- Organized collaborative exchange activities with Duke-related partners, including project presentation and forum-style exhibition, strengthening cross-institution communication and early team coordination.

MASArena: Benchmarking Framework for Multi-Agent Systems

Principal Lead | Systems Architecture & Evaluation Infrastructure

May 2025 – July 2025 [Project]

- Identified a core systems bottleneck in multi-agent research: fragmented evaluation pipelines made reproduction difficult and cross-system comparison unreliable; formulated a unified evaluation workflow around this problem.
- Led the architecture and implementation of **MASArena**, an open-source benchmarking stack that unifies task execution, experiment orchestration, result management, and visual analytics across single-agent and multi-agent systems.
- Built reusable, extensible evaluation infrastructure for standardized comparison across agents, tools, and datasets, strengthening reproducibility and interpretability. [Code contributions]

■ Other Research & Projects

Beyond Right to be Forgotten: Managing Heterogeneity Side Effects Through Strategic Incentives

First Author | Incentive Mechanism Design & Theoretical Analysis

ACM MobiHoc 2025

Skills: **Federated Unlearning** | **Incentive Mechanism Design** | **Game Theory** | **Theoretical Analysis**

- Studied heterogeneity side effects in federated unlearning under non-IID data, where removing a client can disproportionately degrade similarly distributed remaining clients and destabilize participation incentives.
- Developed a theoretical framework and a Stackelberg-game-based payment mechanism to retain crucial clients, improving global stability by up to 6.23%, reducing worst-case client degradation by 10.05%, and achieving up to 38.6% runtime efficiency over full retraining.

FedCampus / FedKit: Cross-Platform Federated Learning System for Smart Campus

- Built a real-world smart-campus mobile system integrating federated learning and differential privacy, and developed a cross-platform pipeline for Android and iOS covering model conversion, hardware-accelerated training, and on-device aggregation; deployed to **100** customized smart watches and Android/iOS clients, with related work accepted at **IEEE INFOCOM 2024 Demo**.

■ Teaching

Teaching Assistant

- ELEC3120: Computer Communication Networks (HKUST, Spring 2024)
- ELEC3300: Introduction to Embedded Systems (HKUST, Fall 2024)
- Vector Space Methods with Applications | ECE 586K (DKU, Spring 2025)