

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

• Supervisor: Prof. Wei Zhang (HKUST) (Mentor: Prof. Bing Luo (DKU))

2023 Fall — Present

The Chinese University of Hong Kong, Shenzhen

Bachelor of Engineering in Electrical and Computer Engineer

• Stream: Computer Engineering

2019 — 2023

Research Focus

Developing **long-horizon LLM agents** from three complementary perspectives: **algorithms** for stable multi-turn optimization, **evaluation** for epistemic competence and reliability, and **systems** for controllable autonomous execution in real-world settings.

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 extended runs.

Representative Research

SeekBench: Benchmarking Epistemic Competence in LLM Search Agents

First Author | Benchmark Design & Evaluation Methodology

ICLR 2026

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

arXiv preprint, 2025

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

- Proposed **FoldAct**, a context-folding algorithm for long-horizon LLM agents, addressing instability caused by non-stationary summary policies in multi-turn reinforcement learning.

- Developed a structured training strategy that stabilizes summary-policy learning and improves optimization efficiency in multi-turn settings.
- Achieved up to $5.19\times$ training speedup on search-agent tasks while maintaining strong long-horizon decision quality.

MorphAgent: Self-Evolving Multi-Agent Collaboration Platform

Co-first Author | System Architecture & Adaptive Collaboration

ICML-MAS, 2025

- Designed a decentralized collaboration framework in which 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 | Framework Architecture

May 2025 – July 2025

- Led the architecture and implementation of an open-source benchmarking framework for standardized experimentation and visual analytics in single-agent and multi-agent systems.
- Established reusable evaluation infrastructure for agents, tools, and datasets through extensible experiment management and visual analysis, improving 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.

FedKit: Cross-Platform Federated Learning for Android and iOS

- Developed a cross-platform federated learning pipeline covering model conversion, hardware-accelerated training, and on-device aggregation; accepted at **IEEE INFOCOM 2024 Demo**.

FedCampus: Privacy-preserving Mobile Application for Smart Campus

- Built a real-world mobile application integrating federated learning and differential privacy, with deployment to **100** customized smart watches and Android/iOS clients.

■ 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)