Hao Qiu
hao.qiu@nus.edu.sg
haoqiu95.github.io/
School of Computing, National University of Singapore
COM2, 15 Computing Drive, Singapore 117418

Research Interests

I am broadly interested in the theoretical foundations of machine learning and the development of algorithms with provable guarantees. Currently, I focus on adaptive and efficient sequential decision-making in **multi-agent systems**. My research topics include:

- **Centralized Learning with Asynchronous Feedback:** I develop adaptive online learning methods that integrate asynchronous feedback and exploit unknown problem structure.
- **Decentralized Learning with Distributed Feedback:** I develop cooperative decision-making methods using local observations and shared information, with guarantees under communication constraints.

I am also interested in applying these methods to **AI agent routing** in multi-agent and LLM-based systems, enabling adaptive and efficient selection of specialized agents, models, or tools for complex tasks.

Appointments

02/2026–Present	**Postdoctoral Researcher**, National University of Singapore, Singapore School of Computing Advisor: Jonathan Scarlett

Education

10/2022–12/2025	**Ph.D. in Computer Science**, University of Milan, Italy Department of Computer Science Supervisor: Nicolò Cesa-Bianchi Thesis: *Distributed and Delayed Online Learning*
09/2017–10/2020	**M.Sc. in Computer Science**, University of Trento, Italy Department of Computer Science Supervisor: Giovanni Iacca
09/2013–06/2017	**B.Sc. in Computer Science**, Wenzhou Business College, China Department of Computer Science

Experience

09/2024–02/2025 06/2023–08/2023	**Visiting Researcher**, Centrum Wiskunde & Informatica (CWI), Netherlands Supervisor: Wouter M. Koolen
08/2021–09/2022	**Data Scientist**, Juneng Environmental, China
02/2021–07/2021	**Data Scientist**, Accenture, China

Publications

Google Scholar † Equal contribution or alphabetical order

Preprints

P1. Juliette Achddou†, Nicolò Cesa-Bianchi† and **Hao Qiu**†. Distributed Online Optimization with Stochastic Agent Availability.

P2. **Hao Qiu**†, Mengxiao Zhang† and Y. Wan. Decentralized Online Exp-Concave Optimization.

P3. **Hao Qiu**†, Andrew Jacobsen†, Emmanuel Esposito and Mengxiao Zhang. Universal Online Convex Optimization with Delayed Feedback.

P4. Ruixu Hu†, **Hao Qiu**† and Wenying Xu. Gradient-Variation and Small-Loss Regret for Decentralized Online Convex Optimization.

P5. **Hao Qiu**, Yu-Jie Zhang and Jonathan Scarlett. Near-Optimal Regret in Decentralized Stochastic Linear and Kernel Bandits.

Conference Papers

C1. Yu-Jie Zhang†, **Hao Qiu**†, Jonathan Scarlett and Kevin Jamieson. Near-Optimal Regret in Adversarial Kernel Bandits. 40th Conference on Neural Information Processing Systems **(NeurIPS 2026)**

C2. Andrew Jacobsen, Emmanuel Esposito, **Hao Qiu** and Mengxiao Zhang. Strongly Adaptive Online Learning with Time-Varying Movement Cost. 40th Conference on Neural Information Processing Systems **(NeurIPS 2026)**

C3. **Hao Qiu**†, Andrew Jacobsen†, Emmanuel Esposito† and Mengxiao Zhang†. Parameter-free Dynamic Regret: Time-varying Movement Costs, Delayed Feedback, and Memory. 43rd International Conference on Machine Learning **(ICML 2026)**

C4. **Hao Qiu**†, Mengxiao Zhang† and Nicolò Cesa-Bianchi. Near-Optimal Regret for Distributed Adversarial Bandits: A Black-Box Approach. 39th Annual Conference on Learning Theory **(COLT 2026)**

C5. **Hao Qiu**, Mengxiao Zhang and Juliette Achddou. Decentralized Online Convex Optimization with Unknown Feedback Delays. 40th AAAI Conference on Artificial Intelligence **(AAAI 2026)**

C6. Jingyuan Liu†, **Hao Qiu**†, Lin Yang and Mengfan Xu. Distributed Multi-Agent Bandits Over Erdős–Rényi Random Networks. 39th Conference on Neural Information Processing Systems **(NeurIPS 2025)**

C7. Haoran Zhang, Yang Xu, Xuchuang Wang, Hao-Xu Chen, **Hao Qiu**, Lin Yang and Yang Gao. Federated Multi-armed Bandits with Efficient Bit-Level Communications. 39th Conference on Neural Information Processing Systems **(NeurIPS 2025)**

C8. Haoran Zhang, Xuchuang Wang, Hao-Xu Chen, **Hao Qiu**, Lin Yang and Yang Gao. Near-Optimal Regret Bounds for Federated Multi-armed Bandits with Fully Distributed Communication. 41st Conference on Uncertainty in Artificial Intelligence **(UAI 2025)**

C9. **Hao Qiu**[†], Emmanuel Esposito[†] and Mengxiao Zhang[†]. Exploiting Curvature in Online Convex Optimization with Delayed Feedback. 42nd International Conference on Machine Learning **(ICML 2025)**

C10. Emmanuel Esposito[†], Saeed Masoudian[†], **Hao Qiu**, Dirk van der Hoeven, Nicolò Cesa-Bianchi and Yevgeny Seldin. Delayed Bandits: When Do Intermediate Observations Help? 40th International Conference on Machine Learning **(ICML 2023)**

C11. Dirk van der Hoeven[†], Ciara Pike-Burke[†], **Hao Qiu** and Nicolò Cesa-Bianchi. Trading-Off Payments and Accuracy in Online Classification with Paid Stochastic Experts. 40th International Conference on Machine Learning **(ICML 2023)**

C12. **Hao Qiu**, Leonardo Lucio Custode and Giovanni Iacca. Black-box Adversarial Attacks Using Evolution Strategies. Genetic and Evolutionary Computation Conference Companion **(GECCO 2021)**

Awards & Honors

2023	ELISE Mobility Grant
2022	Doctoral Scholarship, University of Milan
2019	Erasmus+ Traineeship Scholarship
2016	1st-Class Scholarship for Outstanding Students, Wenzhou Business College
2016	National 2nd Prize, China Undergraduate Mathematical Contest in Modeling
2015	Provincial 1st Prize, China Undergraduate Mathematical Contest in Modeling
2015, 2014	3rd-Class Scholarship for Outstanding Students, Wenzhou Business College

Talks

T1. **Hao Qiu**. Near-Optimal Regret for Distributed Adversarial Bandits: A Black-Box Approach.
39th Annual Conference on Learning Theory. 07/2026.

T2. **Hao Qiu**. Distributed Multi-Agent Bandits Over Erdős–Rényi Random Networks.
Zhejiang University. 04/2026.
Southeast University. 04/2026.
National University of Singapore. 03/2026.

T3. **Hao Qiu**. Distributed Online Optimization with Stochastic Agent Availability.
Inria. 10/2024.

Research Mentoring

2026 **Southeast University** (with Prof. Wenying Xu)
Ruixu Hu (Ph.D. student at Southeast University)
Research topics: Nearly Optimal Regret for Decentralized Online Riemannian Optimization (submitted to AAAI 2027).
Gradient-Variation and Small-Loss Regret for Decentralized Online Convex Optimization.

2025 **University of Milan** (with Prof. Nicolò Cesa-Bianchi)
Luca Zhao (Undergrad. at the University of Milan → Master's student at the University of Milan)
Research topic: Stochastic Bandit Algorithms for Crowdsourcing under Censored Feedback.

2025 **Nanjing University** (with Prof. Lin Yang)
Jingyuan Liu (Undergrad. at Nanjing University → Ph.D. student at the University of California, Berkeley)
Research topic: Distributed Multi-Agent Bandits Over Erdős–Rényi Random Networks (NeurIPS 2025).

2024 **University of Milan** (with Prof. Nicolò Cesa-Bianchi and Juliette Achddou)
Julien Lou (Research intern at the University of Milan → Ph.D. student at Inria)
Research topic: Learning with a prior in bandits.

Academic Service

Area Chair

2027 International Conference on Learning Representations (ICLR)
2027 International Conference on Machine Learning (ICML)

Journal Reviewer

Journal of Machine Learning Research (JMLR)
IEEE Transactions on Signal Processing (TSP)

Conference Reviewer

2026–27 AAAI Conference on Artificial Intelligence (AAAI)
2026 International Conference on Learning Representations (ICLR)
2025–26 Advances in Neural Information Processing Systems (NeurIPS)
2024–26 International Conference on Machine Learning (ICML)
2025 European Workshop on Reinforcement Learning (EWRL)

Volunteer

2026 Annual Summer School on Mathematical Aspects of Data Science
2025 36th International Conference on Algorithmic Learning Theory

---

Last updated: October 21, 2026