

FedKDD/FedMAS 2026: The 2026 International Joint Workshop on Federated Learning for Multi-agent Systems and Data Mining

Haozhao Wang
Huazhong University of Science and
Technology
Wuhan, China
hz_wang@hust.edu.cn

Zhuangdi Zhu
George Mason University
Fairfax, United States
zzhu24@gmu.edu

Zheng Xu
Google Research
Mountain View, United States
xuzheng@google.com

Yi Liu
Department of Computer Science,
City University of Hong Kong
Hong Kong, China
yliu247-c@my.cityu.edu.hk

Dongyuan Li
Center for Spatial Information
Science, The University of Tokyo
Tokyo, Japan
lidy@csis.u-tokyo.ac.jp

Ming Hu
East China Normal University
Shanghai, China
ecnu_hm@163.com

Mingyue Cheng
University of Science and Technology
of China
Hefei, Anhui Province, China
mycheng@ustc.edu.cn

Junyuan Hong
University of Texas at Austin
Austin, United States
mr.junyuan.hong@gmail.com

Nathalie Baracaldo
IBM
San Jose, United States
baracald@us.ibm.com

Neil Shah
Snap
Seattle, United States
nshah@snap.com

Salman Avestimehr
University of Southern California
Los Angeles, United States
avestimehr@gmail.com

Jiayu Zhou
Michigan State University
East Lansing, United States
jiayuz@umich.edu

Carl Yang
Emory University
Atlanta, United States
j.carlyang@emory.edu

Abstract

Multi-agent systems (MAS) enable complex applications in autonomous driving, smart logistics, and distributed sensing, relying on collective intelligence from agents with sensitive data. Federated learning (FL) serves as a decentralized framework for these agents to learn shared models while preserving privacy and autonomy. Despite progress, integrating FL and MAS remains challenging. Key issues include coordinating heterogeneous agents, handling non-IID data, ensuring communication efficiency, and maintaining robustness and fairness in open environments. This workshop brings together researchers to explore this convergence. We aim to foster discussions on foundational advances, real-world deployments, and interdisciplinary opportunities, focusing on scalability, trustworthiness, adaptive coordination, and the societal impact of federated multi-agent intelligence.

CCS Concepts

• Information systems → Data mining.

Keywords

Federated Learning, Distributed Data Mining, Multi-Agent system

ACM Reference Format:

Haozhao Wang, Zhuangdi Zhu, Zheng Xu, Yi Liu, Dongyuan Li, Ming Hu, Mingyue Cheng, Junyuan Hong, Nathalie Baracaldo, Neil Shah, Salman Avestimehr, Jiayu Zhou, and Carl Yang. 2026. FedKDD/FedMAS 2026: The 2026 International Joint Workshop on Federated Learning for Multi-agent Systems and Data Mining. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2 (KDD '26)*, August 09–13, 2026, Jeju Island, Republic of Korea. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3770855.3818248>

1 Introduction

Collaborative intelligence of autonomous agents underpins complex systems such as networked robotics and smart infrastructure. Robust multi-agent system (MAS) demands learning from distributed, privacy-sensitive data. Strict privacy regulations and agent autonomy requirements hinder conventional centralized training. Federated learning (FL) resolves this issue by enabling collaborative training without raw data sharing, securing privacy and supporting



This work is licensed under a Creative Commons Attribution 4.0 International License.
KDD '26, Jeju Island, Republic of Korea
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2259-2/2026/08
<https://doi.org/10.1145/3770855.3818248>

decentralized knowledge fusion. Integrating FL with MAS poses distinct challenges. Multi-agent settings involve heterogeneous objectives, intricate dependencies, non-stationary data and dynamic topologies, which complicate coordination, impede knowledge aggregation, and compromise system robustness, fairness and scalability. Moreover, agent operation in adversarial environments makes trust, security and anti-malicious resilience critical.

This workshop connects FL and MAS, focusing on fundamental progress and practical deployment via four core themes: 1) Coordination and Scalability: Innovative algorithms for scalable coordination of heterogeneous agents, covering communication-efficient protocols, incentive mechanisms and system heterogeneity adaptation. 2) Trustworthiness and Safety: Security, privacy, robustness and ethical alignment frameworks, including adversarial robustness, differential privacy, Byzantine resilience and fair resource allocation. 3) Emergent Behaviors and Adaptive Learning: Research on global intelligence emerging from decentralized FL, such as non-stationary environment learning, adaptive personalization and complex agent relationship modeling. 4) Interdisciplinary Applications and Benchmarking: Explorations in high-stakes scenarios including autonomous fleets, smart grids and healthcare, along with realistic simulations, unified evaluation frameworks and open benchmark platforms. We also target challenges in integrating MAS with foundation models and generative AI. Gathering researchers and practitioners, this workshop fosters innovative research to advance the trustworthiness, scalability and practical impact of federated multi-agent intelligence.

2 Audience and Schedule

Diversity Commitment. During the selection of organizers and speakers, we commit to encouraging diversity. Both groups achieved gender parity and represent diverse affiliations, races, and nationalities across Asia, Europe, North America, and beyond. The full scale of scientific seniority is covered, from PhD candidates to full professors. Our invited talks present broad perspectives on the intersection of federated learning and multi-agent coordination.

Schedule. The format of the workshop is a half-day event (4 hours). We will invite 4 keynote speakers who are leading experts from academia and industry. The potential speakers include Prof. Rui Zhang from HUST, Prof. Zhenzhe Zheng from SJTU, Prof. Gauri Joshi from CMU. We will have contributed papers as both oral and poster presentations for accepted ones. We will highlight outstanding papers in the award ceremony.

Table 1: Workshop schedule.

Time	Event
8:00 am - 8:10 am	Opening Ceremony
8:10 am - 8:50 am	Keynote Talk 1
8:50 am - 9:30 am	Keynote Talk 2
9:30 am - 10:10 am	Keynote Talk 3
10:10 am - 10:50 am	Coffee Break and Poster Presentation
10:50 am - 11:30 am	Keynote Talk 4
11:30 am - 11:50 am	Best Paper Presentation
11:50 am - 12:00 pm	Award Ceremony and Closing Session

3 Submission Guidelines

Paper submissions: June 1st, 2026. Paper notifications: June 24th, 2026. Workshop date: August 9. We invite short technical papers

- up to 5 pages including references and unlimited pages of appendix. All manuscripts should be submitted in a single PDF file including all content, figures, tables, and references, following the new Standard ACM Conference Proceedings Template. Additionally, papers must be in the two-column format, with the recommended setting for Latex file: `\documentclass[sigconf, review]{acmart}`. Papers should be submitted to openreview website at <https://openreview.net/group?id=KDD.org/2026/Workshop/FedKDD>. While all accepted papers will be presented with posters, high-quality accepted papers will also have the opportunity to participate in the oral/spotlight presentation and win our Best Paper Award(s). We will also present accepted papers on our website: <https://fedmas-fedkdd.github.io/fedmas-fedkdd2026/>.

4 Review Process

The double-blind review process will be adopted. We will attract a diverse program committee, and each submission will receive at least 3 reviews. Organizers will make transparent decisions. To discourage finalized works, novelty and scope will be explicitly reviewed. Reviewers must state conflicts of interest to avoid bias.

5 Previous Workshop

This proposal follows directly from the successful FedKDD 2025: International Joint Workshop on Federated Learning for Data Mining and Graph Analytics, held in conjunction with KDD 2025. FedKDD 2025 focused on advancing federated methods for tabular, sequential, and graph-structured data. Building on this foundation, FedKDD/FedMAS 2026 incorporates Multi-Agent Systems. Real-world federated deployments—such as smart cities or autonomous fleets—are not passive clients, but active, goal-driven agents capable of strategic behavior and dynamic collaboration. Integrating MAS concepts with FL and data mining enables robust, adaptive decentralized intelligence.

6 Related Workshops

The landscape of workshops addressing privacy-preserving, decentralized, and collaborative AI has expanded significantly in recent years. At KDD, the Workshop on Privacy-Preserving Machine Learning and Data Mining (PPML-DM) has explored secure computation techniques, while the Workshop on Federated Learning for Data Mining and Graph Analytics has touched on privacy in distributed settings¹. Beyond KDD, the International Workshop on Federated Learning with Generative AI In Conjunction with IJCAI 2025, focuses primarily on algorithmic and systems challenges in FL across AIGC, and healthcare domains². Similarly, AAMAS has hosted workshops on Multi-Agent Reinforcement Learning and Decentralized Coordination, emphasizing agent autonomy and strategic interaction but with limited integration of data mining or federated learning principles³. Notably, none of these forums explicitly bridge federated learning, multi-agent system theory, and large-scale data mining—the precise intersection that FedKDD/FedMAS 2026 aims to catalyze.

¹FedKDD: <https://kdd2024.kdd.org/workshops/>

²FedGenAI-IJCAI'25: <https://federated-learning.org/FedGenAI-ijcai-2025/>

³AAMAS 2025 Workshops: <https://aamas2025.org/index.php/conference/program/accepted-workshops/>