

Introduction to the Mini-track on Technological Advancements in Digital Collaboration with Generative AI and Large Language Models

Jie Tao
Fairfield University
jtao@fairfield.edu

Lina Zhou
The University of North Carolina at
Charlotte
lzhou8@uncc.edu

Philip Maymin
Fairfield University
pmaymin@fairfield.edu

The intellectual foundation of the Collaboration Systems and Technologies track rests on the premise that the value of a technology is realized within the context of a broader "collaboration system". The emergence of Large Language Model (LLM)-powered agentic systems represents a paradigm shift in this context, introducing AI not merely as a tool, but as a semi-autonomous actor within the system. This evolution necessitates a deeper inquiry into the architectural foundations required to support effective human-agent collaboration. This focus marks a deliberate thematic progression from last year's mini-track, which centered on techniques for improving direct human-model interaction. The rapid development of the field now compels a shift in focus from the interface to the architecture, raising new, fundamental questions about how to design and manage the entire systems of autonomous agents. Consequently, this mini-track addresses two central challenges that are critical for the future of collaborative AI: At the individual agent level, how do we establish the shared understanding necessary for a human to trust and effectively partner with an AI? And at the system level, how do we implement effective group governance to ensure the collective behavior of autonomous agents aligns with human goals?

This mini-track consists of two papers. They provide a diverse range of perspectives on leveraging generative AI and LLMs for digital collaboration. They examine and advance various techniques, address practical applications, and shed light on the challenges and opportunities in this rapidly evolving field. Specifically, they cover a rich set of themes in technological advancements in human-AI collaboration:

- Agentic workflow design for human-AI collaborations
- Advancements in AI reasoning models and explainability
- Collaborative frameworks for human-AI teaming
- Ethical considerations in the governance of AI systems

The papers in this mini-track offer complementary and insights into these themes. The first paper, "Agent Reasoning Tools (ARTs): A Novel Approach for Tool Definition to Empower LLM-based Agent Systems" by Jie Tao and Lina Zhou, investigates the challenge of shared understanding. The authors identify the rigidity and implementation overhead of conventional API-based agent tools as a key barrier to transparency. To address this, they introduce Agent Reasoning Tools (ARTs), a novel framework that defines an agent's capabilities using detailed, human-readable natural language docstrings, which are then directly interpreted by the LLM. This approach fosters a more interpretable, "glass-box" agent, enabling human collaborators to better understand the internal logic guiding an agent's actions and thereby facilitating a more robust foundation for human trust and co-creation.

The second paper, "LLM-Based Policy Generation for Distributed Adaptive Systems" by Marco Carvalho and Fitzroy Nembhard, addresses the challenge of group governance in complex, multi-agent environments. The authors address the significant difficulty of manually formulating and verifying the policies that govern agent behavior. They propose a novel framework that leverages pre-trained transformers to automatically generate formal, ontology-driven policies expressed in the Web Ontology Language (OWL). By enabling practitioners to manage system-wide constraints more intuitively, this work provides a scalable solution for human-led governance, simplifying the work of the "human in the loop" and ensuring that a collective of autonomous agents operates in a safe and predictable manner.

Taken together, these papers offer foundational solutions that address both the internal logic of an individual agent and the external governance of the collective. They present a compelling architectural vision for the next generation of human-agent collaboration. They elucidate that progress requires a dual focus on making individual agents more transparent at the micro-level and making agent collectives more governable at the macro-level.

Furthermore, this work surfaces several profound avenues for future inquiry that are critical to the Collaborative Systems and Technologies community.

A first avenue involves examining the tension between agent autonomy and system governance. Future research should investigate the dynamics that arise when an agent's flexible, interpretable reasoning leads to actions that conflict with pre-defined system policies, which may necessitate the development of more adaptive, negotiated governance models. A second critical area involves critically examining the cognitive ergonomics of AI oversight. As the human role evolves from operator to system governor, we must study the

novel cognitive loads this imposes to ensure the human overseer does not become a bottleneck to system performance. Finally, the development of these more socially sophisticated agents necessitates a new "computational organizational sociology" lens to understand the emergent social dynamics, such as trust, influence, and informal power structures, that will inevitably arise in these hybrid human-agent collectives. These profound questions underscore the rich, interdisciplinary challenges ahead and position this mini-track as a critical forum for shaping the future of collaboration systems.