Towards Understanding The Impact of Generative AI Agent Roles In Collaborative Problem-Solving Tasks

ANIRBAN MUKHOPADHYAY, KURT LUTHER, Virginia Tech, USA
KEVIN SALUBRE, SHASHANK MEHROTRA, HIFZA JAVED, TERUHISA MISU, KUMAR AKASH, Honda Research Institute, USA

Collaborative problem-solving under time pressure is common but difficult, as teams must generate ideas quickly, coordinate actions, and track progress. Generative AI offers new opportunities to assist, but we know little about how proactive agents affect the dynamics of real-time, co-located teamwork. We designed digital escape rooms as collaborative problem-solving tasks to study two generative AI agents: a Facilitator, which provided discussion summaries and proposed team structures, and a Peer, which contributed ideas as an imperfect teammate. We conducted a within-subjects user study with 24 participants, comparing team performance and processes across three conditions: No AI, Peer AI, and Facilitator AI. Preliminary results show that the Peer agent occasionally enhanced problem-solving by offering timely hints and memory support, though it also disrupted flow and created over-reliance. In comparison, the Facilitator agent provided light scaffolding but had limited impact on outcomes.

ACM Reference Format:
Anirban Mukhopadhyay, Kurt Luther and Kevin Salubre, Shashank Mehrotra, Hifza Javed, Teruhisa Misu, Kumar Akash. 2026. Towards Understanding The Impact of Generative AI Agent Roles In Collaborative Problem-Solving Tasks. 1, 1 (April 2026), 6 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 Introduction

Complex and time-sensitive problem-solving in the real world is rarely an individual task. From disaster response teams coordinating under time pressure to cybersecurity analysts mitigating an attack, group collaboration is the norm. Co-located collaboration, where team members work together in the same location and at the same time, often enhances productivity in tasks that rely on frequent communication and joint efforts, such as brainstorming, knowledge building, and planning [3, 17]. However, despite decades of CSCW and HCI research, developing technology that supports group processes and outcomes remains difficult.

Two recent research trends offer new possibilities for augmenting teamwork: generative AI and proactive systems. With advances in conversational and reasoning capabilities [19, 20], generative AI has become a powerful collaborator [5, 11, 22]. At the same time, these systems are prone to persuasive but flawed outputs [24], leaving groups vulnerable to over-reliance, anchoring on AI suggestions, or deferring to them to avoid social conflict [6].

In parallel, research has explored how AI and intelligent systems can act proactively [8, 12, 18]. Proactive behaviors have been shown to improve trust, situational awareness, and engagement across diverse contexts, from creativity to

Authors' Contact Information: Anirban Mukhopadhyay, Kurt Luther, Virginia Tech, Blacksburg, Virginia, USA; Kevin Salubre, Shashank Mehrotra, Hifza Javed, Teruhisa Misu, Kumar Akash, Honda Research Institute, San Jose, California, USA.

Manuscript submitted to ACM

decision support [25]. Together, these trends converge in the growing CSCW and HCI framing of human–AI teams (HATs), where AI agents are understood not only as tools but as team members [4, 6, 13, 15, 23].

While both generative AI and proactivity have shown promise independently, little is known about how they intersect in collaborative problem-solving. Should a proactive AI act as a *Peer*, contributing ideas as an imperfect teammate, or as a *Facilitator*, shaping group coordination and communication? How do these proactive roles influence not just team performance but also group processes? Addressing these questions is critical for designing AI that can effectively support collaboration in time-sensitive contexts. To explore this further, we investigate the following research questions in the context of co-located collaboration in time-sensitive problem-solving tasks:

RQ1: How do different AI agent roles (Peer vs. Facilitator) influence team performance?

RQ2: How do these AI agent roles shape team processes, such as workload, communication, and coordination?

In the following sections, we first describe the system design, including the task environment and the design of the peer and facilitator agents. We then detail our methods, outlining the within-subjects user study setup and procedures. Finally, we present preliminary qualitative findings that show how the two AI roles shaped team processes and performance.

2 System Design

Task Environment. We designed digital escape rooms to study teamwork under different AI conditions. Escape rooms are team-based recreational activities in which players work together to solve a series of puzzles or challenges in order to "escape" within a fixed time limit [14]. Traditionally, these are physical, in-person games set in themed environments (e.g., mystery, adventure, or crime-solving scenarios) that require critical thinking, coordination, and communication among team members. More recently, escape rooms have also been designed as digital environments for education and training, for example, to teach cybersecurity skills [7]. In our work, we created digital escape rooms to provide a controlled, engaging, and cognitively demanding environment for studying teamwork with and without AI agents.

Our puzzles were structured to last about 15–25 minutes, ensuring teams spent sufficient time collaborating and interacting with the AI features. Each puzzle contained three interconnected sub-puzzles, distributed across two screens to require active information exchange. The designs were inspired by cooperative online puzzle games such as Acorn Cottage [1] and Alone Together [2].

We tested puzzles against state-of-the-art multimodal and reasoning generative AI models (e.g., GPT-4.1, o3, o4) to evaluate model performance. While these models generated ideas that can lead to the solution by combining elements from the two screens, they failed to provide complete solutions.

Based on this task environment, we designed the two following generative AI powered agents that serve as functional technological probes [10].

Fiona — The Facilitator Agent. We designed the facilitator agent (Fiona) to scaffold groups' meta-cognitive processes during co-located problem-solving [21]. Drawing on prior work on facilitator agents, Fiona offered summaries of ongoing team discussions, time reminders, and prompts for workload division [16, 26]. The user interface is presented in Figure 1 (a) along with the features.

Ava — The Peer Agent. We designed Ava to contribute ideas without knowing the solution. The goal was to spark new directions and mimic peer-like collaboration rather than function as a facilitator or external advisor. This design builds upon a growing body of work demonstrating the value of AI as a creative collaborator [9].

Fig. 1. Figure shows (a) Screen 2 of Puzzle 1 with the Facilitator agent condition; and (b) Screen 1 of Puzzle 1 with the Peer agent condition. In Figure (a), there are two main features of the facilitator agent: (1) Fiona suggested structured collaboration strategies like the 1-2-4-All liberating structure [16], provided time reminders, and asked teams to divide up unsolved parts of the puzzle; (2) Fiona added the summary of ideas discussed by the team every 3 minutes. In Figure (b), there are three main features of the peer agent: (1) Ava proactively shared brainstorming thoughts every 3 minutes, based on puzzle screenshots and contextualized by team conversations; (2) Teams could follow up by asking Ava to explain or vary its ideas; (3) Ava was available as a chat-based partner on each puzzle screen, responding to questions.

To design Ava, we ran a formative study with six participants who each attempted a two-screen puzzle (same as one of the user study puzzles) with ChatGPT. Afterward, we interviewed them to gather what kind of help a peer agent should give, and when and where it would be most useful — assuming the AI doesn't know the solution. We present the details of the agent capabilities in Figure 1 (b).

3 Methods

We ran a within-subjects user study comprised of six groups, with four participants each. We recruited a total of 24 participants (19 men and 5 women) through voluntary convenience sampling from employees at our company.

Prior to the start of the study, we obtained informed consent from participants and administered the pre-study survey to collect basic demographic information. We then provided an overview of the study context, and the overall tasks. Based on the session condition, we presented a brief overview of the AI agent the participants will interact with before starting the puzzles.

The main experiment consisted of three sessions, each with a different puzzle and one of three AI conditions: No AI, Peer Agent, and Facilitator Agent, with a 20-minute time limit for each. The order of conditions was counterbalanced across teams using six Latin-square sequences to control for learning and fatigue effects. After each session, participants completed survey questionnaires about their experience. Once all sessions were finished, we conducted a 25–30 minute focus-group interview to gather feedback on the puzzle-solving process and participants' perceptions of the different AI agents.

4 Preliminary Results

4.1 RQ1: How do different AI agent roles (Peer vs. Facilitator) influence team performance in collaborative problem-solving tasks?

Facilitator agent had minimal impact on problem-solving success. Teams generally felt that the facilitator agent had little effect on puzzle performance. Participants said that the agent's summaries were often redundant, overly lengthy, and poorly timed, which under time pressure reduced their usefulness. P17 explained: *"When I saw the summarization, probably that was like couple minutes ago… by then I already knew which parts I should work on."* P7 described her experience as: *"I personally looked at it only once or twice… it was just reiterating what we were speaking out loud, so didn't really find it helpful."* Participants from teams 4 and 5 described the facilitator agent as having only a "bare minimum" impact on performance but still serving as a stabilizing element that provided focus and subtle structure.

Peer agent had widely varying perceptions. Participants saw the peer agent as a valuable but inconsistent teammate. At its best, it provided timely hints and reminders that directly shaped puzzle progress. As P6 explained, *"Most of the tasks we solved were hints given by the agent… it basically directed us towards correct answers."* Others highlighted its usefulness as a memory aid and calculator, helping recall details and reduce small errors under pressure.

Despite these benefits, the Peer's interventions also disrupted flow or created confusion. Like the Facilitator agent, the Peer's unsolicited suggestions could feel irrelevant, poorly timed, or too vague, leading to frustration. P18 said, *"Most of the time it was slightly distracting… we are solving a particular subtask, but this is suggesting something very different."* Some participants even admitted to over-relying on the AI, only to feel disappointed when its confident outputs proved unhelpful.

4.2 RQ2: How do these AI agent roles shape team processes—such as workload, communication, and coordination—during problem-solving?

Facilitator agent had little impact on how teams worked together. Participants were curious about the Facilitator at first, but quickly lost interest and largely ignored it. P23 put it as, *"Our communication was brainstorming, not strategic. The AI summaries didn't spark anything, they just restated what was already there."* Instead of reducing workload, its long summaries often felt like an added burden, while teamwork and task division continued to be managed by the group itself after being shaped early by the agent. Still, some appreciated that the AI stayed in the background, offering light scaffolding and optional reminders without disrupting communication.

Peer agent had a mixed influence on team processes. Participants found that the peer agent sometimes relieved workload by providing hints or performing low-level tasks, but interaction costs like typing prompts and interpreting long responses often offset these gains. Collaboration was affected as teams shifted from dividing tasks to more joint problem-solving to leverage proactive AI thoughts and responses to queries. However, unclear AI capabilities and parallel private chats occasionally introduced redundancy and confusion. The agent also reshaped communication patterns, often diverting attention away from teammates. Some participants admitted speaking less with the group because they were focused on interacting with the AI, for example, P11 mentioned, *"It definitely changes the dynamic… you have this fifth person that you interact with kind of alone."*

5 Conclusion

Our study presents an early but promising step toward understanding how proactive generative AI agents can enrich real-time, co-located teamwork. By comparing two distinct roles, a facilitator that provided summaries and team structure cues, and a peer that contributed ideas and memory support, we examined how proactive AI influences not only task performance but also group processes such as workload, coordination, and communication. Our results highlight that the value of proactive generative AI lies not in static roles but in the ability to adapt—providing the right kind of support at the right moment.

References

[1] 2025. Acorn Cottage. https://www.quarantini.space/ac-joining-instructions

[2] 2025. Enchambered Escape Room. https://www.enchambered.com/puzzles/alone-together/

[3] Christopher Andrews, Alex Endert, and Chris North. 2010. Space to think: large high-resolution displays for sensemaking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. Association for Computing Machinery, New York, NY, USA, 55–64. doi:10.1145/1753326.1753336

[4] Francisco Maria Calisto, João Fernandes, Margarida Morais, Carlos Santiago, João Maria Abrantes, Nuno Nunes, and Jacinto C. Nascimento. 2023. Assertiveness-based Agent Communication for a Personalized Medicine on Medical Imaging Diagnosis. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–20. doi:10.1145/3544548.3580682

[5] Chun-Wei Chiang, Zhuoran Lu, Zhuoyan Li, and Ming Yin. 2023. Are Two Heads Better Than One in AI-Assisted Decision Making? Comparing the Behavior and Performance of Groups and Individuals in Human-AI Collaborative Recidivism Risk Assessment. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3544548.3581015

[6] Chun-Wei Chiang, Zhuoran Lu, Zhuoyan Li, and Ming Yin. 2024. Enhancing AI-Assisted Group Decision Making through LLM-Powered Devil's Advocate. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*. ACM, Greenville SC USA, 103–119. doi:10.1145/3640543.3645199

[7] Tara N. Cohen, Andrew C. Griggs, Joseph R. Keebler, Elizabeth H. Lazzara, Shawn M. Doherty, Falisha F. Kanji, and Bruce L. Gewertz. 2020. Using Escape Rooms for Conducting Team Research: Understanding Development, Considerations, and Challenges. *Simulation & Gaming* 51, 4 (Aug. 2020), 443–460. doi:10.1177/1046878120907943

[8] Stephanie Houde, Kristina Brimijoin, Michael Muller, Steven I. Ross, Dario Andres Silva Moran, Gabriel Enrique Gonzalez, Siya Kunde, Morgan A. Foreman, and Justin D. Weisz. 2025. Controlling AI Agent Participation in Group Conversations: A Human-Centered Approach. In *Proceedings of the 30th International Conference on Intelligent User Interfaces*. 390–408. doi:10.1145/3708359.3712089 arXiv:2501.17258 [cs].

[9] Kent F. Hubert, Kim N. Awa, and Darya L. Zabelina. 2024. The current state of artificial intelligence generative language models is more creative than humans on divergent thinking tasks. *Scientific Reports* 14, 1 (Feb. 2024), 3440. doi:10.1038/s41598-024-53303-w

[10] Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B Bederson, Allison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, et al. 2003. Technology probes: inspiring design for and with families. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 17–24.

[11] Angel Hsing-Chi Hwang, John Oliver Siy, Renee Shelby, and Alison Lentz. 2024. In whose voice?: examining AI agent representation of people in social interaction through generative speech. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*. 224–245.

[12] Janet G. Johnson, Macarena Peralta, Mansanjam Kaur, Ruijie Sophia Huang, Sheng Zhao, Ruijia Guan, Shwetha Rajaram, and Michael Nebeling. 2025. Exploring Collaborative GenAI Agents in Synchronous Group Settings: Eliciting Team Perceptions and Design Considerations for the Future of Work. doi:10.48550/arXiv.2504.14779 arXiv:2504.14779 [cs].

[13] Patricia K. Kahr, Gerrit Rooks, Chris Snijders, and Martijn C. Willemsen. 2024. The Trust Recovery Journey. The Effect of Timing of Errors on the Willingness to Follow AI Advice.. In *Proceedings of the 29th International Conference on Intelligent User Interfaces (IUI '24)*. Association for Computing Machinery, New York, NY, USA, 609–622. doi:10.1145/3640543.3645167

[14] Erica Kleinman and Casper Harteveld. 2024. The Untapped Potential of Escape Rooms as Gamified Research Environments. In *Companion Proceedings of the 2024 Annual Symposium on Computer-Human Interaction in Play (CHI PLAY Companion '24)*. Association for Computing Machinery, New York, NY, USA, 276–278. doi:10.1145/3665463.3678865

[15] Shuai Ma, Xinru Wang, Ying Lei, Chuhan Shi, Ming Yin, and Xiaojuan Ma. 2024. "Are You Really Sure?" Understanding the Effects of Human Self-Confidence Calibration in AI-Assisted Decision Making. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–20. doi:10.1145/3613904.3642671

[16] Keith McCandless. 2020. Liberating Structures: Change Methods for Everybody Every Day. https://keithmccandless.medium.com/liberating-structures-change-methods-for-everybody-every-day-648e9c0d04a7

[17] Gary M. Olson and Judith S. Olson. 2000. Distance Matters. *Human–Computer Interaction* 15, 2-3 (Sept. 2000), 139–178. doi:10.1207/S15327051HCI1523_4 Publisher: Taylor & Francis _eprint: https://doi.org/10.1207/S15327051HCI1523_4.

[18] Mohammad Amin Samadi, Spencer JaQuay, Jing Gu, and Nia Nixon. 2024. The AI collaborator: Bridging human-AI interaction in educational and professional settings. *arXiv preprint arXiv:2405.10460* (2024).

[19] William Seymour and Emilee Rader. 2024. Speculating About Multi-user Conversational Interfaces and LLMs: What If Chatting Wasn't So Lonely?. In *Proceedings of the 6th ACM Conference on Conversational User Interfaces (CUI '24)*. Association for Computing Machinery, New York, NY, USA, 1–4. doi:10.1145/3640794.3665888

[20] Yang Shi, Tian Gao, Xiaohan Jiao, and Nan Cao. 2023. Understanding Design Collaboration Between Designers and Artificial Intelligence: A Systematic Literature Review. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW2 (Oct. 2023), 368:1–368:35. doi:10.1145/3610217

[21] Leigh Thompson and Taya R. Cohen. 2012. Metacognition in Teams and Organizations. In *Social Metacognition*. Psychology Press. Num Pages: 20.

[22] Simone van den Broek, Supraja Sankaran, Jan de Wit, and Alwin de Rooij. 2024. Exploring the Supportive Role of Artificial Intelligence in Participatory Design: A Systematic Review. In *Proceedings of the Participatory Design Conference 2024: Exploratory Papers and Workshops - Volume 2*

(PDC '24, Vol. 2). Association for Computing Machinery, New York, NY, USA, 37–44. doi:10.1145/3661455.3669868

[23] Mengyao Wang, Jiayun Wu, Shuai Ma, Nuo Li, Peng Zhang, Ning Gu, and Tun Lu. 2025. Adaptive Human-Agent Teaming: A Review of Empirical Studies from the Process Dynamics Perspective. doi:10.48550/arXiv.2504.10918 arXiv:2504.10918 [cs].

[24] Désirée Zercher, Ekaterina Jussupow, and Armin Heinzl. 2023. When AI joins the team: a literature review on intragroup processes and their effect on team performance in team-AI collaboration. (2023).

[25] Rui Zhang, Wen Duan, Christopher Flathmann, Nathan McNeese, Guo Freeman, and Alyssa Williams. 2023. Investigating AI Teammate Communication Strategies and Their Impact in Human-AI Teams for Effective Teamwork. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW2 (Sept. 2023), 1–31. doi:10.1145/3610072

[26] Zheng Zhang, Weirui Peng, Xinyue Chen, Luke Cao, and Toby Jia-Jun Li. 2025. LADICA: A Large Shared Display Interface for Generative AI Cognitive Assistance in Co-located Team Collaboration. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–22. doi:10.1145/3706598.3713289