There is no (A)I in Team? Implications of AI Agents for Group Decision-Making.

PATRICIA KAHR, University of Zurich, Switzerland

Group decision-making (GDM) is a well-studied phenomenon, with decades of research exploring how groups can both enhance and undermine collective judgment. As agentic AI systems increasingly enter collaborative workflows, they become more than tools: they act as structural participants that shape how groups deliberate, converge, and assign responsibility. Translating this rich body of knowledge to settings where AI agents act as group participants represents a timely and open research challenge. Drawing on established GDM theory, this position paper maps key opportunities and pitfalls of integrating AI agents into group decision-making processes. I argue that while AI agents can improve participation equity, informational diversity, and process facilitation, the same capabilities risk accelerating groupthink, diffusing accountability, and fostering long-term skill atrophy in human decision-makers. I conclude with four research questions intended to inspire discussion at the AutomationXP26 workshop on how future work can support empowering, effective, and ethically sound collaborations in group settings with AI agents.

CCS Concepts: • **Human-centered computing → Human computer interaction (HCI)**; **Collaborative and social computing**; • **Computing methodologies → Artificial intelligence**; • **Social and professional topics → Computing profession**.

Additional Key Words and Phrases: AI-assisted group decision-making, agentic AI, automation bias, human-AI collaboration

ACM Reference Format:

Patricia Kahr. 2026. There is no (A)I in Team? Implications of AI Agents for Group Decision-Making.. In *Proceedings of AutomationXP26 Workshop of the 2026 CHI Conference on Human Factors in Computing Systems (AutomationXP26).* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 Introduction

Although group decision-making (GDM) is rarely effortless, combining the efforts of individuals with different experiences, skills, or perspectives can often lead to more creative or robust solutions than individuals working alone – not only through shared wisdom but also through mutual checking and accountability. For these reasons, AI support in GDM is increasingly relevant, widening the research focus from dyadic human-AI interactions to collective decision-making contexts in which AI agents are part of the workflow.

Research on human GDM in sociology and psychology has developed over decades and offers rich foundations for understanding these settings. Prior work has examined how groups conform [1], how group dynamics can undermine judgment [11], how information cascades shape collective beliefs [2], and how group interaction patterns relate to collective performance [16]. Agentic AI systems, however, differ fundamentally from the human participants this literature was built around. Rather than executing bounded tasks, they may coordinate multiple sub-agents, initiate actions autonomously, and increasingly be configured to take on flexible roles within specific teams and contexts — raising the prospect that AI may soon participate in committees, panels, or multi-stakeholder decision processes as a structural member rather than a background tool.

Author's Contact Information: Patricia Kahr, patricia.kahr@uzh.ch, University of Zurich, Department of Informatics, Zurich, Switzerland.

Manuscript submitted to ACM

This distinction matters because it means established GDM theory cannot simply be carried forward unchanged. When an AI agent shapes who speaks, what information is surfaced, or how quickly a group converges, it intervenes in the very mechanisms (such as conformity pressure, information sharing, or accountability) that decades of GDM research have identified as critical. Whether AI participation amplifies or disrupts these mechanisms, and under what conditions, is an open question. Early empirical work suggests the picture is mixed: AI assistance can alter group dynamics in ways that differ meaningfully from individual AI-assisted decision-making [3], but the broader sociotechnical implications remain underexplored.

This position paper briefly explores considerations for integrating AI agents into GDM, mapping both opportunities and pitfalls, and offers four research questions as a foundation for discussion at the AutomationXP26 workshop.

2 Group Decision Making: Wisdom or Defection of Crowds

Decades of work in organizational behavior, social psychology, and cognitive science have explored the conditions under which groups make better or worse decisions than individuals. On the one hand, collective processes can enhance decision quality by improving deliberation and aggregating diverse perspectives [9]. On the other hand, groups are susceptible to well-documented failure modes, such as biased information sharing, in which members disproportionately discuss information already known to all rather than surfacing unique knowledge [15], which can systematically skew group judgments even when better information is available within the group.

Two structural dimensions are particularly relevant when considering how AI agents enter these dynamics. The first is group structure, which refers to the composition of a group, the roles that individuals in a group take on in the decision-making scenario, as well as their hierarchical standing, the diversity, and expertise that shape group dynamics. The second is group procedures, which refers to the (explicit and implicit) behavioral rules, communication strategies, and information-sharing patterns that govern how decisions unfold in practice. AI agents can intervene at both levels: at the structural level by taking on specific roles and thereby influencing participation dynamics and group composition. At the procedural level, AI agents are implemented in GDM settings to actively shape how information is introduced, framed, and weighted during group deliberation.

The following two sections examine the implications of such interventions, first as potential opportunities and then as potential risks.

2.1 AI Agents as a Boost for GDM

In many group decision scenarios, vocal minorities or dominant personalities can create unbalanced group dynamics that hinder optimal decision-making. Introverted individuals or those lower in the hierarchy may withhold concerns or defer to more assertive voices – a phenomenon sometimes described as HiPPO effects (Highest Paid Person's Opinion). AI could promote more equitable participation by encouraging contributions from quieter group members or balancing participation through turn-taking strategies to prevent dominance effects before they take hold.

A related disadvantage of groups is strong alignment, one-sidedness, or gaps in the information available for a given decision. In such scenarios, AI could act as a devil's advocate: representing a missing perspective, surfacing neglected data, offering alternative framings, or contributing cross-domain insights that expand the informational diversity of the discussion. This kind of structured dissent may improve reasoning and decision quality [13] and problem-solving through cognitive diversity [10].

Finally, not all groups are equally capable of creating effective decision-making environments. Unresolved group dynamics, the absence of a neutral moderator, or an unproductive discussion culture could undermine decision outcomes.

As a facilitator, AI agents could take on a neutral coordinating role, supporting systematic comparison of alternatives, counterfactual reasoning, or documentation of trade-offs, thereby improving the procedural quality of deliberation and, ultimately, group performance [5, 16].

2.2 AI Agents as a Liability for GDM

While the capabilities described above suggest AI can benefit GDM, the same capabilities of AI may also introduce structural risks. Efficiency gains, such as faster alignment toward a common goal, come with a corresponding risk: AI outputs may anchor group discussions and shorten deliberation cycles, leading groups to converge prematurely rather than adequately explore the decision space. Early compliance without critical engagement is well-documented as automation bias [14] and, at the group level, may increase the likelihood of groupthink [4, 11]. Efficient convergence can simultaneously suppress minority perspectives, whereas maintaining diverse viewpoints typically improves decision quality [13].

A second risk is that AI agents may be instrumentalized by group members to legitimize pre-existing opinions rather than genuinely inform deliberation. Although AI is often perceived as a neutral participant, it can be selectively invoked to reinforce particular positions [7]. A connected consequence is the diffusion of decision responsibility: when AI input is woven into the deliberative process, individuals may struggle to locate accountability. Hence, decisions appear to emerge from a human-AI collective rather than identifiable agents, potentially weakening both individual accountability and the group's sense of ownership over outcomes.

Finally, beyond immediate effects on control and responsibility, repeated AI-assisted decision-making may cause longer-term skill atrophy. What appears in the moment as reduced cognitive load may in practice lead to disengagement: approving AI output without meaningful evaluation, a pattern sometimes described as rubber-stamping [6, 8]. In long-term or recurring decision-making contexts, systematic cognitive offloading of this kind may therefore risk degrading collective evaluative capacities over time [12].

3 Research Questions to Inspire Workshop Discussions

The opportunities and challenges outlined above raise interconnected questions about the role AI agents should play in GDM, the conditions under which their involvement helps or harms, and how these dynamics unfold when AI-supported GDM becomes a repeated or long-term practice. The following four research questions are intended to inspire discussion at the workshop:

- **RQ1:** How does the specific role that AI takes on (e.g., participation equalizer, devil's advocate, neutral facilitator) alter GDM outcomes at the group level (e.g., argument diversity, decision quality, group cohesion) and at the individual level (e.g., perceived fairness, autonomy, and efficacy)? And how do these effects shift as groups interact with AI agents over longer periods?
- **RQ2:** How does responsibility attribution evolve in repeated or long-term AI-supported GDM, and what design strategies can mitigate the dilution of decision responsibility and accountability and the erosion of individual decision efficacy?
- **RQ3:** What design interventions can preserve human agency and critical engagement while still leveraging AI's coordination and facilitation capabilities in GDM settings?
- **RQ4:** What governance and design guardrails are needed to ensure that AI participation in GDM does not undermine group accountability structures, suppress minority voices, or erode the legitimacy of group decisions?

4 Conclusion

As AI agents increasingly enter collaborative workflows, the question is no longer whether they will participate in group decision-making, but how. The underlying position paper has argued that established GDM theory offers a valuable but insufficient foundation for understanding these settings, and that translating it into contexts where AI agents should act as effective and accountable team members requires both empirical investigation and deliberate design. The opportunities are considerable: AI agents can promote participation equity, expand informational diversity, and support procedural quality in groups that lack the structure or roles to do so themselves. But so are the risks: premature convergence, diffused accountability, and long-term skill atrophy are not hypothetical concerns but extensions of well-documented group dynamics into new territory. The four research questions proposed here are intended to initiate discussions at the AutomationXP26 workshop, with the broader goal of informing human-AI collaboration in GDM that is empowering, effective, and ethically sound.

5 AI Use Declaration

The preparation of this manuscript involved the use of digital AI tools to support the research and writing process. Google Scholar Labs and Ai2 Asta were used to assist in streamlining and cross-checking literature. Grammarly and Claude were used to improve grammatical accuracy and conciseness. The research concept, all arguments, interpretations, and editorial decisions remain my own intellectual work and responsibility.

References

[1] Solomon E. Asch. 1951. Effects of Group Pressure upon the Modification and Distortion of Judgments. (1951), 177–190.

[2] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. 1992. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy* 100, 5 (1992), 992–1026. doi:10.1086/261849

[3] 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*. ACM, New York, NY, USA. doi:10.1145/3544548.3581015

[4] Chun-Wei Chiang, Zhuoran Lu, Zhuoyan Li, and Ming Yin. 2025. The Amplifying Effect of Explainability in AI-assisted Decision-making in Groups. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA. doi:10.1145/3706598.3713534

[5] Gerardine DeSanctis and R. Brent Gallupe. 1987. A Foundation for the Study of Group Decision Support Systems. *Management Science* 33, 5 (1987), 589–609. doi:10.1287/mnsc.33.5.589

[6] M. C. Elish. 2019. Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction. *Engaging Science, Technology, and Society* 5 (2019), 40–60. doi:10.17351/ests2019.260

[7] Tarleton Gillespie. 2014. *The Relevance of Algorithms.* MIT Press, Cambridge, MA. 167–194 pages.

[8] Kate Goddard, Abdul Roudsari, and Jeremy C. Wyatt. 2012. Automation Bias: A Systematic Review of Frequency, Effect Mediators, and Mitigators. *Journal of the American Medical Informatics Association* 19, 1 (2012), 121–127. doi:10.1136/amiajnl-2011-000089

[9] Jacob K. Goeree and Leeat Yariv. 2011. An Experimental Study of Collective Deliberation. *Econometrica* 79, 3 (2011), 893–921. doi:10.3982/ECTA8852

[10] Lu Hong and Scott E. Page. 2004. Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers. *Proceedings of the National Academy of Sciences* 101, 46 (2004), 16385–16389. doi:10.1073/pnas.0403723101

[11] Irving L. Janis. 1972. *Victims of Groupthink: A Psychological Study of Foreign-Policy Decisions and Fiascoes.* Houghton Mifflin, Boston, MA.

[12] Weng Marc Lim et al. 2024. Does Using Artificial Intelligence Assistance Accelerate Skill Decay and Hinder Skill Development Without Performers' Awareness? *European Journal of Education* (2024). doi:10.1111/ejed.12610 PMC11239631.

[13] Charlan Jeanne Nemeth. 1986. Differential Contributions of Majority and Minority Influence. *Psychological Review* 93, 1 (1986), 23–32. doi:10.1037/0033-295X.93.1.23

[14] Raja Parasuraman and Victor Riley. 1997. Humans and Automation: Use, Misuse, Disuse, Abuse. *Human Factors* 39, 2 (1997), 230–253. doi:10.1518/001872097778543886

[15] Garold Stasser and William Titus. 1985. Pooling of Unshared Information in Group Decision Making: Biased Information Sampling During Discussion. *Journal of Personality and Social Psychology* 48, 6 (1985), 1467–1478. doi:10.1037/0022-3514.48.6.1467

[16] Anita Williams Woolley, Christopher F. Chabris, Alex Pentland, Nada Hashmi, and Thomas W. Malone. 2010. Evidence for a Collective Intelligence Factor in the Performance of Human Groups. *Science* 330, 6004 (2010), 686–688. doi:10.1126/science.1193147