LEX LOCALIS-JOURNAL OF LOCAL SELF-GOVERNMENT
ISSN:1581-5374 E-ISSN:1855-363X
VOL. 24, NO. 1(2026)

AI Agents in Operational Managerial Decision-Making: Opportunities, Risks, and Human Oversight in Modern Organizations

Dr Patryk Chmielarz

University of the National Education Commission, Krakow, Poland

ORCID: 0000-0003-0286-741X

e-mail: chmielarzp13@gmail.com

Abstract

Artificial intelligence (AI) agents are another recent advancement in organizational management, where AI systems are able to process information, plan tasks, engage with digital products and tools and carry out operational tasks in varying degrees of autonomy. While traditional AI systems are more often used to make predictions or recommendations, AI agents can be integrated into the day-to-day operations of an organization. This paper delves into the opportunities and challenges that arise with the application of AI agents in operational managerial decision-making, specifically regarding inventory management, procurement, and workflow coordination. The study is based on the bounded rationality theory, the human–AI collaboration and automation–augmentation view and uses a structured narrative review method for synthesizing the literature found in the field. Four major findings emerge from the thematic analysis: opportunities for operational efficiency, issues of decision reliability, ongoing need for human managerial control and organizational governance's influence on the level of agent

autonomy. The results indicate that AI agents could enhance managerial decision-making to minimize repetitive tasks and streamline information processing, provided they operate within the framework of well-defined authority and organizational controls while the quality of the data they access is high. The paper makes several contributions to management scholarship, including highlighting the difference between AI-supported decision making and agent-driven decision making, and suggesting bounded autonomy as a way to responsibly adopt AI in organizations.

Keywords: AI agents, operational management, managerial decision-making, human oversight, procurement, inventory management, agentic AI

1. Introduction and Background

Artificial intelligence is one of the pivotal technological innovations that impact organizational management and decision-making. AI applications are becoming a critical tool for organizations to manage information, detect patterns, forecast operating demands, and help managers with their tasks. The developments have sparked interest in the possibility of making intelligent technologies useful in improving the efficiency of organizations and the quality of their decisions. The development of AI applications, however, brings up questions about the connection between human managerial judgment and organizational accountability (Jarrahi, 2018, Raisch & Krakowski, 2021).

Recent progress in agentic AI has enabled the development of AI systems capable of performing complex, interconnected tasks rather than just producing stand-alone analytical results.Progress in agentic AI has brought forth systems that can perform interconnected tasks, not limited to producing stand-alone analytical outputs. AI agents can execute activities towards specified objectives using computational reasoning, planning mechanisms, memory and external tools. Wang et al. (2024) describe how autonomous agents that are based on large language models (LLMs) are typically composed of planning, memory, and action-related units that allow them to interact with digital environments. Analogously, Wang et al. (2024) show reasoning and action can be combined into language-models based systems that can be used to retrieve information and interact with external tools.

These are important to operational managerial decision-making because companies have to deal with inventory control, procurement needs, supplier data, and departmental operations regularly. There are many times when decisions have to be made on a large amount of structured data, reoccurring administration tasks and coordination across various functions within organisations. Intelligent agents can help with these processes, gathering information, assessing alternatives, detecting inconsistencies, or planning or performing pre-defined actions.

Practical experience in organisations shows that AI tools are more effectively implemented in specific business processes than in blanket automation. Organizations' AI capabilities, as outlined by Benbya et al. (2021), include automation, engagement, decision insights, and innovation, providing a valuable reference when choosing bounded operational applications to explore. Their analysis is also from an earlier wave of tools that use AI to assist with the procurement process, which does not enable them to verify the reliability of autonomous procurement or inventory execution.

In the case of an organization with multiple warehouses, the AI agent could track inventory levels, alert whoever is responsible for the warehouse to restock when supplies are running low, and compare the amount of inventory available at each location to suggest replenishment tasks. Likewise, an AI-based procurement system can compare supplier quotes, evaluate the purchasing needs, and draft purchase requests for management approval. These kinds of activities illustrate the possibilities of agent-based technologies to assist operational coordination and to eliminate repetitive administrative tasks. However, these applications are still very much reliant on accurate organizational information and well-defined operating conditions (Parasuraman et al., 2000; Wang et al., 2024).

To grasp the potential benefits of AI agents, it is essential to understand the concept of bounded rationality. Limited information, time and computational power are management decisionmakers faced with limitations, according to Simon (1979). This means that managers can't check out all the options, and often have to make decisions without knowing what the future will bring. AI agents can potentially mitigate some of these drawbacks by handling large amounts of information and optimizing the consideration of decision options.

But because of better computer power, managerial uncertainty is not necessarily gone. Contextual information, informal knowledge, ethical considerations, and unforeseen events are among the many factors that can influence an organizational decision and that cannot necessarily be captured effectively in a digital system. Jarrahi (2018) suggests that AI and human decision makers have complementary strengths, and that the use of technology is a part of a multi-faceted decision-making process.

Additionally, the use of AI agents raises issues of self-directed decision execution. Unlike traditional decision-support systems, agents can also have the technical capabilities to directly interact with organizational applications and take action without having to get approval at each point. Questions are raised about the delegation of authority, accountability for mistakes and boundaries of technological autonomy. Increasingly agentic information systems demands more focus on task and authority sharing and delegating between human actors and technological artifacts as highlighted by Baird and Maruping (2021).

As other countries are adopting AI tools and solutions, a critical issue arises about the amount of operational managerial decisions that should be handed over to AI agents. The difference between recommending and making decisions is particularly relevant because if their autonomous actions were inappropriate, they could have financial, operational and/or organizational implications before managers can have any chance of intervening (Parasuraman et al., 2000; Floridi et al., 2018). This study is therefore limited to the specific aspects of Inventory Management, procurement and workflow coordination.

1.1 Research Problem

While there are discussions on the use of artificial intelligence in organizational decision-making, human–AI collaboration and algorithmic management, the new operational capability of AI agents raises new concerns about reliability and managerial authority (Baird & Maruping, 2021; Jarrahi, 2018). The core research question is how to ensure that AI agents can assist decision making in operations, without leading to unacceptable risks arising from inaccurate information, inappropriate actions, or reduced organizational accountability.

1.2 Research Objectives

The study pursues two objectives: (1) to examine opportunities and operational risks associated with AI agents in inventory management, procurement, and workflow coordination; and (2) to identify how human managerial oversight and organizational governance can support responsible agent use.

1.3 Research Questions

RQ1: What opportunities and operational risks arise from using AI agents in inventory management, procurement, and workflow coordination?

RQ2: How can human managerial oversight support the responsible use of AI agents in these operational decisions?

2. Literature Review

2.1 AI Agents in Operational Decision-Making

AI agents are computational systems that can perform certain tasks aimed at achieving specific goals by means of information processing, planning, and interacting with digital environments. Large language models can be integrated with memory, reasoning mechanisms, external applications, and action execution (Wang et al., 2024), to form contemporary agent architectures. The field of technological development of autonomous agents is an extension of classical applications of the Artificial Intelligence. While conventional analytical systems typically provide information and/or recommendations, agentic systems can execute series of activities that involve information retrieval, planning, and interaction with the external tools. According to Benbya et al. (2021), there are three types of organizational applications of AI: automation, insight, and decision-making. These features can be used in a business context for managing inventory, researching purchase options and managing administrative workflows. An agent can, for instance, gather stock data, determine stock replenishment needs, and create purchasing paperwork. Its accuracy and usefulness, however, are dependent on the quality of organizational records, and the appropriateness of the agent's operating instructions. Decision structures in the organization can be made on the basis of the nature of a task. Shrestha et al. (2019) classify the forms of human delegation as human-to-AI, sequential human-to-AI and aggregative human-to-AI. For operational managers, this typology means that the delegation of duties for quotation comparison or anomaly detection can be different from the authorisation/purchasing of the supplier. Furthermore, Mikalef and Gupta (2021) conclude that organizational AI capability is dependent on technological, human and intangible resources, which explains the need for more than just an agent to be effective in the operation. Evaluation and reliability also remain challenges in the study of autonomous agents as noted by Wang et al. (2024). This indicates that technical competence is not sufficient for appropriate autonomous deployment in organisational

context. Additionally, managers need to reflect on whether the agent can be relied upon under real-world factors and if the agent's actions align with the goals of the organisation.

2.2 Bounded Rationality and Managerial Decision-Making

The concept of bounded rationality offers a valuable lens to examine how AI agents can aid organizations. Simon (1979) states that decision makers have constraints with regard to the amount of information they can process, the amount of time they have available, and the ability to effectively compare options. These limitations are felt by operational managers when they are in charge of coordinating inventory needs, supplier relations, procurement expenses, and work processes. Often multiple factors need to be taken into account when making these decisions. AI agents can help managers compile information and provide them with other possible answers. For instance, an inventory-management agent might check sales records, current inventory and restocking plans, and track products that are in need of managerial attention. However, having computational assistance does not ensure proper decisions. AI agents rely on data from organisational databases, and may not possess a complete understanding of exceptional situations. The notion of bounded rationality (Simon, 1979) is still relevant, as uncertainty can never be fully overcome by the improvement of computational processing. Jarrahi (2018) also highlights the complementary role of intuition and context interpretation by human decision-makers to analytical AI.

2.3 Human–AI Collaboration and Managerial Authority

Human-AI collaboration is the term used for systems in which humans and AI work together to perform organizational work. According to Dellermann et al. (2019), hybrid intelligence is the

utilization of human and machine intelligence to work together in the learning and problem-solving process. This is important since managers and AI agents have different strengths. While agents gather standardized information and repetitive computational tasks, human managers provide organizational experience and contextual knowledge and judgment about exceptional conditions (Dellermann et al., 2019; Jarrahi, 2018). Trust calibration also determines the appropriateness of human supervision. Glikson and Woolley (2020) review the empirical studies that show that trust in AI can be shaped by representations of the AI systems and their perceived abilities.

At the same time, Dietvorst et al. (2015) show that algorithm aversion occurs following observed errors, while Logg et al. (2019) present evidence of algorithm appreciation in a number of contexts in which advice is taken. The conflicting results suggest that managers could either ignore beneficial suggestions from their agents or even take them for granted, thus, it becomes particularly significant to verify their agents' suggestions and have calibrated trust. This relationship can be understood with the help of the automation/augmentation paradox as described by Raisch and Krakowski (2021). They believe that automation and augmentation are not mutually exclusive processes, but rather are interconnected. Some managerial functions may be automated, allowing managers to take on more complex duties, but too much automation can cause human capability and organizational control issues. When AI agents take action, the question of delegation of authority is crucial. Baird and Maruping (2021) refer to agentic information-system artifacts as ones that lead to evolving relationships between human actors and technologic systems, based on task delegation. In a workflow, there can be varying degrees of technological autonomy, according to Parasuraman et al. (2000), who identify four types of

automation: information acquisition, information analysis, decision selection, and action implementation

2.4 Risks Associated With Agent-Based Decisions

AI agents can pose issues regarding information quality, access to the systems, accountability, and managerial reliance. Poor inventory data, missing or incorrect buying data, or obsolete workflow instructions can cause agents to make inappropriate recommendations or steps.

Another issue is the implications of self-determination execution. The model proposed by Parasuraman et al. (2000) illustrates the importance of examining the potential consequences of delegating the decision-selection and action-implementation functions to technological systems when designing the automation.

It is also important to have organizational accountability. Algorithmic technologies shape the way organizational activities are organised, assessed and managed, as explained by Kellogg et al. (2020). While not exclusive to modern-day AI agents, their study provides a relevant insight into shifts in managerial power and organizational dynamics.

There are also concerns about data privacy and cybersecurity as AI agents may need to access sensitive organizational data. The Responsible AI scholarship focuses on the challenges of recognising, assessing and managing AI risks across the lifecycle of the technology (Floridi et al., 2018). Explainability is also a key concern in the field of AI, as demonstrated by research on the topic (Shin, 2021).

Usable interfaces and documented accountability also provide good oversight. Shin (2021) studies the role of explainability in user trust & acceptance. Organizational implications of learning algorithms and evolving boundaries of expertise are studied by Faraj et al. (2018). These contributions help to maintain records of agent actions, analyze workflow exceptions, and give managers a way to meaningfully intervene.

.5 Research Gap

While the reviewed literature provides recognised viewpoints on decision-making with the support of AI, intelligent agents, organisational automation and governance, these viewpoints are generally independent from one another. Technical research is mainly about the agent architecture and capability (Wang et al., 2024; Wang et al., 2024), while management research is about human judgment, authority and consequences of automation (Baird & Maruping, 2021; Raisch & Krakowski, 2021).

This paper merges these views in the ongoing operational managerial decisions including those cases where agents move from the role of analytical backup to autonomous operation. It's all about the connection between operational efficiency, decision reliability, delegated authority, and human oversight.

3. Methodology

3.1 Research Design

The current research utilizes a structured narrative review of literature to explore the opportunities and obstacles of AI agents in the process of operational managerial decision making. A narrative review makes it possible to bring together theoretical and interdisciplinary approaches that are relevant to a new research phenomenon. It is based on bounded rationality, human–AI collaboration, automation, organizational control, and AI governance (Baird & Maruping, 2021; Jarrahi, 2018; Simon, 1979).

The review is not a primary survey, organizational experiments or statistical analysis. Its findings are a thematic analysis of the existing scholarship and not empirical findings.

Students are expected to select literature and to explore the scope of their work.

The review includes scholarly publications on AI agents, managerial decision making, autonomous systems, human–AI collaboration, and responsible AI governance. Sources may be peer reviewed journal articles, conference research, and institutional risk-management guidance. The retention of foundation studies is based on explanations of managerial concepts that are enduring, while recent studies interpret agent capabilities.

The synthesis also incorporates more general evidence on the performance of organizations with AI, but excludes direct evidence of autonomous agents. According to Noy and Zhang (2023), generative AI can boost productivity in professional applications of text generation in writing, while Brynjolfsson et al. (2025) cite improvements in productivity among customer-support workers using an AI assistant to complete text-generation tasks. These studies are not about AI-driven inventory or procurement agents themselves, so they provide insights into potential benefits but do not conclusively show AI-driven agentic operational effectiveness.

The literature is limited in scope to three functions of the operations: inventory management, procurement, and coordination of the operations workflow. This is not a systematic paper, a systematic database search or a bibliometric analysis, but rather a structured conceptual narrative review.

3.3 Thematic Analysis

The literature has been organized based on the two research questions, thematic synthesis. Five starting points – operational efficiency, decision reliability, delegated authority, human oversight, and organizational governance – are merged into four themes: efficiency and coordination, autonomous execution risks, human oversight and governance requirements. The results presented in this study can be seen as theory-driven patterns and implications that need to be tested empirically.

4. Results

4.2 Learning Through Practice

The first theme relates to the possibility of AI agents to enhance information processing and coordination in the operational activities. Wang et al. (2024) suggest that planning, memory, and action capabilities are key components in autonomous-agent architectures, and Benbya et al. (2021) bring forward the functions of automation and AI-supported organizational decision processes. These capabilities create a technical infrastructure for systems that are managing multiple operational tasks.

For inventory management, agents can fetch inventory details, identify inventory needs, and prepare data for management analysis. They can assist in quotation comparisons, documentation examinations, and purchasing procedures in procurement. They can help in coordinating tasks, sharing information, and sending reminders about unfinished tasks within the organization.

Some of the opportunities can be seen from the bounded rationality point of view; structuring the relevant information could remove some of the cognitive load of managers (Simon, 1979). However, the literature surveyed does not support an absolute improvement in operational efficiency for all of these functions as a result of modern AI agents. The finding, therefore, points to potential organizational opportunities and not to actual causal effects.

This is important when reading productivity evidence. Noy and Zhang (2023) explore AI-assisted writing, and Brynjolfsson et al. (2025) look at AI-assisted customer support; neither leaves the execution of fully autonomous agents completely untouched in the inventory/ procurement domain. The appropriate take-home message is that, while AI can support particular work processes, there is no reason to believe that delegating decisions will necessarily be better without precautions.

Risks associated with 4.2 Decision Reliability and Autonomous Execution:

The second theme is about the errors or omissions in the information given to AI agents for their operational decisions. If the information is incomplete or inconsistent, or inappropriate, AI systems can generate inaccurate outputs. Evaluation and reliability are important challenges in the development of autonomous agents according to Wang et al, (2024).

If inventory records are wrong, an inventory agent can suggest the wrong amounts. A procurement agent might focus on price prices without considering the dependability of the delivery or contractual obligations. These issues become critical when suggestions become implemented decisions. Parasuraman et al. (2000) separate out information analysis, decision selection, and implementation; managers might catch an incorrect recommendation prior to taking action, while autonomous execution can have immediate consequences.

The level of autonomy that can be granted to an agent should therefore be affected by the level of reliability of the decision. especially if the activity involves investment or situation that cannot be returned.

4.3 Human Managerial Oversight and Decision Authority

The third theme is that of human managerial judgment. Jarrahi (2018) speaks about complementary capacities of human and AI, and Dellermann et al. (2019) about combined intelligence in problem solving. Managers can still be a crucial component in situations where decisions need to be made in context or with regard to unusual situations.

A procurement agent might be able to find another supplier who can provide cheaper goods, and a manager may be aware of quality issues, contractual matters or short-term delivery delays. Superficial requirements for approval are not sufficient to fulfil the obligations to effective oversight; information, authority and opportunities for meaningful challenge and intervention are necessary.

Raisch and Krakowski (2021) describe how automation and augmentation create an intertwined set of organizational tensions. Therefore, AI agents should be woven into workflows using the right responsibility sharing, not assuming that they can replace all human work.

Managers' commitment to supervision is not guaranteed to be stable. The findings reported by Dietvorst et al. (2015) and Logg et al. (2019) indicate there are different kinds of algorithm aversion and appreciation, and there are trust factors discussed by Glikson and Woolley (2020) that should be taken into account in the supervision of these arrangements.

The flexibility of the organizational governance and autonomy of the agent.

The fourth theme is related to governance arrangements that govern activities of AI-agents. The key principles of ethical AI scholarship are governance, transparency, responsibility and the consideration of unintended consequences (Floridi et al., 2018). In addition, Shin (2021) list more information reliability and security issues associated with generative AI.

There should be clear policies about information agents that can be accessed, the actions they can take and when they will need human approval to take these actions. An inventory agent may have access to the inventory and make recommendations for replenishment, but must go through some sort of approval process to move inventory. A procurement agent could be limited in the amount of organizational funds that he can put at risk, but may still be able to compare several suppliers.

Baird and Maruping (2021) highlight the relationships of delegation with agentic information systems. Governance should be part of the workflow design, permissions and monitoring – not bolted on after deployment.

Interface-level safeguards can thus be used in conjunction with independent review in an operational control regime. Shin (2021) offer evidence on explainability and trust in AI decisions, and Faraj et al. (2018) explore the way that algorithms remake coordination and organizational control. When applied to purchasing agents, these ideas help achieve approval thresholds, trackable procurement activities, exception logs, and periodical review of purchasing agent decisions.

4.5 Summary of Results

Theme	Main finding	Supporting literature
Operational efficiency	Potential information	Simon (1979); Wang et al.

	processing and workflow coordination benefits	(2024); Yao et al. (2023); Brynjolfsson et al. (2025); Noy and Zhang (2023)
Decision reliability	Inaccurate information plus autonomous execution creates operational risks	Parasuraman et al. (2000); Wang et al. (2024); Shrestha et al. (2019)
Human oversight	Contextual judgment and review remain important	Dellermann et al. (2019); Jarrahi (2018); Raisch and Krakowski (2021); Glikson and Woolley (2020); Logg et al. (2019)
Governance	Authority limits and monitoring support responsible delegation	Autio et al. (2024); Baird and Maruping (2021); Tabassi (2023); Amershi et al. (2019); Raji et al. (2020)

5. Discussion

The study findings highlight the complex relationship among technological skills, managerial power, and operational governance, revealing the organizational implications of AI agents. While the agentic systems can be used to process information and coordinate tasks, their effectiveness is not only determined by the autonomous level of the system.

The first is the link between the operational efficiency and the quality of managerial decision making. Agents can decrease information handling demands by gathering relevant records, arranging options, and coordinating activities. This potential contribution is in keeping with Simon's bounded rationality, in which he notes constraints on rational decision-making of humans. But when inputs are wrong or factors in the context are left out, speed isn't necessarily appropriate. Efficiency should be assessed in conjunction with reliability. Huang and Rust (2021) also differentiate between repetitive automation and more complex analytical functions, providing a conceptual separation that can be transferred with their marketing framework—but not actual evidence on procurement agents.

Shrestha et al. (2019) also illustrate how organizational decisions can take various forms of human–AI structures as a function of the situation. This re-emphasizes the difference between the advantages of delegating information-retrieval and the significantly different dangers of delegating purchasing power. A similar argument is advanced by Benbya et al. (2021) which explains why managerial use cases should be selected based on business processes instead of technology enthusiasm.

The second one brings in the difference between analytical support and autonomous action execution. Traditional decision-support systems usually offer information for interpretation, and agents select, and take, actions in digital workflows. Parasuraman et al. (2000) agree with the separation between levels and stages of automation. Therefore, an agent can request information without a purchase order approval and/or inventory transfer but must receive approval prior to a purchase order or inventory transfer.

The third discovery relates to the synergy between humans and AI. Jarrahi (2018) and Dellermann et al. (2019) agree that there should be a combination of computational and human judgment. Agents can perform repetitive tasks, while managers deal with exceptions, coordination issues and context-related decisions. However, according to Raisch and Krakowski (2021), there can be tensions regarding skills, control, and roles of the workforce as a result of automation and augmentation. Managerial review should then be backed up with some competency and authority.

The fourth finding relates to accountability. Baird, V., & Maruping, S. (2021) explain the theory behind delegating activities to agentic systems of information. If agents are carrying out tasks that previously were carried out by staff, the responsibility must still be clearly apparent. With regard to governance and accountability, Floridi et al. (2018) agreed that this is a right thing to do, and Shin (2021) indicated explainability as a precondition for trust.

These findings collectively imply that bounded autonomy is one possible organizing principle. This model consists of agents having certain tasks to be accomplished under specific conditions, and the final or exceptional ones that need to be approved by humans. This does not mean that human decision making is always more accurate, it is merely recommended that delegation is proportionate to the consequences of error, the reliability of information and the ability of organizations to monitor.

The empirical evidence presented should be interpreted with care. Mikalef and Gupta (2021) link organizational AI with performance related outcomes, whilst Brynjolfsson et al. (2025) and Noy and Zhang (2023) investigate productivity gains in specific work contexts that are aided by AI. Indirect evidence of readiness and potential process gains can be drawn from these studies, but this is an empirical question regarding the impact of autonomous AI agents on procurement accuracy, inventory availability, and exception handling has yet to be answered.

The paper thus continues the conversation on automation versus augmentation to contexts where an AI agent not only provides advice but also performs actions. Empirical studies should be conducted to compare the human-led, AI-assisted and bounded-autonomous approaches in terms of decision accuracy, time to complete, managerial workload and organizational trust.

6. Conclusion and Recommendations

6.1 Conclusion

This paper explored the potential and limitations of using AI agents in real-world decision-making processes within enterprise management, specifically in inventory management, procurement, and workflow coordination. Opportunities were identified in the thematic synthesis in information processing, task coordination, and repetitive administrative activities. The opportunities are aligned with research on planning and tool-use skills (Wang et al., 2024; Wang et al., 2024).

Other analyses revealed risks from an error in data, inappropriate autonomous actions, unclear accountability and ineffective managerial monitoring. Such issues are consistent with those found in the literature of automation and organisational governance (Kellogg et al., 2020; Parasuraman et al., 2000; Floridi et al., 2018). Human judgment will continue to be critical in situations that require context and interpretation (Jarrahi, 2018; Dellermann et al., 2019).

The paper calls for bounded autonomy as a defensible way of adopting AI agents, allowing for well-defined operations while maintaining human accountability for decisions that have a significant impact. Such conceptual results need to be tested by means of empirical organizational research.

6.2 Recommendations

The first step is to have differentiated operational tasks such as stock control, documentation of procurement process and tracking of workflows. These tasks allow for technological reliability to be evaluated prior to expanding authority.The tasks allow for the evaluation of technological reliability before authority is expanded (Parasuraman et al., 2000; Wang et al., 2024).

Secondly, autonomy limits should be based on operational risks. Routine, reversible activities can be automated, while significant purchasing commitments, out-of-the-ordinary inventory transfers, and consequential decisions should be signed by management (Baird & Maruping, 2021; Floridi et al., 2018).

Third, the reliable organizational records should be regarded as a prerequisite. To minimize inaccurate recommendations, inventory databases, procurement data and workflow records should be validated and updated (Shin, 2021; Wang et al., 2024).

Fourth, managers need to be trained to interpret outputs from agents, to be aware of the limitations of automation, errors or to be able to make independent judgments (Raisch & Krakowski, 2021; Jarrahi, 2018).

A topic for training should also be the calibration of trust, for managers should have procedures for challenging promising-sounding recommendations, for verifying the information that led to them, and for correcting their mistakes. There are some guidelines on human–AI interaction and empirical research on the appreciation and aversion of algorithms as a basis for this kind of training (Shin, 2021; Dietvorst et al., 2015; Logg et al., 2019).

Fifth, organisations should establish formal accountability in the form of documented permissions, records of carried out actions, performance monitoring and responsibilities for error correction (Baird & Maruping, 2021; Floridi et al., 2018).

Lastly, future research should evaluate the accuracy, speed, cost, and management trust of the human-led, AI-assisted, and bounded-autonomous operational decisions. A longitudinal study on the impact of multiple interactions with agents on managerial knowledge and organizational learning would be interesting.

References

Baird, A., & Maruping, L. M. (2021). The next generation of research on IS use: A theoretical framework of delegation to and from agentic IS artifacts. MIS Quarterly, 45(1), 315–341. https://doi.org/10.25300/MISQ/2021/15882

Benbya, H., Pachidi, S., & Jarvenpaa, S. L. (2021). Special issue editorial: Artificial intelligence in organizations: Implications for information systems research. Journal of the Association for Information Systems, 22(2), 281–303. https://doi.org/10.17705/1jais.00662

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. The Quarterly Journal of Economics, 140(2), 889–942. https://doi.org/10.1093/qje/qjae044

Dellermann, D., Ebel, P., Söllner, M., & Leimeister, J. M. (2019). Hybrid intelligence. Business & Information Systems Engineering, 61(5), 637–643. https://doi.org/10.1007/s12599-019-00595-2

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. Journal of Experimental Psychology: General, 144(1), 114–126. https://doi.org/10.1037/xge0000033

Faraj, S., Pachidi, S., & Sayegh, K. (2018). Working and organizing in the age of the learning algorithm. Information and Organization, 28(1), 62–70. https://doi.org/10.1016/j.infoandorg.2018.02.005

Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. Minds and Machines, 28(4), 689–707. https://doi.org/10.1007/s11023-018-9482-5

Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. Academy of Management Annals, 14(2), 627–660. https://doi.org/10.5465/annals.2018.0057

Huang, M.-H., & Rust, R. T. (2021). A strategic framework for artificial intelligence in marketing. Journal of the Academy of Marketing Science, 49, 30–50. https://doi.org/10.1007/s11747-020-00749-9

Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. Business Horizons, 61(4), 577–586. https://doi.org/10.1016/j.bushor.2018.03.007

Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. Academy of Management Annals, 14(1), 366–410. https://doi.org/10.5465/annals.2018.0174

Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. Organizational Behavior and Human Decision Processes, 151, 90–103. https://doi.org/10.1016/j.obhdp.2018.12.005

Mikalef, P., & Gupta, M. (2021). Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance. Information & Management, 58(3), Article 103434. https://doi.org/10.1016/j.im.2021.103434

Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. Science, 381(6654), 187–192. https://doi.org/10.1126/science.adh2586

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans, 30(3), 286–297. https://doi.org/10.1109/3468.844354

Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. Academy of Management Review, 46(1), 192–210. https://doi.org/10.5465/amr.2018.0072

Shin, D. (2021). The effects of explainability and causability on perception, trust, and acceptance: Implications for explainable AI. International Journal of Human-Computer Studies, 146, Article 102551. https://doi.org/10.1016/j.ijhcs.2020.102551

Shrestha, Y. R., Ben-Menahem, S. M., & von Krogh, G. (2019). Organizational decision-making structures in the age of artificial intelligence. California Management Review, 61(4), 66–83. https://doi.org/10.1177/0008125619862257

Simon, H. A. (1979). Rational decision making in business organizations. The American Economic Review, 69(4), 493–513. https://www.jstor.org/stable/1808698

Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Xu, C., Lin, Y., Zhao, W. X., Wei, Z., & Wen, J. (2024). A survey on large language model based autonomous agents. Frontiers of Computer Science, 18(6), Article 186345. https://doi.org/10.1007/s11704-024-40231-1