

A Study of the Impact of Generative AI on Managerial Decision Making

¹Mr. Nikhil Suryakant Salve, ²CMA Shivaji Rajmane

¹Engineering Lead, ²Assistant Professor

¹MBA(BA) Department, ²BBA/MBA Department

Zeal Institute of Management and Computer Application (ZIMCA), Pune, India

¹nikhilsalve.2015@gmail.com, ²cmashivajirajmane@gmail.com

Abstract—

Generative Artificial Intelligence (GenAI) is increasingly being used as a practical support tool for managers, especially for summarizing information, preparing reports, and exploring alternative courses of action. This paper studies how GenAI influences key aspects of managerial decision-making: decision-cycle speed, perceived decision quality, trust and transparency, adoption barriers, and perceived strategic return on investment (ROI). A quantitative survey was conducted with 100 mid-to-senior managers who have professional exposure to GenAI tools. The findings indicate that higher frequency of GenAI use, by itself, did not translate into faster decisions ($r = 0.01$; $p > 0.90$). In contrast, using GenAI for information synthesis was associated with higher perceived decision confidence and analytical depth (t-test $p < 0.05$; correlation $r = 0.25$; $p < 0.01$). The results also suggest that AI literacy and structured Human-in-the-Loop (HITL) practices are linked with greater trust and decision transparency. Concerns about hallucinations and data privacy remain major adoption constraints, reinforcing the need for verification and governance. Overall, GenAI appears to add the most value when treated as an assistant that improves the quality of managerial thinking, while responsibility and final judgment remain with the manager.

Index Terms: *Generative AI, managerial decision-making, large language models, Human-in-the-Loop, adoption barriers, decision confidence*

1. Introduction

Managerial decisions are increasingly made under tight timelines and high information volume. In many roles, the core challenge is not the lack of data but the difficulty of converting dispersed inputs into a usable judgment. Alongside structured dashboards, managers must interpret unstructured signal emails, documents, meeting notes, customer feedback, and market commentary—that are hard to scan and integrate at scale (Simon, 1955). Generative Artificial Intelligence (GenAI), especially large language model (LLM) tools, changes the nature of decision support by working directly with language: it can produce drafts, summaries, comparisons, and option outlines from natural-language prompts (Vaswani et al., 2017).

From a manager's perspective, GenAI often functions as a “cognitive assistant” that supports sense-making and communication rather than replacing judgment (Mollick, 2024; Davenport & Miller, 2022). In day-to-day work, it is commonly used to (1) compress lengthy text into key points and potential action items, (2) generate initial scenarios and trade-offs for “what-if” thinking, and (3) accelerate business writing (e.g., converting rough notes into an executive summary or a first draft of a decision memo). Used this way, GenAI can reduce effort in early-stage information gathering and structuring, while the manager retains responsibility for validating critical facts and applying organizational context.

At the same time, GenAI introduces operational risks that are relevant to managerial accountability. LLMs may generate fluent but incorrect content (“hallucinations”), requiring verification before outputs are used in decision briefs

or stakeholder communication (Bender et al., 2021). Bias is also a practical concern: if the underlying data and modeling choices encode inequities, generated recommendations can reproduce unfair patterns in high-impact domains (European Union, 2024; NIST, 2023). In addition, limited explainability can make it difficult to justify why an AI-assisted option was preferred, especially when decisions must be defensible to auditors or stakeholders (NIST, 2023). Finally, privacy and compliance constraints influence whether managers can use public tools or must rely on controlled enterprise deployments and governance (NIST, 2023; IBM Institute for Business Value, 2023).

Because of these benefits and risks, the dominant recommendation in prior work is augmentation: GenAI can assist with drafting, summarization, and option generation, while the manager remains the decision owner (Davenport & Miller, 2022; IBM Institute for Business Value, 2023). Human-in-the-Loop (HITL) practices operationalize this idea by combining fast AI output with human review checkpoints for fact validation, feasibility, ethics, and documentation of assumptions. In managerial terms, GenAI may speed up analysis and communication, but accountability does not transfer to the tool.

Research Gap:

Although GenAI is widely discussed as a disruptive capability, organizations still lack sufficient empirical clarity on what is changing in day-to-day managerial decision making. It is not always clear whether GenAI adoption leads to faster decisions, better decision quality, or improved strategic outcomes and which factors (such as AI proficiency, trust, governance, and risk perceptions) shape those outcomes. This study addresses that gap by examining managers' experiences and reported outcomes from using GenAI in decision-relevant work.

Study Objectives:

Based on the above gap, this study focuses on managerial use of GenAI and evaluates its influence through five specific objectives:

1. Analyze the impact of GenAI use on the speed and efficiency of managerial decision-making cycles.
2. Evaluate whether GenAI-assisted decision support improves the quality, depth, and confidence of decisions compared to traditional approaches.
3. Identify key organizational and psychological barriers (like hallucination risk, bias, or lack of trust) that affect GenAI adoption for decisions.
4. Assess the perceived strategic ROI (benefits versus costs) from implementing GenAI in managerial workflows.
5. Propose a practical Human-in-the-Loop framework for responsible GenAI adoption that balances AI assistance with human governance and accountability.

By empirically investigating these objectives with data from practicing managers, the study aims to derive actionable insights for both researchers and organizations on how to leverage GenAI effectively in managerial decision-making.

2. Literature Review

2.1 Evolution of AI in Managerial Decision Support

Managerial decision support has traditionally relied on decision support systems (DSS) and later business intelligence tools, which were built to work well with structured data and predefined models (Keen & Scott Morton, 1978). As organizational work digitized, however, a growing share of decision-relevant evidence moved into narrative channels (emails, documents, reports, and other text). This shift matters because structured dashboards do not automatically help managers interpret context, connect qualitative signals, or summarize large volumes of narrative material. LLM-based GenAI represents a different capability: transformer architectures enable models to process language context and produce outputs such as summaries, comparisons, and draft briefs that resemble a first-pass analyst write-up (Vaswani

et al., 2017; Mariani & Dwivedi, 2024).

Generative AI vs Traditional and Predictive AI:

Traditional AI approaches in organizations have often been rule-based, while predictive analytics focuses on estimating outcomes from historical patterns. GenAI differs in what it produces: it generates new language artifacts (e.g., a decision brief, an explanation of options, or alternative strategies) that can be directly used in managerial communication (Mollick, 2024). This is relevant because managerial work is strongly shaped by framing, interpretation, and stakeholder communication. In that sense, GenAI shifts AI from a back-end analytical engine to an interactive assistant that managers can query in everyday language (Davenport & Miller, 2022).

2.2 GenAI's Contributions to Decision Quality and Speed

Empirical and practitioner-oriented studies increasingly describe GenAI as useful for both efficiency and exploration in knowledge work (Brynjolfsson et al., 2023; Deloitte, 2024; Noy & Zhang, 2023). Instead of supporting a single calculation, it can rapidly propose alternative framings, draft multiple response options, and outline trade-offs for managerial review (Agrawal et al., 2022). This support is particularly relevant under bounded rationality: when time and attention are limited, managers may narrow the set of options they evaluate (Simon, 1955). GenAI can widen the initial option space and help organize information, but the manager still needs to validate assumptions and select the final course of action.

Managerial Decision Quality:

Decision quality may improve when GenAI helps managers surface more relevant evidence, connect related inputs, and reduce omissions in an initial analysis (Mariani & Dwivedi, 2024). At the same time, fluent output is not equivalent to verified truth, and organizational context—constraints, trade-offs, and tacit knowledge—remains primarily human-held (Davenport & Miller, 2022; Bender et al., 2021). For this reason, the literature generally frames GenAI outputs as decision inputs to be evaluated rather than answers to be adopted without scrutiny (IBM Institute for Business Value, 2023).

Managerial Decision Speed:

Many organizations expect GenAI to shorten parts of the decision cycle by speeding up summarization and first-draft writing (McKinsey Global Institute, 2023; Gartner, 2024). However, time savings are not automatic. Prior work suggests that productivity gains depend on how the tool is embedded into workflows and how much additional review and rework is required (Brynjolfsson et al., 2023). In practice, decision-cycle speed is more likely to improve when GenAI outputs flow into meetings and approvals through a repeatable process with defined validation steps.

2.3 Human–AI Collaboration and Trust

Across research and practice reports, a recurring position is that the best managerial outcomes come from collaboration rather than full automation (Davenport & Miller, 2022; IBM Institute for Business Value, 2023). Human-in-the-Loop (HITL) models place GenAI in roles such as reading support, drafting, and option generation, while managers perform feasibility checks, risk screening, ethical judgment, and final accountability. With repeated use and explicit review checkpoints, managers can develop calibrated trust—knowing when GenAI output is useful and when it requires closer validation (NIST, 2023).

Trust and AI Literacy:

AI literacy influences how managers interpret and act on GenAI output. When users understand typical strengths and failure modes, they are less likely to treat fluent text as guaranteed accuracy and more likely to verify high-impact

claims (Bender et al., 2021; NIST, 2023). At the same time, familiarity can reduce resistance because the tool becomes easier to control through better prompting and structured review habits (Davis, 1989). This “calibrated trust” framing supports managerial transparency: GenAI can be used as a hypothesis generator, but decision justification still requires documented assumptions and corroboration.

2.4 Adoption Barriers: Hallucinations, Bias, and Governance

Even when managers recognize value in GenAI, scaling it across an organization remains challenging. Industry and research sources repeatedly highlight barriers related to output reliability, bias and fairness, explainability, and governance requirements (NIST, 2023; Deloitte, 2024; European Union, 2024). The most frequently discussed constraints include:

- **Hallucinations:** LLMs can generate plausible statements that are not factually grounded, creating a verification burden for managerial use (Bender et al., 2021). In decision contexts, this implies that outputs used in briefs, reports, or recommendations should be validated against trusted sources before they influence action.
- **Algorithmic Bias:** Where models and data reflect historical inequities, generated content can reinforce unfair patterns in sensitive decisions (European Union, 2024; NIST, 2023). For managerial adoption, this motivates governance controls such as bias monitoring, review for protected attributes, and human approval for high-impact outputs.
- **Lack of Explainability:** Limited transparency in how outputs are produced can weaken managerial accountability, especially when decisions require justification to stakeholders (NIST, 2023). Governance guidance therefore emphasizes traceability and documentation—capturing prompts, key assumptions, and supporting evidence when AI suggestions inform a decision (NIST, 2023).
- **Privacy and Compliance:** Managers may be restricted from sharing sensitive or regulated data with external GenAI services, which can slow adoption in compliance-heavy industries (NIST, 2023; IBM Institute for Business Value, 2023). Organizations respond with usage policies, access controls, and—in some cases—private or on-premises deployments to reduce data leakage risk (NIST, 2023).

Synthesizing the above streams, the literature indicates that GenAI may strengthen managerial decision work by improving how quickly managers can interpret large text inputs and by broadening the options considered early in analysis (McKinsey Global Institute, 2023; Mariani & Dwivedi, 2024). At the same time, the same sources emphasize that value depends on governance—verification routines, bias controls, explainability practices, and clear data policies are needed to manage hallucination, fairness, and privacy risks (NIST, 2023; European Union, 2024). Building on this foundation, the present study tests these ideas empirically by examining how managers’ report GenAI’s impact and which factors shape adoption and trust in day-to-day decision contexts.

3. Research Methodology

3.1 Research Design and Approach

This study uses a quantitative, cross-sectional survey approach to capture how managers are currently using GenAI and what outcomes they associate with that use. The design is descriptive (to map usage patterns and perceptions) and explanatory (to examine hypothesized relationships between GenAI-related factors and decision outcomes), consistent with the MBA research expectations and reporting approach (Savitribai Phule Pune University, n.d.).

A structured questionnaire was developed to measure:

- **GenAI usage patterns** (frequency and types of tasks performed with GenAI)

- **Decision-making outcomes** (perceived speed of decisions, confidence in decisions, quality of outputs)
- **Adoption enablers** (e.g. manager's AI proficiency, presence of human oversight practices)
- **Adoption barriers** (concerns about AI hallucinations, data privacy)
- **Perceived strategic ROI** of GenAI integration.

These constructions correspond to the conceptual model: independent variables including frequency of GenAI use, AI literacy, usage of HITL practices, and risk perceptions (hallucination risk, privacy risk); and dependent variables including decision speed, decision quality/confidence, adoption extent, and perceived ROI.

3.2 Sample and Data Collection

The target respondents were mid-level and senior professionals in decision-influencing roles within Pune's corporate environment who had exposure to GenAI tools (e.g., ChatGPT, Gemini, Microsoft Copilot). Purposive sampling was applied so that responses came from participants who could realistically comment on AI- assisted decision work. A total of 100 complete survey responses were collected and used for analysis.

- Being in a decision-making role (team leads, managers, department heads, etc., with at least two years of managerial experience),
- Having used a GenAI tool at work for tasks like analyzing data, writing reports, or generating content.

The questionnaire was administered online and anonymously to encourage honest responses. **100 respondents** provided complete and usable data for analysis. For an MBA-level exploratory study using correlation and group comparisons, this sample size was treated as practically sufficient for initial hypothesis testing, while recognizing that larger samples would provide greater statistical power and generalizability (Davis, 1989).

3.3 Analysis Strategy and Hypotheses Testing

Responses were organized and analyzed using spreadsheet and basic statistical tools. Five hypotheses (H1-H5) were evaluated to test relationships between GenAI usage, decision outcomes, trust, perceived risks, and perceived ROI.

- **H1:** Frequent GenAI use → Faster decision-making cycle.
- **H2:** GenAI-based information synthesis → Higher quality decisions (accuracy, depth, confidence).
- **H3:** Higher AI literacy → Greater trust and lower resistance to AI in decisions.
- **H4:** High perceived AI risk (hallucination, privacy) → Lower adoption extent (less usage across the enterprise).
- **H5:** Using a HITL decision model → Higher ROI from AI adoption (compared to human-only or AI-only decision approaches).

Appropriate **statistical tests** were applied:

- **Descriptive statistics** (frequencies, means, SD) to profile GenAI usage and perceptions.
- **Correlation analysis** (Pearson's r) to measure linear relationships among key continuous variables (e.g., GenAI usage frequency with decision outcomes).
- **Group comparison tests** (e.g., t-tests or ANOVA) to compare means of outcome variables between different groups (e.g., managers who use GenAI vs. those who don't, or high-frequency vs low-frequency users).
- **Chi-square tests** for associations between categorical groupings (like high/medium/low usage categories and levels of perceived transparency).
- **Regression analysis** (linear regression) to examine the predictive power of multiple independent variables (e.g., risk perceptions) on a dependent variable (e.g., GenAI adoption level), controlling for others.

All hypothesis tests used a significance threshold of $\alpha = 0.05$, with **p-values** reported to indicate support or lack

thereof for each hypothesis.

3.4 Validity and Reliability Considerations

To address internal validity, clear operational definitions for each construct were used, and a pilot test of the survey was conducted to refine ambiguous questions. The self-reported nature of the survey means results might be influenced by common method bias and social desirability bias (managers could overstate success or understate difficulties). To mitigate this, respondents were assured of confidentiality, and questions were carefully worded to be neutral. The cross-sectional design limits causal claims, so results are interpreted in terms of associations rather than definitive cause-effect proof.

4. Results

4.1 GenAI Usage among Managers

Overall, GenAI usage was common in the sample. About 65% of respondents reported using GenAI frequently (once a week or more). The most common use cases were summarizing text, drafting emails and reports, and preparing outlines for analysis. Adoption appeared stronger in information-heavy functions such as IT, strategy, and marketing. Many respondents described a shift in effort: less time spent on first pass reading and writing, and more time spent reviewing, correcting, and deciding based on AI-assisted drafts.

4.2 Hypothesis Testing and Key Findings

H1 (Decision Speed):

The hypothesis proposed that more frequent GenAI use would be linked to faster decision cycles. The analysis did not support this relationship. The correlation between GenAI use frequency and self-rated decision speed was essentially zero ($r = 0.01$) and not statistically significant ($p > 0.90$). Within this sample, higher frequency of GenAI use was not associated with faster decisions.

Conclusion:

“H1 not supported”. In this dataset, decision speed does not vary with GenAI usage frequency.

H2 (Decision Quality/Confidence):

The second hypothesis focused on GenAI as a synthesis tool whether it improves perceived accuracy, depth, and confidence. This hypothesis was supported. Managers who used GenAI extensively for research, summarization, and drafting reported higher confidence and perceived analytical depth. Group comparisons showed a statistically significant difference (t-test $p < 0.05$). In addition, GenAI usage frequency showed a positive association with decision confidence ($r = 0.25$; $p < 0.01$).

Conclusion:

“H2 is supported”. Managers who reported heavier use of GenAI for synthesis-related tasks also reported higher confidence and perceived analytical depth, consistent with the observed t-test ($p < 0.05$) and correlation ($r = 0.25$; $p < 0.01$).

H3 (Trust and AI Literacy):

The third hypothesis examined whether AI literacy supports trust and transparency in AI-assisted decisions. The results supported this expectation. Respondents with higher self-rated AI proficiency were more likely to report that they could explain and justify how AI influenced their decisions. A significant positive relationship was observed between perceived proficiency and decision transparency ($r = 0.28$; $p < 0.01$). A chi-square test comparing transparency ratings across GenAI usage groups was also significant ($p < 0.001$), suggesting that more engaged users experience greater perceived transparency.

Conclusion:

“H3 is supported”. Higher self-rated AI proficiency is associated with higher perceived decision transparency ($r = 0.28$; $p < 0.01$), and transparency ratings differed across GenAI usage groups (chi-square $p < 0.001$).

H4 (Risk Perception and Adoption):

The fourth hypothesis expected that stronger concerns about hallucinations and privacy would reduce GenAI adoption across the enterprise. In this dataset, the quantitative evidence was inconclusive. Although respondents described risk concerns as a barrier to scaling GenAI, the statistical model did not show a clear, significant negative association between risk-perception scores and adoption extent. Therefore, this study cannot confirm H4 based on the current sample.

Conclusion:

“H4 is not clearly supported by the current sample”. Risk concerns were reported qualitatively, but the quantitative analysis did not establish a statistically significant negative relationship with adoption extent.

H5 (HITL and ROI):

The fifth hypothesis proposed that organizations using a Human-in-the-Loop approach would report higher strategic ROI. In this study, the difference in ROI perceptions across HITL-related groupings was not statistically significant (ANOVA $p = 0.37$). While some respondents reported higher ROI perceptions in settings with more formal review practices, the quantitative comparison did not provide sufficient evidence to confirm H5 in this sample.

Conclusion:

“H5 not confirmed in the present analysis”. In this dataset, the ROI comparison did not reach statistical significance (ANOVA $p = 0.37$).

To summarize the hypothesis outcomes in plain terms, two of the five hypotheses (H2 and H3) are supported, one (H1) is not supported, and two (H4 and H5) remain inconclusive based on the current evidence.

remain unresolved due to insufficient evidence. Table 1 provides a brief overview:

Table 1. Summary of Hypotheses and Outcomes

Hypothesis	Outcome
H1: Frequent GenAI use reduces decision cycle time	Not supported (no significant relationship)
H2: GenAI info synthesis improves decision quality/confidence	Supported (significant positive effect)
H3: Higher AI literacy increases trust & transparency	Supported (significant positive effect)
H4: Perceived AI risks reduce GenAI adoption	Not clearly supported (inconclusive)
H5: HITL model yields higher strategic ROI	Not clearly supported (inconclusive)

4.3 Additional Findings

Beyond the hypotheses, qualitative feedback from respondents provided context:

- Many managers noted GenAI is excellent for first drafts or brainstorming, but they always “fact-check and tweak” outputs before using them in final decisions or official documents.
- Hallucination incidents were reported by ~30% of users (e.g., AI generating non-existent statistics), reinforcing the need for fact verification.
- Some respondents from highly regulated industries (finance, healthcare) indicated strict internal policies limiting GenAI usage due to compliance they expected slower adoption until these tools could be deployed within a secure environment.

These findings paint a picture of cautious optimism: managers are embracing GenAI where it clearly adds value and saves time, but they are mindful of its pitfalls.

5. Discussion

Based on our findings, GenAI's most consistent contribution is linked to perceived decision quality rather than speed. In this study, managers who used GenAI for synthesis and drafting reported higher confidence and greater perceived analytical depth (H2 supported). This result is directionally consistent with evidence that GenAI can improve performance and work quality in knowledge tasks, although impacts vary by context and user capability (Brynjolfsson et al., 2023; Cui et al., 2024).

This study also clarifies an important non-finding: GenAI usage frequency was not associated with faster decision speed (H1 not supported). While industry narratives often anticipate cycle-time reduction from automation of summarization and drafting (McKinsey Global Institute, 2023), our results indicate that frequency of use alone is not a reliable predictor of perceived decision speed in this sample. This suggests that any time savings may depend on how GenAI is integrated into decision workflows rather than on usage intensity by itself.

The study shows that AI literacy is associated with higher perceived transparency in AI-assisted decisions (H3 supported). Managers with higher self-rated proficiency reported greater ability to explain and justify how AI influenced their decisions. This finding supports the broader argument that calibrated use—understanding limitations and applying verification—matters for responsible deployment (NIST, 2023).

Risk concerns remain prominent in how managers describe adoption decisions. Although H4 was not statistically supported in this sample, respondents repeatedly highlighted hallucinations and privacy constraints as reasons for cautious rollout. These concerns mirror governance themes emphasized in risk and regulatory guidance, which stress validation, privacy controls, and accountability practices for trustworthy AI use (NIST, 2023; European Union, 2024).

The ROI-related hypothesis (H5) was not confirmed (ANOVA $p = 0.37$). Within this cross-sectional survey, perceived ROI did not differ significantly by HITL-related groupings. This indicates that, in this dataset, ROI perceptions are not clearly explained by HITL practices alone, and more detailed organizational measures may be needed to evaluate ROI relationships.

Overall, this study supports a practical position: GenAI contributes most when used as an assistant for drafting and synthesis, while managers retain responsibility for validation, context-setting, and final judgment.

6. Conclusion

This study shows that GenAI is already influencing managerial work, particularly where decisions depend on reading, synthesizing, and communicating large amounts of information. Within our sample, the strongest observed benefit is higher perceived decision confidence and analytical depth when GenAI is used for synthesis and drafting (H2 supported). In contrast, GenAI usage frequency alone did not show a measurable relationship with perceived decision speed (H1 not supported). These results suggest that adoption intensity is not sufficient by itself; the value managers perceive depends on how GenAI is applied to decision-relevant tasks.

The study also highlights a condition for sustainable value: responsible use. Because managers remain accountable for outcomes, organizations need practical support such as AI literacy programs, clear guidance on data handling, and Human-in-the-Loop checkpoints for high-impact outputs. This recommendation is consistent with risk management and governance guidance that emphasizes accountability, privacy safeguards, and traceability for trustworthy AI use (NIST, 2023; IBM Institute for Business Value, 2023).

In summary, the evidence from this research supports treating GenAI as a capability that strengthens decision preparation and managerial communication rather than as an autonomous decision-maker. When combined with verification habits, privacy safeguards, and clear accountability, managers can use GenAI to improve the quality of decision inputs without weakening governance.

7. Managerial Implications

For practicing managers and organizations, several actionable implications emerge from this study:

- **Embed GenAI thoughtfully into decision workflows:** Don't treat AI tools as isolated add-ons. Instead, redesign key processes (e.g. report preparation, project planning) so that GenAI can handle preliminary analysis, while clear human review checkpoints ensure quality control.
- **Invest in AI training and literacy:** Organizations should provide training programs to help managers learn effective prompting techniques, understand GenAI's limitations, and interpret outputs. AI-literate managers are more likely to use GenAI in a transparent and controlled way, consistent with our H3 findings and enterprise guidance on responsible AI adoption (NIST, 2023; Deloitte, 2024).
- **Establish verification protocols:** Create standard validation checklists for AI-generated content. For example, require that any critical fact or figure suggested by AI is cross verified by a human or a trusted source. This reduces the likelihood of decisions being influenced by incorrect outputs (Bender et al., 2021; NIST, 2023).
- **Implement AI governance and data policies:** Develop guidelines on GenAI usage, including what types of data can be input to AI tools (to protect privacy) and where human approval is mandatory. Aligning internal policy with external regulatory expectations and risk guidance helps maintain accountability and compliance (European Union, 2024; NIST, 2023; IBM Institute for Business Value, 2023).
- **Measure and monitor ROI of GenAI use:** Track metrics like time saved on routine tasks, improvements in decision turnaround, and instances of AI-enabled insights that influenced strategy. Use these measures to refine how GenAI is deployed and to build a business case for further adoption (McKinsey Global Institute, 2023; PwC, 2024; Stanford Institute for Human-Centered Artificial Intelligence, 2024).

By following these practices, organizations can better harness GenAI's capabilities while avoiding pitfalls, ultimately creating a synergy where human expertise and AI insights together drive superior decisions.

8. Limitations and Future Research

This study has limitations typical of survey-based research. Because the measures are self-reported, responses may reflect social desirability or individual optimism/pessimism about technology. The cross-sectional design captures a snapshot in time and cannot establish how effects evolve as GenAI usage matures. In addition, the sample is concentrated in Pune and may not reflect adoption dynamics in other regions or industries.

Future studies can strengthen evidence by using larger and more diverse samples, combining surveys with objective indicators (time saved, error rates, quality ratings), and tracking outcomes over time. Experimental designs comparing human-only, AI-only, and HITL decision conditions would help isolate causality. There is also scope to evaluate how enterprise safeguards such as retrieval-augmented generation, private deployments, and formal review workflows change hallucination rates, trust, and decision outcomes.

In summary, while our study provides initial empirical insights on GenAI in managerial decision-making, there is ample scope for further investigation to guide best practices and to keep pace with this rapidly evolving technological landscape.

References

- [1] A. Agrawal, J. S. Gans, and A. Goldfarb, *Prediction Machines, Updated and Expanded: The Simple Economics of Artificial Intelligence*. Harvard Business Review Press, 2022.
- [2] E. Brynjolfsson, D. Li, and J. J. Raymond, *Generative AI at Work*. National Bureau of Economic Research, 2023.
- [3] T. H. Davenport and S. M. Miller, *Working with AI: Real Stories of Human Machine Collaboration*. MIT Press, 2022.
- [4] Deloitte, *The State of Generative AI in the Enterprise: Now Decides Next*. Deloitte AI Institute, 2024.
- [5] Gartner, *Top Strategic Technology Trends for 2024: AI-Augmented Development and Management*, 2024.
- [6] M. M. Mariani and Y. K. Dwivedi, Generative artificial intelligence in management: A research agenda for the era of large language models, *International Journal of Information Management*, 2024.
- [7] McKinsey Global Institute, *The economic potential of generative AI: The next productivity frontier*, 2023.
- [8] E. Mollick, *Co-Intelligence: Living and Working with AI*. Portfolio/Penguin, 2024.
- [9] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, On the dangers of stochastic parrots: Can language models be too big?, in *Proc. ACM Conf. Fairness, Accountability, and Transparency*, pp. 610-6123, 2021.
- [10] R. Cui, D. Li, and J. Zhang, The impact of generative AI on productivity and work quality: Evidence from knowledge work tasks, *SSRN Electronic Journal*, 2024.
- [11] A. R. Doshi and O. P. Hauser, Generative AI can help evaluate strategic decisions by leveraging diverse perspectives, *Strategic Management Journal*, 2024.
- [12] European Union, *Regulation (EU) 2024/ laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*, 2024.
- [13] IBM Institute for Business Value, *The CEOs guide to generative AI: Strategy, governance, and value creation*. IBM, 2023.
- [14] P. G. W. Keen and M. S. Scott Morton, *Decision support systems: An organizational perspective*. Addison-Wesley, 1978.
- [15] S. Noy and W. Zhang, Experimental evidence on the productivity effects of generative artificial intelligence, working paper, 2023.
- [16] PwC, *AI Jobs Barometer 2024*. PricewaterhouseCoopers, 2024.
- [17] H. A. Simon, A behavioral model of rational choice, *The Quarterly Journal of Economics*, vol. 69, no. 1, pp. 99-118, 1955.
- [18] Stanford Institute for Human-Centered Artificial Intelligence, *AI Index Report 2024*. Stanford University, 2024.
- [19] A. Vaswani *et al.*, Attention is all you need, in *Advances in Neural Information Processing Systems*, 2017.
- [20] F. D. Davis, Perceived usefulness, perceived ease of use, and user acceptance of information technology, *MIS Quarterly*, vol. 13, no. 3, pp. 319-340, 1989.
- [21] National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. U.S. Department of Commerce, 2023.
- [22] P. Lewis *et al.*, Retrieval-augmented generation for knowledge-intensive NLP tasks, in *Advances in Neural Information Processing Systems*, 2020.
- [23] Savitribai Phule Pune University, *MBA Program Curriculum and Research Guidelines*, n.d.