AI Assisted Executive Decision Intelligence for Proactive Enterprise Performance Management

Lee Ting Feng[1], Toh Yang Jian[1] , Ismail, A.O[1], Ismael, S. N.[1]

[1]Azman Hashim International Business School, Universiti Teknologi Malaysia

DOI: https://dx.doi.org/10.47772/IJRISS.2026.101000692

Received: 30 October 2026; Accepted: 04 November 2026; Published: 16 November 2026

ABSTRACT

Fragmented enterprise information presents challenges for executives seeking to identify performance changes, interpret emerging operational risks, and determine appropriate managerial responses. This paper presents the design and development of the X Decision Intelligence Platform, an AI-assisted executive decision-support prototype developed by the research team. The prototype focuses on the Food and Beverage quick-service restaurant (QSR) segment as its demonstrator, while the broader architecture encompasses executive, operations, finance, customer, workforce, and strategy and sustainability analytics. Adopting a design science-informed approach, the study documents the development of the working prototype and examines its architecture, functional components, AI-assisted development process, and proposed decision-governance mechanisms. The platform integrates KPI monitoring, anomaly alerts, and a prompt-driven large language model decision engine designed to support contextual interpretation and recommended managerial actions. The analysis describes a four-layer architecture and a two-phase AI-assisted development process, while identifying requirements for evidence traceability, uncertainty management, knowledge updates, and human oversight. The study contributes a documented prototype development case and a conceptual pathway towards accountable enterprise decision support. Further empirical evaluation is required to establish the platform's effectiveness in improving managerial decision quality and operational performance.

Keywords - Decision intelligence; executive decision support; large language models; early warning systems; enterprise analytics; human oversight; AI-assisted development

I. INTRODUCTION

Enterprise executives must interpret operational and strategic information that originates in different organisational functions. A sales decline may need to be examined alongside product availability, staffing, operating costs and customer feedback. Merely presenting separate measures does not resolve the managerial task of deciding which changes matter, how they may be connected and what evidence is needed before action. Dashboard research shows that information layout, context and interaction require deliberate design choices (Bach et al., 2023). Reviews of AI business value similarly show that technical capabilities must be considered together with organisational resources, adoption conditions and strategic alignment (Enholm et al., 2022; Perifanis & Kitsios, 2023).

The project identifies distributed enterprise information, slow reporting cycles, limited cross-functional visibility and reactive monitoring as the motivating problems. The present study focuses on the Food and Beverage quick-service restaurant (QSR) segment, while its intended enterprise scope covers additional functions and management levels.

In this paper, decision intelligence means an integrated workflow connecting performance information, warning signals, contextual interpretation, recommendations and accountable managerial review. AI modelling offers techniques relevant to pattern recognition and decision support (Sarker, 2022), but a useful executive platform requires more than generating persuasive explanations. The system must preserve the distinction between observed evidence, a model’s interpretation and an action that management has authorised. Research

on human - AI collaboration cautions that combining human and AI capabilities does not automatically improve performance (Vaccaro et al., 2024).

1.1 Purpose and Scope of the Study

This study aims to examine the design and development of an AI-assisted executive decision intelligence prototype intended to support proactive enterprise performance management. Specifically, it investigates how fragmented enterprise performance information can be organised into an integrated decision-support architecture, how AI-assisted development tools contribute to prototype construction, and how the resulting design supports the progression from performance monitoring towards contextual interpretation and managerial decision support.

The study addresses three research questions:

RQ1: How can enterprise performance monitoring, anomaly identification, and AI-assisted interpretation be integrated into a unified executive decision-support architecture?

RQ2: How was the AI-assisted executive decision intelligence prototype developed, and what architectural and functional components resulted from the development process?

RQ3: What governance requirements, limitations, and evaluation considerations emerge from the developed prototype in supporting accountable managerial decision-making?

The study focuses on the Food and Beverage QSR segment as the primary demonstrator within a broader enterprise decision intelligence architecture. Its contribution lies in documenting the developed prototype, explaining the integration of its components, and identifying the requirements for subsequent empirical evaluation.

II. LITERATURE BACKGROUND

2.1 From Dashboards to Contextual Decision Support

Dashboards organise information for recurring analytical and monitoring tasks. Bach et al. (2023) identify design patterns that support structured choices about information presentation and interaction. In the project context, this suggests that a unified workspace should not become an indiscriminate collection of charts. Indicators should be presented with definitions, relevant comparisons and pathways to inspect supporting records. The platform’s intended contribution is to connect such views with alerts and proposed responses.

Enholm et al. (2022) synthesise enablers, inhibitors and value-generating mechanisms associated with organisational AI use. Perifanis and Kitsios (2023) further discuss AI in relation to business and IT strategy. Together, these reviews justify examining whether the platform’s analytical functions fit real decision responsibilities.

2.2 Anomaly Detection and Early Warning

Time series anomaly detection seeks patterns that differ from expected behaviour. A comprehensive evaluation by Schmidl et al. (2022) demonstrates why method comparison matters, while Darban et al. (2024) review deep learning approaches, applications and unresolved challenges. Multivariate methods additionally examine dependencies among variables over time (Wang et al., 2025). These research streams are relevant to a platform intended to monitor linked enterprise indicators, although the project document does not specify a particular statistical or machine learning detector.

Benchmark selection and evaluation procedures can substantially affect apparent performance. Paparrizos et al. (2022) provide a benchmark suite designed to support reproducible univariate anomaly evaluation. Garg et al.

(2022) discuss the importance of scoring and event-sensitive assessment in multivariate detection, and Baldán and García-Gil (2024) examine benchmarking using the UCR archive. Accordingly, a dashboard alert should not be described as a proven forecast of business failure. A threshold breach, an unusual historical pattern and a prospective risk estimate are different outputs that require different validation.

2.3 Explainability and Language Model Outputs

The project places a prompt-driven language model at the centre of contextual interpretation. Explainability research identifies challenges across system design, development and deployment (Saeed & Omlin, 2023). For an executive user, a readable explanation is helpful only if it remains connected to reliable evidence. A narrative that cites relevant indicators is not, by itself, proof that the system has discovered a causal mechanism.

Natural language generation can produce fluent content that is unsupported by the available evidence (Ji et al., 2023). This risk is relevant when a platform summarises performance, proposes likely causes or recommends actions. The explanation should distinguish verified measurements from hypotheses and should identify missing information. The platform's rationale can then support managerial investigation without being mistaken for causal confirmation.

2.4 Human Oversight and Organisational Learning

Human-centred AI emphasises preserving meaningful user control as automation increases (Shneiderman, 2022). In a field study of medical diagnosis, Lebovitz et al. (2022) show how professionals' engagement with opaque AI depended on practices for interrogating its outputs. Although that domain differs from executive performance monitoring, the findings motivate designing ways for managers to question and inspect AI recommendations. Vaccaro et al. (2024) provide broader evidence that human-AI combinations can underperform their stronger component, reinforcing the need for task-specific evaluation.

Organisational learning also depends on the quality of the information retained. Sambasivan et al. (2021) document how neglected data work can generate compounding downstream problems. For a platform with an enterprise memory, this means that recommendations, decisions and outcomes need provenance, review and version control.

2.5 Challenges and Future Directions

The literature and project description identify related challenges: integrating inconsistent information; selecting indicators that reflect decisions; validating alerts; preventing unsupported generated explanations; protecting sensitive records; and maintaining human accountability. Scalability also concerns organisational processes, not only technical capacity. Additional modules require agreed definitions, data ownership and review responsibilities. Future work should therefore examine the complete information-to-action workflow alongside model performance (Enholm et al., 2022; Saeed & Omlin, 2023).

III. METHODOLOGY

3.1 Research Design

This study adopts a design science-informed, artifact-centred case study approach to examine the development of an AI-assisted executive decision intelligence platform. Design science research is relevant to information systems studies involving the construction of technological artifacts to address identified organisational problems (Peffers et al., 2007).The research focuses on three interrelated activities: examining the business problem and intended decision-support requirements, documenting the design and development of a working prototype, and analysing the resulting architecture, functional components, and governance considerations.

The developed prototype serves as the primary unit of analysis. The study uses the prototype's documented design and development information to explain how its components are organised to support enterprise performance monitoring and managerial interpretation. Although the prototype was developed and executed by the research team, the study does not claim that its effectiveness in improving managerial decision-making has been empirically established. Accordingly, the research contributes an artifact-centred development case rather than a completed experimental evaluation of decision-support effectiveness.

3.2 Prototype Development Context

The X Decision Intelligence Platform was developed by the research team to address fragmented enterprise performance information and the need for more proactive management support. The Food and Beverage QSR segment served as the primary demonstrator, using simulated KPI data. The wider platform design encompasses six business domains: executive, operations, finance, customer, workforce, and strategy and sustainability analytics.

3.3 AI-Assisted Development Process

Prototype development involved two documented phases. During the first phase, WorkBuddy was used to generate underlying data, anomaly logic, and reasoning materials; Miora supported the initial interface concept; and CodeBuddy integrated the outputs into an application. During the second phase, Claude Code supported interface revisions, additional modules, feature enhancements, and data-integrity improvements, while Codex contributed to deployment and hosting activities. These activities are examined as part of the artifact development process.

3.4 Artifact and Document Analysis

The developed prototype was examined through its documented architecture, functional descriptions, and AI-assisted development workflow. These materials were analysed to identify the relationship between the intended business requirements, system components, development activities, and decision-support functions.

The analysis was organised into six categories: (1) business problems and target users, (2) business modules and intended decision functions, (3) architecture and data flow, (4) AI-assisted decision engine, (5) prototype development process, and (6) decision governance and expected managerial value.

The findings were organised according to these categories and interpreted in relation to relevant literature on dashboard design, artificial intelligence, anomaly detection, explainability, and human oversight.

3.5 Study Boundaries

Although the research team developed and ran the prototype, the present study does not report controlled experiments, systematic managerial user testing, or measured improvements in decision quality. Consequently, the findings are limited to the documented artifact, its development process, and its potential contribution to executive decision support. Empirical validation is proposed as a subsequent research stage.

IV. RESULTS

4.1 Business Requirements and Domain Coverage

The documented value proposition is to transform fragmented enterprise data into explainable insights, earlier risk identification and actionable recommendations. The six domains are the Executive Command Center, Operations Analytics, Finance Analytics, Customer Analytics, Workforce Analytics, and Strategy and Sustainability Analytics. The command centre is intended to consolidate priorities, while the domain modules provide more detailed views for investigation.

Table 4.1 maps the six domains to their intended decision focus.

Domain	Intended decision focus
Executive command centre	Enterprise priorities and performance exceptions
Operations	Operating performance and emerging risks
Finance	Financial trends and business impact
Customer	Customer indicators and service issues
Workforce	Staffing and resource planning
Strategy and sustainability	Initiatives KPIs and ESG monitoring

According to Table 4.1, for the QSR demonstrator, the design can be understood through an illustrative investigation: declining outlet performance could be reviewed alongside operational, customer and staffing indicators.

4.2 Four Layer Architecture

The architecture distinguishes business, application, data and technology layers as shown in Figure 1.

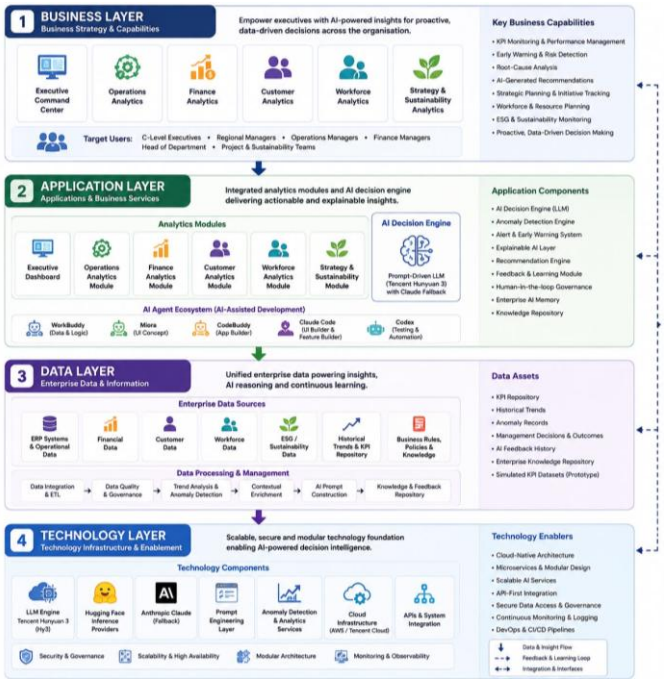


Figure 1: Four Layer Architecture of the X Decision Intelligence Platform (developed for this project)

As illustrated in Figure 1, The business layer defines the management domains, users and intended capabilities. It connects performance monitoring with early warnings, strategic tracking, workforce planning and sustainability oversight. Its central purpose is to express why the platform is needed and who should use its outputs.

The application layer combines domain analytics with an AI decision engine, anomaly detection, alerts, recommendations, explanations, feedback and human review. The presence of these components in the architecture establishes a design arrangement, not independent confirmation of their operational maturity. Separating functions is nevertheless useful because an explanation service, an anomaly rule and an approval record need distinct testing and responsibilities.

The data layer contains KPI repositories, historical trends, anomaly records, policies, knowledge and management feedback. It explicitly includes simulated KPI datasets for the prototype. The architecture also

identifies enterprise sources such as ERP, financial, customer, workforce and sustainability data. The depicted processing sequence moves through integration, quality checks, analysis, contextual enrichment, prompt construction and feedback storage.

The technology layer identifies language model services, inference access, a fallback provider, prompt engineering, analytics services, cloud infrastructure and integration interfaces. Security, availability, modularity and observability appear as cross-cutting enablers. These labels should be translated into deployment controls and testable requirements before claims of enterprise readiness are made.

4.3 Prompt Driven Decision Engine

The prototype incorporates a prompt-driven large language model decision engine intended to support the interpretation of enterprise performance information. Its documented design combines simulated KPI values, historical trends, organisational policies, and business context to generate contextual explanations and suggested managerial responses.

The decision-support workflow connects performance information with the identification of potential anomalies, interpretation of possible contributing factors, assessment of business implications, and formulation of recommended next steps. This arrangement extends the intended role of the platform beyond descriptive KPI presentation towards AI-assisted managerial interpretation.

However, the generation of a plausible explanation does not establish the actual cause of a performance change. AI-generated interpretations should therefore be treated as hypotheses requiring verification against relevant business evidence. Similarly, the identification of unusual KPI patterns should not be interpreted as validated predictive performance without appropriate testing.

The prototype consequently illustrates an approach to integrating performance monitoring and AI-assisted interpretation, while the accuracy, reliability, and managerial usefulness of its outputs remain subjects for subsequent empirical evaluation (Ji et al., 2023; Saeed & Omlin, 2023).

4.4 Outcomes of the AI-Assisted Development Process

The two-phase AI-assisted development process resulted in a working executive decision intelligence prototype incorporating a multi-domain interface, an organised application structure, and components intended to support performance monitoring and AI-assisted managerial interpretation.

The first phase established the initial prototype foundation through the preparation of simulated data, anomaly-related logic, reasoning materials, interface concepts, and application integration. The second phase focused on refining the interface, extending modules and features, addressing identified data-integrity issues, and supporting application deployment.

The resulting artifact demonstrates how different AI-assisted development tools were used across complementary activities, including data preparation, interface design, application integration, feature refinement, and deployment support.

These development outcomes provide evidence of artifact construction and execution. However, the contribution of individual tools to development productivity, software quality, or managerial decision-making effectiveness was not independently measured. The findings therefore concern the documented development process and resulting prototype rather than comparative performance advantages of the development tools.

4.5 Knowledge Evolution and Decision Governance

The design proposes incremental updates to business rules, policies, KPI definitions and organisational knowledge so that recommendations can reflect changing priorities. This is conceptually different from retraining the underlying model. Context can be updated without establishing that model parameters have learned new organisational knowledge. The paper therefore treats continuous learning as a proposed knowledge-management and feedback mechanism rather than demonstrated online model learning.

The enterprise memory is intended to capture recommendations, management decisions and outcomes, including acceptance, modification or rejection. To be useful, these records need an identified owner, dated evidence and an auditable link between the recommendation and the decision. Outcome changes may have multiple causes; the memory should not automatically assign success to the AI. Concerns about downstream data problems reinforce the need for quality review before feedback is reused (Sambasivan et al., 2021).

Human review is a central design principle. It describes supporting evidence, confidence indicators and business impact assessments intended to help managers validate and approve actions. This aligns with the emphasis on retaining user control in human-centred AI (Shneiderman, 2022). However, an approval interface alone does not establish effective oversight. Users must be able to challenge assumptions, identify missing evidence and decline recommendations without losing access to the underlying information.

V. DISCUSSION

5.1 Managerial Value and Cross Functional Interpretation

The findings address the first research question by demonstrating, at the architectural level, how enterprise performance monitoring, anomaly-related information, and AI-assisted interpretation can be organised within a unified decision-support platform. The four-layer architecture provides a structured relationship between business requirements, application functions, data resources, and enabling technologies. The resulting design establishes a foundation for connecting performance information with contextual explanations and proposed managerial responses. Nevertheless, the presence of these components does not independently demonstrate that managerial decisions have become more accurate, timely, or effective.

The developed prototype brings together executive monitoring, domain-specific analytics, AI-generated explanations, and proposed managerial actions within a unified design. Its potential value is not the number of domains displayed, but whether those domains help users investigate a consequential question. A regional manager, for instance, needs sufficient context to decide whether a flagged change warrants local investigation or escalation. This application of dashboard principles is consistent with Bach et al. (2023), while the expected organisational benefits remain contingent on strategic and operational fit (Enholm et al., 2022; Perifanis & Kitsios, 2023).

The QSR focus provides a bounded setting for evaluation. It is more defensible to test a defined monitoring and investigation task in one segment than to infer enterprise-wide effectiveness from the existence of six modules. Subsequent expansion should be justified by decision needs, data quality and ownership. Finance, workforce and sustainability outputs also have different interpretation requirements, so a shared interface should not erase domain-specific meaning.

5.2 Explainability Does Not Establish Correctness

A plausible language model explanation can help users articulate hypotheses, but it can also create unwarranted certainty. The engine should distinguish observed change, possible contributors, unavailable evidence and proposed action. Ji et al. (2023) document hallucination risks in generated language, while Saeed and Omlin (2023) explain why interpretability challenges extend across the system life cycle. These studies support testing factual consistency and user understanding separately from presentation quality.

The case also raises a distinction between human approval and meaningful human judgement. Evidence from medical AI shows that professionals may need practices for interrogating outputs before incorporating them into critical decisions (Lebovitz et al., 2022). Transfer to executive settings requires further study, but the implication is practical: users need accessible source evidence and opportunities to challenge recommendations. Human-AI task performance should be measured because benefits cannot be inferred from collaboration alone (Vaccaro et al., 2024).

5.3 Reproducibility

Modularity is useful when changes can be isolated and their effects tested. Adding a domain should involve documented indicator definitions, controlled prompt changes, regression checks and explicit permissions. Knowledge updates should be versioned so an investigator can reconstruct which policies and context informed a recommendation. These practices would make the artifact easier to evaluate and are consistent with treating demonstration and evaluation as distinct design science activities (Peffers et al., 2007).

The documented development process also highlights the importance of traceability in AI-assisted software construction. The involvement of multiple AI development tools creates a need to preserve information about generated artifacts, subsequent modifications, integration decisions, and deployment configurations. Such documentation is necessary to distinguish the contributions of different development activities and to support future replication or extension of the prototype.

5.4 Evaluation Agenda

A future evaluation should begin with data integrity and functional correctness. For simulated data, document how values and anomalies were generated and what the expected results are. Where live data is introduced, examine missingness, inconsistent units, timestamps, duplicates and authorised access. Synthetic scenarios are useful for testing functions but do not establish real-world predictive validity. Dataset selection and benchmark design should be justified using the concerns identified in anomaly evaluation research (Paparrizos et al., 2022; Baldán & García-Gil, 2024).

The warning component should then be tested against simple, transparent baselines. Report detection precision, missed events, false alerts and time to detection when these measures fit the use case. For linked indicators, specify whether the model actually evaluates multivariate dependencies or whether it simply presents variables together. Scoring functions and event definitions require documentation because they can change apparent performance (Garg et al., 2022; Schmidl et al., 2022). A deep learning extension is a possible research direction, not a necessary condition for a useful prototype (Darban et al., 2024; Wang et al., 2025).

User evaluation should involve the intended management roles and realistic tasks. Relevant measures include time to locate an issue, accuracy of interpretation, quality of the proposed next step and ability to identify an incorrect AI suggestion. Where feasible, compare the current reporting workflow with the platform and distinguish AI assistance from interface improvements. Human-only and combined conditions can test whether assistance improves outcomes in this particular setting (Vaccaro et al., 2024). Realised benefits such as reduced consolidation time or earlier intervention require a measured baseline, a defined observation period and careful attribution.

In relation to the third research question, the prototype analysis identifies several requirements for progressing towards accountable executive decision support. These include the reliability of underlying KPI data, accuracy of anomaly-related outputs, factual consistency of AI-generated explanations, traceability of supporting evidence, and meaningful human oversight of recommended actions. These requirements provide a basis for a subsequent empirical evaluation of the prototype.

VI. CONCLUSION

This study examined the design and development of the X Decision Intelligence Platform, an AI-assisted executive decision-support prototype intended to facilitate proactive enterprise performance management. The

prototype was developed and executed by the research team, with the Food and Beverage QSR segment serving as the primary demonstrator within a broader six-domain enterprise architecture.

The findings highlight three contributions. First, the four-layer architecture provides a structured approach to integrating enterprise performance information, anomaly-related functions, AI-assisted interpretation, and intended managerial decision-support capabilities. Second, the documented two-phase development process illustrates how multiple AI-assisted development tools were employed in constructing and refining a working prototype. Third, the analysis identifies governance and evaluation requirements associated with evidence traceability, explanation reliability, organisational knowledge management, and human oversight.

The study contributes to the understanding of how an AI-assisted enterprise performance monitoring prototype can be developed and organised to support the progression towards contextual and accountable managerial decision support. However, prototype development and execution do not establish improvements in decision quality or organisational performance. Future research should therefore evaluate the platform using realistic managerial scenarios, functional testing, and systematic user assessment to determine its effectiveness, reliability, and practical value.

REFERENCES

1. Bach, B., Freeman, E., Abdul-Rahman, A., Turkay, C., Khan, S., Fan, Y., & Chen, M. (2023). Dashboard design patterns. IEEE Transactions on Visualization and Computer Graphics, 29(1), 342–352. https://doi.org/10.1109/TVCG.2022.3209448
2. Baldán, F. J., & García-Gil, D. (2024). Benchmarking anomaly detection methods: Insights from the UCR Time Series Anomaly Archive. Expert Systems. Advance online publication. https://doi.org/10.1111/exsy.13767
3. Darban, Z. Z., Webb, G. I., Pan, S., Aggarwal, C. C., & Salehi, M. (2024). Deep learning for time series anomaly detection: A survey. ACM Computing Surveys. https://doi.org/10.1145/3691338
4. Enholm, I. M., Papagiannidis, E., Mikalef, P., & Krogstie, J. (2022). Artificial intelligence and business value: A literature review. Information Systems Frontiers, 24, 1709–1734. https://doi.org/10.1007/s10796-021-10186-w
5. Garg, A., Zhang, W., Samaran, J., Savitha, R., & Foo, C.-S. (2022). An evaluation of anomaly detection and diagnosis in multivariate time series. IEEE Transactions on Neural Networks and Learning Systems. https://doi.org/10.1109/TNNLS.2021.3105827
6. Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Dai, W., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. ACM Computing Surveys, 55(12), Article 248. https://doi.org/10.1145/3571730
7. Lebovitz, S., Lifshitz-Assaf, H., & Levina, N. (2022). To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. Organization Science, 33(1), 126–148. https://doi.org/10.1287/orsc.2021.1549
8. Paparrizos, J., Kang, Y., Boniol, P., Tsay, R. S., Palpanas, T., & Franklin, M. J. (2022). TSB-UAD: An end-to-end benchmark suite for univariate time-series anomaly detection. Proceedings of the VLDB Endowment, 15(8), 1697–1711. https://doi.org/10.14778/3529337.3529354
9. Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. Journal of Management Information Systems, 24(3), 45–77. https://doi.org/10.2753/MIS0742-1222240302
10. Perifanis, N.-A., & Kitsios, F. (2023). Investigating the influence of artificial intelligence on business value in the digital era of strategy: A literature review. Information, 14(2), Article 85. https://doi.org/10.3390/info14020085
11. Saeed, W., & Omlin, C. (2023). Explainable AI (XAI): A systematic meta-survey of current challenges and future opportunities. Knowledge-Based Systems, 263, Article 110273. https://doi.org/10.1016/j.knosys.2023.110273

12. Sambasivan, N., Kapania, S., Highfill, H., Akrong, D., Paritosh, P., & Aroyo, L. (2021). "Everyone wants to do the model work, not the data work": Data cascades in high-stakes AI. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, Article 39. https://doi.org/10.1145/3411764.3445518
13. Sarker, I. H. (2022). AI-based modeling: Techniques, applications and research issues towards automation, intelligent and smart systems. SN Computer Science, 3, Article 158. https://doi.org/10.1007/s42979-022-01043-x
14. Schmidl, S., Wenig, P., & Papenbrock, T. (2022). Anomaly detection in time series: A comprehensive evaluation. Proceedings of the VLDB Endowment, 15(9), 1779–1797. https://doi.org/10.14778/3538598.3538602
15. Shneiderman, B. (2022). Human-centered AI: Ensuring human control while increasing automation. Proceedings of the 33rd ACM Conference on Hypertext and Social Media. https://doi.org/10.1145/3538882.3542790
16. Vaccaro, M., Almaatouq, A., & Malone, T. W. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. Nature Human Behaviour, 8, 2293–2303. https://doi.org/10.1038/s41562-024-02024-1
17. Wang, F., Jiang, Y., Zhang, R., Wei, A., Xie, J., & Pang, X. (2025). A survey of deep anomaly detection in multivariate time series: Taxonomy, applications, and directions. Sensors, 25(1), Article 190. https://doi.org/10.3390/s25010190