

Responsible Automation in E-business: A Framework for Agentic AI Integration and Governance

BIMTECH Business Perspectives

I-24

© The Author(s) 2026

DOI: 10.1177/25819542261479531

bsp.bimtech.ac.in



Mohammad Talha Siddiqui¹  and Yusuf Kamal¹ 

Abstract

This study proposes the Responsible Autonomy Framework (RAF), a layered governance architecture for agentic artificial intelligence (AI) in e-business. The RAF integrates operational efficiency, strategic adaptability, and responsible oversight into a single layered model. Using a conceptual research design grounded in an integrative literature review, agency theory, and sociotechnical systems thinking, the study develops a three-layer model comprising operational autonomy, strategic autonomy, and responsible autonomy. In which responsibility acts as the governing layer rather than an afterthought. The review finds that existing governance frameworks address transparency, accountability, or fairness largely in isolation, with none spanning the full path from operational automation to embedded governance, while the RAF addresses this by treating responsible design as a condition for sustainable autonomy rather than a competing value to be balanced against it. The framework offers managers actionable guidance for deploying agentic AI responsibly and offers regulators a structured basis for balancing innovation with oversight, while extending agency theory and sociotechnical systems theory through a reconceptualisation of agentic AI as a sociotechnical principal-cum-agent.

Keywords

Responsible autonomy framework (RAF), agentic artificial intelligence (AI), e-business, operational autonomy, strategic autonomy, explainable AI (XAI)

¹ University of Lucknow, Uttar Pradesh, India

Corresponding author:

Mohammad Talha Siddiqui, University of Lucknow, Lucknow, Uttar Pradesh 226003, India.

E-mail: siddiqui_talha@lkouniv.ac.in



Creative Commons Non Commercial CC BY-NC: This article is distributed under the terms of the Creative Commons Attribution-NonCommercial 4.0 License (<http://www.creativecommons.org/licenses/by-nc/4.0/>) which permits non-Commercial use, reproduction and distribution of the work without further permission provided the original work is attributed.

Introduction

The digital transformation of organisations has evolved through several distinct phases from the early computerisation and process automation to the current integration of intelligent, adaptive and autonomous systems. Within this evolution, e-businesses have been at the forefront of adopting automation technologies across domains such as customer service, supply chain management, marketing and regulatory compliance. Automation through chatbots, robotic process automation (RPA) and large language models (LLMs) has enabled faster query handling, reduced costs and improved personalisation in customer engagement (Duesterwald et al., 2024; Hughes et al., 2025; Van et al., 2025). In the supply chain and logistics, predictive maintenance, dynamic route optimisation and AI-driven demand forecasting have redefined the efficiency benchmarks (Gong, 2025; Nicoletti & Appolloni, 2025). The marketing functions have increasingly used artificial intelligence (AI) recommendation systems to create personalised consumer experiences, while compliance activities such as auditing, fraud detection and transaction verification have become progressively automated (Hnatyshyn et al., 2025; Sallem et al., 2024).

Despite these developments, most automation has stayed focused on deterministic and rule-based systems designed primarily for efficiency, speed and accuracy. The emergence of agentic AI represents a marked shift in this pattern, moving attention from pre-programmed automation to autonomous and goal-directed decision-making. The agentic systems combine perception, reasoning and adaptive learning, allowing them to operate as semi-independent entities capable of handling the complex interactions within both the digital and physical contexts (Duesterwald et al., 2024; Hughes et al., 2025). This shift carries major implications for the electronic business. The AI agents now contribute to strategic decision-making, manage the coordination across the supply chains, personalise the customer experiences in real time and handle the compliance requirements with limited human involvement (Gong, 2025; Van et al., 2025).

The integration of autonomy into business operations introduces new challenges related to governance and accountability. In traditional automated systems, the decision-making stays under human supervision, but the autonomous AI agents make the boundaries of responsibility less clear when mistakes, ethical issues or unintended outcomes occur. The questions of transparency, interpretability and trust become central, particularly in areas such as finance, healthcare and international electronic commerce (Biswas & Talukdar, 2025; Hughes et al., 2025). Regulatory frameworks such as the European Union Artificial Intelligence Act and the General Data Protection Regulation (GDPR) highlight the importance of strong accountability and compliance structures, reinforcing that autonomy must develop alongside responsibility (Manda et al., 2024; Vatankhah et al., 2024).

Thus, while the e-businesses benefit from the efficiency and innovation through the agentic AI, the rapid evolution of autonomy necessitates a framework that integrates operational automation, strategic adaptability and responsible governance. The Responsible Autonomy Framework (RAF), introduced below,

positions responsibility not as an afterthought but as the structural layer of autonomy, ensuring that the digital transformation of the e-businesses evolves in alignment with societal expectations and regulatory norms.

Research Problem

The rise of agentic AI in e-business introduces new and significant opportunities and risks, necessitating critical reflection on its governance and accountability. The traditional automation frameworks stressed efficiency, consistency and cost reduction, with human supervision ensuring responsibility for the outcomes (Bohnsack & De Wet, 2025; Thapliyal & Thapliyal, 2024). The growing autonomy of the AI agents capable of independently interpreting the environments, adapting to the feedback and making strategic decisions raises significant challenges regarding transparency, explainability and liability (Biswas & Talukdar, 2025; Hughes et al., 2025). Unlike deterministic systems, agentic AI operates with degrees of unpredictability, creating the ‘black box’ effect that obscures the decision processes and complicates both managerial oversight and regulatory compliance (Alshaketheep et al., 2024; Vannostrand et al., 2024).

The challenge becomes more serious in the fields where the AI decisions have financial, legal or ethical consequences. In financial services, autonomous agents now assist with credit evaluation, fraud detection and investment planning. The mistakes or bias in these settings can cause reputational harm, financial loss or legal liability (Munoko et al., 2020; Sallem et al., 2024). Within the supply chain networks, agentic AI improves procurement and logistics efficiency, yet its independent actions can expose weaknesses in the coordination and accountability between the organisations (Duesterwald et al., 2024; Gong, 2025). In marketing, agentic systems such as recommendation tools strengthen personalisation but may also reinforce bias, compromise data privacy or weaken consumer trust (Biswas & Talukdar, 2025; Van et al., 2025). The autonomy without governance erodes confidence across these domains.

The current governance and ethical frameworks stay fragmented. The frameworks stressing risk management (Wirtz et al., 2022), digital ethics (Ashok et al., 2022) or holistic responsibility (Kottur, 2024) each provide valuable insights, yet they primarily address either compliance, trust or ethics in isolation. None sufficiently reconcile the tension between empowering the AI agents with autonomy and embedding accountability throughout their decision lifecycles. The policy developments such as the GDPR and the EU AI Act highlight the growing regulatory concern, but the operationalisation in the fast-evolving agentic AI environments remains unclear (Manda et al., 2024; Vatankhah et al., 2024).

The research problem can therefore be summarised as the governance gap in transitioning from automation to autonomy in e-business. The AI agents promise adaptability, scalability and innovation, but the absence of integrated frameworks risks creating opaque, unaccountable and potentially harmful systems. This article responds to this gap by proposing the RAF, which positions responsibility as the embedded structural element of autonomy rather than the external corrective

measure. In doing so, it aims to bridge the divide between technological advancement and ethical governance, ensuring that the agentic AI strengthens rather than undermines the foundations of trust in digital business ecosystems.

Objectives and Research Questions

The primary objective of this study is to conceptualise and develop the RAF, which is a layered model that integrates operational efficiency, strategic adaptability and responsible governance into a unified approach for e-business autonomy. Unlike the prior frameworks, the RAF positions responsibility as the structural component of autonomy itself. The aim is to advance the theory, provide managerial guidance and inform regulatory practice by ensuring that autonomy in the digital business ecosystems evolves in tandem with accountability and trust.

To achieve this objective, the study addresses the following research questions:

- RQ1: How can the existing approaches to automation, governance and digital ethics be synthesised to conceptualise the unified framework for responsible autonomy in e-business?
- RQ2: In what ways does the RAF contribute to advancing the theory, managerial practice and policy in the governance of agentic AI for e-business?

Literature Review

The academic study of automation in electronic commerce has developed across the three largely separate streams: the operational adoption of information technologies for the efficiency and process standardisation; the governance frameworks designed to manage the risks associated with the digital systems; and the emerging body of work on responsible AI. While each stream has produced valuable insights, they have evolved with limited cross-pollination, leaving the gap at precisely the intersection that the agentic AI now occupies. This review traces the development of each stream and identifies the specific limitations that the RAF is designed to address.

Evolution of E-business Automation

The early research on electronic commerce automation focused primarily on aligning the information technology governance with organisational performance. Liu et al. (2010) proposed the holistic governance framework for e-business success, arguing that the effective management of risk, performance and compliance was central to realising the strategic potential of the digital systems. Their framework was grounded in the assumption that human oversight remained primary and that automation served a supporting role in standardising the routine

processes. Similar themes appeared in the work on enterprise resource planning and supply chain coordination, where the efficiency gains of the automated data pipelines were offset by the concerns about the system rigidity and the limited capacity of the rule-based models to handle the contextual variation (Hnatyshyn et al., 2025; Murphy et al., 2024).

The scope of the automation expanded substantially with the emergence of AI-driven customer engagement tools. Adam et al. (2021) demonstrated that AI-based chatbots in customer service could improve response consistency and reduce handling time, though their study also identified the tension between scalability and the quality of the complex, context-sensitive interactions. The deployment of LLMs extended these capabilities, enabling the conversational agents to engage with a broader range of queries and to support personalisation at scale (Duesterwald et al., 2024; Van et al., 2025). In logistics and supply chain management, predictive maintenance systems and dynamic routing algorithms began to demonstrate measurable reductions in disruption exposure and operational cost (Gong, 2025; Nicoletti & Appolloni, 2025). In accounting and compliance, the automated auditing tools improved the accuracy and speed of anomaly detection while reducing the workload associated with the manual review (Ramon-Poma et al., 2024; Sallem et al., 2024).

Despite these advances, a persistent limitation ran through the automation literature: the systems described remained fundamentally reactive and rule-bound. They could execute the predefined workflows efficiently but struggled to adapt to unfamiliar conditions, exercise contextual judgement or act strategically in dynamic environments. Hughes et al. (2025) characterise this as the distinction between task automation and genuine intelligence, arguing that the former improves efficiency within the known parameters while the latter requires the capacity for reasoning, adaptation and goal-directed behaviour. The transition from one to the other is the defining challenge that agentic AI addresses and that the existing governance frameworks have been slow to catch up with.

Rise of Agentic AI in E-business

The limitations of conventional automation have driven the growing interest in agentic AI systems that combine autonomy, adaptability and continuous learning. Unlike the rule-based predecessors, the agentic systems can decompose the goals, interpret the feedback and improve their performance over time, acting as semi-independent participants within digital ecosystems rather than as passive execution engines (Biswas & Talukdar, 2025; Hughes et al., 2025). The implications for the e-business are substantial. In e-commerce, recommendation systems have moved beyond the fixed association rules to employ reinforcement learning, generating personalised suggestions in real-time based on the evolving customer behaviour (Biswas & Talukdar, 2025; Van et al., 2025). In supply chain management, the adaptive agents now simulate disruption scenarios and autonomously reconfigure the logistics networks to maintain operational continuity (Gong, 2025; Nicoletti & Appolloni, 2025). In financial services, agentic systems support credit

assessment, fraud detection and portfolio optimisation using continuous learning to balance risk and responsiveness (Vannostrand et al., 2024).

The integration of natural language processing and multimodal reasoning has expanded the role of agentic systems beyond background data processing into active organisational participation. The conversational agents now contribute to decision-making in the customer support and strategic management contexts, not merely retrieving the information but synthesising it into recommendations and adapting their approach based on conversational cues (Duesterwald et al., 2024; Hughes et al., 2025). The use of reinforcement learning from human feedback has further aligned the agentic behaviour with the ethical principles and the contextual priorities, enabling the organisations to embed the value-based constraints into the learning process itself (Biswas & Talukdar, 2025; Huang, 2025).

These developments represent the qualitative shift in the relationship between AI and organisational decision-making. The systems now in use are no longer the adjuncts to human judgement but the active co-participants whose decisions carry the real operational and ethical weight. As Bohnsack and De Wet (2025) argue, AI is increasingly becoming the strategy itself rather than the tool for executing the strategies defined by humans. This shift demands governance approaches that extend beyond risk mitigation to address accountability, transparency and the distribution of responsibility in the systems where human agency and AI agency are increasingly intertwined.

Challenges in Responsible AI Deployment

The governance challenges associated with the agentic AI deployment centre on the four interrelated dimensions: transparency, accountability, fairness and regulatory compliance. Each has received substantial attention in the literature, though typically in isolation from the others.

Transparency, which is the capacity of the AI systems to render their decision processes interpretable to the users, managers and regulators, is widely recognised as foundational to trust. Explainable AI (XAI) methods such as LIME and SHAP have demonstrated the technical feasibility of generating post-hoc explanations for the complex model outputs, showing which variables drove the specific recommendation or classification (Jahan et al., 2025; Sanjammagari et al., 2026). In fraud detection, the transparent decision trails have been shown to improve both the detection of false positives and the maintenance of consumer confidence (Ansari et al., 2023). However, the interpretability of the individual decisions becomes progressively harder to guarantee as the models grow more complex and as the agentic systems engage in the multi-step reasoning processes that resist simple post-hoc explanation (Petrenko, 2025).

Accountability, which is the attribution of responsibility for the AI-generated decisions, raises distinct challenges in the agentic contexts where the decisions emerge from the interaction of multiple agents rather than from a single identifiable system. Franklin et al. (2022) propose a causal framework for attributing responsibility to artificial autonomous agents, incorporating elements such as

intent, autonomy and foreseeability. Triantafyllou and Radanovic (2023) extend this to the decentralised multi-agent settings, proposing computational methods for approximating the degrees of responsibility using Monte Carlo Tree Search. These contributions are technically sophisticated, but their integration into the operational governance practices in the e-business remains largely unexplored.

Fairness concerns the risk that the AI systems replicate or amplify the structural inequities embedded in the training data. The research has documented how the hiring tools, credit scoring models and customer segmentation systems can reproduce the discriminatory patterns that undermine both the ethical legitimacy and the regulatory compliance (Sallem et al., 2024; Shankar, 2024). The challenge is not merely technical; it is also organisational, requiring that fairness be treated as the design criterion rather than the post-deployment audit item. In e-business, fairness lapses carry the reputational, legal and ethical consequences that extend beyond the immediate transaction to affect long-term consumer trust.

Regulatory compliance is becoming an increasingly structured requirement. The GDPR mandates data minimisation, purpose limitation and the right to explanation for the automated decisions affecting individuals (Hughes et al., 2025; Vatankhah et al., 2024). The European Union Artificial Intelligence Act introduces the risk-based categorisation framework that classifies the AI systems according to their potential for harm, imposing higher transparency and oversight requirements on the high-risk applications in sectors such as financial services, healthcare and employment. The organisations operating agentic AI systems in these sectors must demonstrate not only regulatory compliance at a point in time but the ongoing adaptability to the regulatory environment that continues to evolve. The pace of regulatory development often outstrips the capacity of the organisations to adjust their governance structures, creating what Manda et al. (2024) describe as the persistent compliance lag.

Taken together, these challenges indicate that responsible AI deployment in e-business requires more than the sum of its parts. Addressing transparency in isolation does not resolve accountability; achieving regulatory compliance does not guarantee fairness. What is needed is an integrated framework that treats these dimensions as the structurally connected requirements rather than the discrete risk categories to be managed separately.

Existing Governance Frameworks and Their Limitations

This subsection examines five representative governance frameworks individually, building on the dimensions introduced above, to identify precisely where each falls short for agentic AI environments.

Liu et al. (2010) established one of the earliest structured approaches to IT governance in e-business, integrating risk management, performance monitoring and organisational alignment within the unified framework. While foundational, their model was designed for human-supervised automation and does not address the governance challenges that arise when AI systems act with the degrees of independent agency.

Al-Mushayt (2019) extended the governance thinking into the AI-augmented e-government contexts, identifying privacy, transparency and citizen trust as the central concerns in the automation of public services. The framework represented the important early recognition that AI deployment requires trust-building mechanisms beyond technical performance, but its scope remained limited to supervised automation rather than autonomous decision-making.

Wirtz et al. (2022) proposed a more comprehensive risk-based governance model organising the AI risks into the six categories: technological, data-related, ethical, legal, organisational and societal. Their integrative approach brought a useful structure to the governance field and provided the basis for the systematic risk assessment. However, its primary orientation is towards risk control rather than the responsible enablement of autonomy. It offers limited guidance on how organisations can use the adaptive and strategic capabilities of the agentic AI while maintaining accountability.

Ashok et al. (2022) addressed the ethical dimension more directly, proposing an ontological framework for digital ethics that translates values such as intelligibility, accountability, fairness, autonomy and privacy into operational guidelines. This contribution provided the principled moral basis for responsible AI, but the framework's applicability to the genuinely autonomous systems where the AI agents make the complex decisions without the direct human involvement remains limited. The challenge of governing the emergent agent behaviour that was not explicitly anticipated at the design time is not fully resolved.

Kottur (2024) offered the most holistic of the existing models integrating fairness, transparency, accountability, privacy and robustness within the framework that emphasises multi-stakeholder cooperation. Like its predecessors, however, the framework is primarily designed for the contexts in which human oversight is readily available. It does not provide the specific mechanisms for governing the agents that operate with significant independence, particularly in the contexts where the speed and complexity of the agentic decision-making exceed the capacity for real-time human review.

The cumulative limitation of these frameworks is their failure to conceptually integrate the operational capabilities of the agentic AI with the governance mechanisms required to ensure responsible deployment. As Table 1 illustrates, each framework addresses a subset of the relevant dimensions, but none provides the end-to-end model that spans from the operational automation through the strategic agentic capabilities to the embedded responsible governance.

Columns represent: IT Governance: alignment of AI with the organisational performance; Trustworthy Automation: trust and transparency in the AI-assisted processes; Risk Management: structured identification and mitigation of the AI-related risks; Ethical Operationalisation: translating the ethical principles into the actionable guidelines; Holistic Responsibility: integration of the multiple responsibility dimensions; Agentic Autonomy: governance of the semi-independent AI agents.

Table 1. Comparison of Existing Governance Frameworks with the Proposed RAF.

Framework	IT Governance	Trustworthy Automation	Risk Management	Ethical Operational	Holistic Responsibility	Agentic Autonomy
Liu et al. (2010)	✓	Partial	✗	✗	✗	✗
Al-Mushayt (2019)	✓	✓	✓	Partial	Partial	✗
Wirtz et al. (2022)	✓	✓	✓	Partial	Partial	✗
Ashok et al. (2022)	Partial	✓	Partial	✓	Partial	✗
Kottur (2024)	✓	✓	✓	✓	✓	✗
RAF (This Study)	✓	✓	✓	✓	✓	✓

Notes: ✓: addressed; Partial: partially addressed; ✗: not addressed.

Research Gap and Positioning of This Study

The review of the existing frameworks shows the consistent pattern: the governance approaches in e-business have developed in response to the challenges of the automated and AI-assisted systems operating under meaningful human oversight but have not kept pace with the shift toward agentic AI that operates with degrees of independence that fundamentally alter the governance requirements. Three specific gaps are identifiable.

First, no existing framework provides the integrated model that connects the operational automation, strategic agentic capabilities and the responsible governance within a single, coherent architecture. The governance literature and the agentic AI literature have developed largely in parallel with limited cross-referencing. The RAF addresses this by treating the responsibility not as the external control layer but as the structural component of the autonomy model itself.

Second, the existing frameworks address the governance primarily as a matter of risk management, compliance or ethical principles. They do not account for the possibility that responsibility and autonomy can be designed to reinforce each other; that building explainability, accountability and privacy-by-design into the agentic systems from the outset may enhance rather than constrain their operational and strategic effectiveness. The RAF makes this argument explicitly and provides the framework through which it can be operationalised.

Third, none of the reviewed frameworks provides actionable guidance for the managers deploying the agentic AI in the e-business contexts. The transition from the theoretical principles to the operational decisions about the system design, human oversight structures and the regulatory alignment requires a prescriptive model oriented toward design and action, which the RAF provides in line with the Gregor (2006) Type V theory classification.

Research Methodology

The conceptual research design is adopted with the primary aim of developing the RAF for e-business autonomy. The conceptual research does not rely on empirical data collection but advances the knowledge by synthesising, extending and reconfiguring the existing theoretical perspectives into the new contribution (MacInnis, 2011). In line with Gregor's (2006) typology of theory contributions, the RAF constitutes the Type V theory 'theory for design and action' in that it prescribes how agentic AI systems can be structured and governed responsibly in e-business contexts.

Conceptual Theory-building Approach

The conceptual approach is particularly appropriate given the evolving and unsettled nature of AI autonomy in e-business where the empirical consensus is still limited. Following MacInnis (2011), the contributions are made through the integration and extension: the integration by drawing together the multiple strands of literature—automation, agentic AI and responsible governance—that have largely developed in isolation; and the extension by reconfiguring these literatures into the unified framework that positions the responsibility not as an add-on but as the structural layer of autonomy.

Rather than describing the agentic AI merely as the technical evolution of automation, the RAF positions the AI agents as semi-autonomous actors situated within sociotechnical systems. This conceptualisation aligns with agency theory, which examines accountability and decision delegation and with sociotechnical systems theory, which emphasises the interdependence of technological artefacts and organisational structures. By drawing on these perspectives, the RAF advances beyond the descriptive accounts of the AI applications to propose a prescriptive model for responsible autonomy in e-business.

Literature-informed Conceptual Synthesis and Mechanism Extraction

The RAF was constructed through the literature-informed synthesis process rather than the systematic coding exercise. The prior governance and automation frameworks (e.g., Al-Mushayt, 2019; Ashok et al., 2022; Kottur, 2024; Liu et al., 2010; Wirtz et al., 2022) were critically examined to identify the mechanisms they emphasised, such as efficiency, risk management, accountability and ethical safeguards.

From this analysis and the rigorous literature review, the mechanisms relevant to e-business autonomy were extracted and clustered abductively into the three interlinked domains:

The operational autonomy: mechanisms that automate and optimise the routine processes for efficiency and scalability.

The strategic autonomy: mechanisms that support adaptive decision-making, coordination and foresight in the changing markets.

The responsible autonomy: mechanisms that embed governance, accountability and trust as the cross-cutting requirements.

This mechanism extraction process ensured that each element of the RAF is grounded in the prior studies while being reorganised into the layered model that integrates efficiency, adaptability and responsibility. The abductive reasoning approach, moving iteratively between the theoretical insights and the observed organisational challenges, enabled the framework to reconcile the fragmented literatures and highlight the underexplored connections.

Design and Action Orientation

Although conceptual in nature, the RAF is oriented towards design and action. Following Gregor (2006), such theories do not merely explain or predict the phenomena but offer guidance for practice. Accordingly, the RAF is presented as the prescriptive artefact that operates across the three domains:

For the theory, it extends the sociotechnical and agency theory by embedding responsibility into autonomy.

For the practice, it provides the managers with a structured roadmap for deploying the agentic AI responsibly in digital commerce.

For the policy, it offers the regulators insights into balancing innovation with governance in the e-business ecosystems.

Conceptual Framework

The RAF: Overview and Logic

The RAF is built on the single governing argument: that the responsible deployment of agentic AI in e-business is not achieved by adding governance as an afterthought to the autonomous systems but by treating the responsibility as the structural layer of the autonomy model from the outset. This argument distinguishes the RAF from the prior frameworks which have tended to address efficiency, risk and ethics as the separate concerns to be balanced against one another. In the RAF, the three layers—operational autonomy, strategic autonomy and responsible autonomy are logically ordered and functionally interdependent. The operational autonomy provides the efficiency substrate; the strategic autonomy enables the adaptive, goal-directed behaviour; and the responsible autonomy supplies the governance architecture that gives the first two layers their legitimacy, accountability and societal acceptability.

The framework draws on the two theoretical traditions. The agency theory originating with Jensen and Meckling (1976) provides the conceptual vocabulary for understanding the delegation of the decision-making authority from the principals to the agents and the governance mechanisms required to align the agent behaviour with the principal interests. In the context of agentic AI, the principal-agent relationship is no longer exclusively human: the AI agents act on behalf of the organisational principals and the governance challenges associated with this delegation the information asymmetry, moral hazard, and misaligned incentives

take on new technical and ethical dimensions. The RAF addresses these by embedding the accountability and transparency mechanisms directly into the agent architecture rather than relying solely on post-hoc monitoring.

The sociotechnical systems theory, which emphasises the mutual constitution of the social structures and the technological artefacts (Orlikowski, 1992; Trist & Bamforth, 1951), provides the second theoretical foundation. From this perspective, the AI agents are not merely technical tools but sociotechnical entities whose behaviour is shaped by and shapes the organisational and social contexts in which they operate. The responsible autonomy layer of the RAF instantiates this perspective by requiring that human oversight, ethical alignment and regulatory compliance be embedded within the system design rather than imposed externally after deployment. The three layers of the RAF thus represent not the technical stack but the sociotechnical architecture in which the capacity for autonomous action and the structures of accountability are co-designed.

Structure of the Framework

The RAF is organised as a three-layer model in which each layer represents the distinct domain of the AI capability and the governance responsibility. The layers are cumulative rather than sequential: the operational autonomy operates continuously as the foundation; the strategic autonomy is built upon and extends the operational layer; and the responsible autonomy functions as the governing super-structure that shapes and constrains the activities of both. Figure 1 illustrates the structure and the relationships between the layers.




Figure 1. Proposed Responsible Autonomy Framework for E-business.

The relationships between the layers operate in both directions. The outputs of the operational autonomy the data processing, language understanding, knowledge navigation and workflow execution provide the informational and technical substrate upon which the strategic autonomy depends. The strategic decisions about the pricing, procurement, market engagement and the multi-agent coordination in turn generate the conditions that the responsible autonomy layer must govern: the higher the degree of strategic autonomy exercised by the AI agents, the more critical it becomes that explainability, accountability, privacy and human oversight are properly instantiated. At the same time, the governance mechanisms of the responsible autonomy layer actively shape how the operational and strategic functions are designed and constrained. The explainability requirements, for example, influence which AI architectures are appropriate for the high-stakes operational tasks; the accountability frameworks determine how the autonomous procurement decisions must be logged and auditable; and the human-in-the-loop structures define the thresholds at which the strategic decisions require human review before the execution.

Operational Autonomy

The operational autonomy encompasses the AI capabilities that enable the e-businesses to automate the routine, high-volume processes with minimal human intervention. Four mechanisms constitute this layer: data processing and analysis, language understanding, autonomous knowledge navigation and robotic and workflow automation. Together, these mechanisms provide the efficiency and scalability substrate upon which the more complex agentic capabilities depend.

The data processing and analysis involves the automated extraction, organisation and interpretation of the large and varied data sets, converting the raw information into actionable insights for tasks including financial reporting, demand forecasting and compliance monitoring (Hnatyshyn et al., 2025; Sallem et al., 2024).

Language understanding enables the AI systems to interpret and generate natural language supporting customer engagement, internal communication and document processing through conversational agents and NLP-based tools (Adam et al., 2021; Hughes et al., 2025).

Autonomous knowledge navigation extends the informational reach of the AI agents by enabling them to retrieve and synthesise information from distributed sources, such as external databases, regulatory repositories and competitor intelligence platforms, without requiring manual query formulation (Petrenko, 2025; Sifakis et al., 2025).

Robotic and workflow automation translates the AI capabilities into the physical and digital operational execution, encompassing RPA in the back-office functions and the physical robotics in the fulfilment and logistics (Mahler et al., 2019; Nicoletti & Appolloni, 2025).

Strategic Autonomy

Strategic autonomy represents the capacity of the AI agents to engage in goal-directed, adaptive and forward-looking behaviour that goes beyond the execution of predefined processes. This layer encompasses adaptive decision-making, multi-agent coordination, strategic foresight and autonomous market

engagement, and it is at this level that the potential of agentic AI to act as a genuine strategic partner rather than merely an operational tool is most fully realised.

The adaptive decision-making involves the continuous modification of strategies, processes and resource allocations in response to real-time data, enabling e-businesses to respond dynamically to market changes, customer behaviour shifts and competitive pressures (Shili & Anwar, 2025; Wang et al., 2025).

The multi-agent coordination enables the distributed AI systems, with specialised agents managing the supply chains, customer service, financial transactions and analytics, to collaborate through shared protocols, producing the collective intelligence that exceeds the capabilities of any individual system (Pal et al., 2018; Xu et al., 2023).

Strategic foresight equips the AI systems to anticipate future risks and opportunities through predictive analytics and scenario simulation, supporting proactive planning rather than reactive adjustment (Carayannis et al., 2025).

Autonomous market engagement describes the capacity of the AI agents to interact directly with external parties, the customers, suppliers and competitors, negotiating, forming relationships and shaping the market dynamics without requiring human initiation of each interaction (Kwon et al., 2024; Minarsch et al., 2022).

Responsible Autonomy

Responsible autonomy is the governing layer of the RAF. It does not add new AI capabilities to the operational and strategic layers but instead provides the governance architecture that ensures both layers operate in alignment with the organisational values, regulatory requirements and societal expectations. Four mechanisms constitute this layer: explainability, accountability, privacy-by-design and human-in-the-loop governance.

Explainability (XAI) requires that the AI systems be capable of rendering their decision processes interpretable to the relevant stakeholders. In the e-business contexts, this means providing both the users and the managers with transparent accounts of why the product was recommended, the credit decision was reached, or the pricing adjustment was made (Chernyaeva et al., 2025; Sanjammagari et al., 2026).

Accountability establishes the mechanisms through which the responsibility for the AI-generated decisions can be attributed, traced and reviewed, including the audit trails, digital signature systems and the blockchain-based records that provide the permanent, tamper-evident logs of the agent actions (Franklin et al., 2022; Mashatan et al., 2021).

The privacy-by-design embeds the data protection principles into the technical architecture of the AI systems from the outset rather than treating compliance as a post-deployment concern, consistent with the requirements of the GDPR and aligned with the risk-based classification approach of the EU AI Act (Danatzis et al., 2024; Vatankhah et al., 2024).

The human-in-the-loop governance ensures that human judgement remains integral to the high-stakes decisions, providing the oversight thresholds at which the autonomous action is paused and the human review is triggered, thereby preserving the meaningful human agency within the AI-augmented decision processes (Holzinger et al., 2025; Hughes et al., 2025).

Integrated Use Case: Agentic AI Deployment in a Mid-scale E-commerce Firm

To illustrate how the three layers of the RAF operate in practice and how the governance mechanisms of the responsible autonomy layer shape the behaviour of the operational and strategic AI functions, this section presents the integrated use case based on the hypothetical mid-scale e-commerce firm operating across multiple national markets. The scenario is designed to be representative of the deployment challenges faced by organisations in this sector, drawing on the patterns documented in the empirical literature reviewed above. The firm in question deploys the suite of AI agents across its customer engagement, pricing, procurement and compliance functions.

The scenario proceeds through the five stages, each activating the different RAF mechanisms and illustrating the interdependencies between the layers. Table 2 summarises the full scenario; the narrative below provides the analytical context.

Table 2. RAF Use Case: Agentic AI Deployment in an E-commerce Firm.

RAF Layer	Mechanism Activated	System Action	Governance Safeguard	Outcome
Operational autonomy	Data processing and analysis	AI pipeline detects the anomalous drop in the checkout conversions; RPA flags it for review	Audit log generated; data access governed by the Privacy-by-Design protocols	Rapid identification of the suspected fraudulent promotion code
Operational autonomy	Language understanding	NLP-enabled customer service agent handles the surge in complaint messages with personalised responses	HITL override available; escalation threshold triggers the human agent review for the complex cases	Customer queries resolved without the human bottleneck; escalations logged
Strategic autonomy	Adaptive decision-making	AI dynamically adjusts the pricing and inventory allocation based on the real-time demand signals	XAI module generates the explanation of the pricing decision for the compliance review	Revenue optimised; pricing decisions traceable and auditable
Strategic autonomy	Multi-agent coordination	Procurement agent renegotiates the supplier contracts autonomously as the supply chain disruption is detected	Accountability framework assigns the decision responsibility; blockchain records the contract modifications	Supply continuity maintained; agent actions documented and reversible
Responsible autonomy	Explainability (XAI) + HITL	Senior manager reviews the AI-generated explanation of the high-value credit decision and exercises override	Human override exercised; decision rationale stored for the regulatory submission	Regulatory compliance maintained; customer trust preserved

In Stage 1, the operational anomaly detection system drawing on the data processing and analysis mechanism identifies the statistically unusual decline in the checkout conversions over the 4-hour window. The system flags the pattern consistent with the misuse of the promotional discount code and generates the incident report automatically logged to the audit trail governed by the privacy-by-design protocols to ensure that the customer data involved is handled in compliance with the GDPR data minimisation requirements. The efficiency gain here is clear: the anomaly is identified within minutes rather than the hours or days that manual monitoring would require.

Stage 2 illustrates the language understanding mechanism operating within the responsible autonomy guardrails. As the promotional anomaly generates customer complaints, the NLP-enabled customer service agent processes the incoming messages and provides personalised responses resolving the majority without human escalation. The responsible autonomy layer requires, however, that the human-in-the-loop threshold be set: any complaint involving the claimed financial loss above the specified value is automatically escalated to the human agent with the AI draft response presented alongside the relevant context. This design preserves operational efficiency while ensuring that the consequential decisions retain human oversight.

In Stage 3, the strategic autonomy layer is activated as the adaptive decision-making agent receives the anomaly report and recalibrates the firm's dynamic pricing model. Drawing on the real-time demand data and the identified fraud pattern, the agent adjusts the prices across the affected product categories and reallocates the inventory to mitigate the revenue impact. Critically, the explainability mechanism of the responsible autonomy layer requires the agent to generate a structured explanation of each pricing adjustment specifying the variables that drove the decision, the alternatives considered and the expected revenue impact for review by the compliance function. This explanatory output serves both internal governance purposes and satisfies the interpretability requirements that would apply under the EU AI Act's high-risk categorisation of the AI systems used in the commercial decision-making.

Stage 4 shows the multi-agent coordination under the accountability framework. As the anomaly investigation shows that the fraudulent promotional code originated with the compromised third-party affiliate, the procurement agent autonomously initiates the contract renegotiation with the affiliate management system, suspending the affiliate's access and triggering the review of the associated supplier agreements. Each autonomous action taken by the procurement agent is recorded on the blockchain-based audit ledger, providing the immutable record of the decision sequence that can be reviewed by the legal counsel, compliance officers or the regulators if required. The accountability framework thus transforms the potential governance failure of the autonomous agent taking the consequential commercial actions without the traceable decision trail into a fully auditable process.

Stage 5 captures the human-in-the-loop mechanism at the strategic level. The senior manager is presented with the AI's recommendation to permanently delist the compromised affiliate and adjust the firm's affiliate programme terms.

The responsible autonomy layer provides the manager with the structured decision support package: the AI's recommendation, the explanatory rationale, the relevant contractual and regulatory context and the summary of the potential financial and reputational consequences of each available option. The manager exercises the human override, modifying the AI's recommendation to include the monitored probationary period rather than the immediate permanent delisting, which is the careful judgement that the AI's optimisation logic had not fully anticipated. This decision is documented, and the AI's learning model is updated to incorporate the revised outcome weighting.

Taken as a whole, this use case shows the three properties of the RAF that distinguish it from the existing governance approaches. First, the responsible autonomy layer operates continuously across all stages, not as the discrete compliance checkpoint but as the embedded governance architecture that shapes how the operational and strategic AI agents are designed and constrained. Second, the governance mechanisms—explainability, accountability, privacy-by-design and the human-in-the-loop are not alternatives but complements: they address the different aspects of the governance challenge and collectively provide the coverage that no single mechanism could offer. Third, the framework does not restrict the AI agency in ways that compromise the operational efficiency or the strategic adaptability; it channels that agency into the forms that are transparent, traceable and amenable to human review at the appropriate thresholds. The RAF's central argument that responsibility and autonomy can reinforce rather than constrain each other is illustrated concretely by the scenario.

Theoretical and Practical Contribution

The RAF's central contribution lies not in introducing a new theoretical perspective but in integrating three previously disconnected streams of literature on operational automation, strategic AI capability, and responsible governance into a single coherent framework. At the theoretical level, this integration is achieved by reconceptualising agentic AI agents through the combined lens of agency theory and sociotechnical systems theory, as entities that exercise delegated authority on behalf of organisational principals while remaining embedded within, and shaped by, the social and institutional structures of the organisations they serve. This synthesis does not displace either theory; it shows how their combination addresses a gap that neither fully resolves alone.

At the framework level, the RAF's primary contribution is the integration. As Table 3 illustrates, no existing framework spans the full range from operational automation through the strategic agentic capabilities to responsible governance. The RAF fills this gap not by adding the fourth dimension to the existing frameworks but by reconceiving the relationship between autonomy and responsibility: rather than treating them as competing values to be balanced, the RAF treats responsible design as the condition of possibility for sustainable autonomy. The AI system that is not accountable, transparent and aligned with human oversight is not fully autonomous in any meaningful sense; it is merely uncontrolled.

Table 3. Mechanisms at Different Layers of the Proposed Responsible Autonomy Framework.

RAF Layer	Mechanism	Supporting Literature
Operational autonomy	Data processing & analysis	Gong, 2025; Hnatyshyn et al., 2025; Murphy et al., 2024; Sallem et al., 2024; Van et al., 2025; Zamain & Subramanian, 2024
	Language understanding	Adam et al., 2021; Duesterwald et al., 2024; Hughes et al., 2025; Van et al., 2025
	Autonomous knowledge navigation	Nicoletti & Appolloni, 2025; Petrenko, 2025; Sifakis et al., 2025
	Robotic/workflow automation	Hughes et al., 2025; Mahler et al., 2019; Nicoletti & Appolloni, 2025
Strategic autonomy	Adaptive decision-making	Kaur & Prashar, 2025; Puhakainen et al., 2003; Shili & Anwar, 2025; Wang et al., 2025
	Multi-agent coordination	Pal et al., 2018; Qiu & Yang, 2021; Xu et al., 2023
	Strategic foresight	Alhyasat et al., 2025; Carayannis et al., 2025; Goncalves et al., 2025; How & Cheah, 2024
	Autonomous market engagement	Kwon et al., 2024; Minarsch et al., 2022; Rosaci & Sarne, 2012; Sümer et al., 2025; Zhao et al., 2025
Responsible autonomy	Explainability (XAI)	Chernyaeva et al., 2025; Jahan et al., 2025; Oprea & Bâra, 2025; Sanjammagari et al., 2026
	Accountability	Chen, 2015; Franklin et al., 2022; Mashatan et al., 2021; Saha et al., 2014; Triantafyllou, 2023; Triantafyllou & Radanovic, 2023
	Privacy-by-design	Bussone et al., 2020; Cuzzocrea & Soufargi, 2025; Danatzis et al., 2024; Patel et al., 2020; Xu et al., 2012
	Human-in-the-loop (HITL)	Chen et al., 2025; Egbengwu et al., 2025; Holzinger et al., 2024, 2025; Mitsiaki et al., 2025; Syed et al., 2025

At the practical level, the RAF provides the managers and policymakers with a structured model for the responsible scaling of agentic AI in digital commerce. The use case presented in Section ‘Integrated Use Case: Agentic AI Deployment in a Mid-scale e-commerce Firm’ shows that the framework is not merely prescriptive in the abstract but actionable in the specific deployment contexts. By mapping the specific governance mechanisms to the specific operational and strategic functions, the RAF reduces the distance between the governance theory and the governance practice, the distance that, as the literature review shows, has been the persistent limitation of the existing approaches.

Conclusion

The RAF has been developed as an integrative conceptual framework advancing the integration of AI in e-business. The framework moves beyond the narrow efficiency-driven automation to articulate the layered progression of operational

autonomy, strategic autonomy and responsible autonomy. The operational autonomy captures the automation of the repetitive and efficiency-oriented processes such as transaction processing, logistics optimisation and conversational interfaces, thereby establishing the foundations of scalability and cost reduction (Duesterwald et al., 2024; Gong, 2025). The strategic autonomy builds on this base by highlighting the proactive, adaptive and collaborative functions of the AI agents, particularly through multi-agent systems (MAS), advanced scenario modelling and autonomous decision-making in fast-moving environments (Huang, 2025; Queiroz et al., 2024). The responsible autonomy layer embeds governance, transparency and ethics directly into the architecture of the autonomy, ensuring that the e-businesses achieve not only the agility and efficiency but also the accountability, fairness and societal trust (Dignum, 2019; Hughes et al., 2025).

By unifying these layers, the RAF contributes to understanding the shift from automation to autonomy in e-business. It shows how the AI agents are no longer passive tools but increasingly act as proactive strategic partners while staying embedded within the ethical and regulatory guardrails. This integration positions the RAF as a marked theoretical and practical advancement in the ongoing digital transformation of commerce. Responsibility is not a constraint on autonomy; it is what makes autonomy sustainable.

Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship and/or publication of this article.

Declaration of Generative AI and AI-assisted Technologies in the Manuscript Preparation Process

During the preparation of this work, the authors used QuillBot and ChatGPT to improve the writing style and to correct grammatical errors. After using these tools, the authors thoroughly reviewed and edited the content as needed and take full responsibility for the content of the published article.

Funding

The authors received no financial support for the research, authorship and/or publication of this article.

ORCID iDs

Mohammad Talha Siddiqui  <https://orcid.org/0000-0001-7165-4317>
Yusuf Kamal  <https://orcid.org/0000-0002-1557-8775>

References

- Adam, M., Wessel, M., & Benlian, A. (2021). AI-based chatbots in customer service and their effects on user compliance. *Electronic Markets*, 31(2), 427–445. <https://doi.org/10.1007/s12525-020-00414-7>
- Alhyasat, W. B. A. K., Alhyasat-Al-Balqa, E. B., & Khattab, S. A. (2025). Technology trends in strategic management in the AI era: Systematic literature review. *Human Systems Management*, 44(5), 764–781. <https://doi.org/10.1177/01672533251322396>

- Al-Mushayt, O. S. (2019). Automating e-government services with artificial intelligence. *IEEE Access*, 7, 146821–146829. <https://doi.org/10.1109/ACCESS.2019.2946204>
- Alshaketheep, K., Mansour, A. M., Al-Ma'aitah, M. M. L., Dabaghia, M. N., & Dabaghie, Y. M. (2024). Leveraging AI predictive analytics for marketing strategy: The mediating role of management awareness. *Journal of System and Management Sciences*, 14(2), 71–89. <https://doi.org/10.33168/JSMS.2024.0205>
- Ansari, M., Ali, S. A., Alam, M., Chaudhary, K., & Rakshit, S. (2023). Unlocking the power of explainable AI to improve customer experiences in e-commerce. In *AI-based data analytics: Applications for business management* (pp. 31–48). CRC Press. <https://doi.org/10.1201/9781032614083-3>
- Ashok, M., Madan, R., Joha, A., & Sivarajah, U. (2022). Ethical framework for artificial intelligence and digital technologies. *International Journal of Information Management*, 62, Article 102433. <https://doi.org/10.1016/j.ijinfomgt.2021.102433>
- Biswas, A., & Talukdar, W. (2025). *Building agentic AI systems: Create intelligent, autonomous AI agents that can reason, plan, and adapt*. Packt Publishing.
- Bohnsack, R., & De Wet, M. (2025). *AI is the strategy: From agentic AI to autonomous business models onto strategy in the age of AI* (arXiv:2506.17339). arXiv. <https://arxiv.org/abs/2506.17339>
- Bussone, A., Kasadha, B., Stumpf, S. C., Durrant, A. C., Tariq, D. S., Gibbs, J. J., Lloyd, K. C., & Bird, J. (2020). Trust, identity, privacy, and security considerations for designing a peer data sharing platform between people living with HIV. In *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2), Article 148. <https://doi.org/10.1145/3415244>
- Carayannis, E. G., Dumitrescu, R., Falkowski, T., Papamichail, G., & Zota, N. R. (2025). Enhancing SME resilience through artificial intelligence and strategic foresight: A framework for sustainable competitiveness. *Technology in Society*, 81, Article 102835. <https://doi.org/10.1016/j.techsoc.2025.102835>
- Chen, J., Lim, C. P., Tan, K. H., Govindan, K., & Kumar, A. (2025). Artificial intelligence-based human-centric decision support framework: An application to predictive maintenance in asset management under pandemic environments. *Annals of Operations Research*, 350(2), 493–516. <https://doi.org/10.1007/s10479-021-04373-w>
- Chen, R. (2015). Autonomous tracing system for backward design in food supply chain. *Food Control*, 51, 70–84. <https://doi.org/10.1016/j.foodcont.2014.11.004>
- Chernyaeva, O., Lee, O. K. D., & Hong, T. (2025). From fake to AI-generated: Leveraging information manipulation theory and explainable AI for robust detection of manipulated reviews. *Decision Support Systems*, 197, Article 114522. <https://doi.org/10.1016/j.dss.2025.114522>
- Cuzzocrea, A., & Soufargi, S. (2025). AB-DOM: An algorithmic framework for supporting privacy-preserving big data publishing in big data lakes. *IEEE Transactions on Big Data*. Advance online publication. <https://doi.org/10.1109/TBDATA.2025.3570081>
- Danatzi, I., Chandler, J. D., Akaka, M. A., & Ng, I. C. L. (2024). Designing digital platforms for social justice: Empowering end users through the Dataswyft platform. *MIS Quarterly*, 48(4), 1771–1802. <https://doi.org/10.25300/MISQ/2024/18334>
- Dignum, V. (2019). Ensuring responsible AI in practice. In *Responsible artificial intelligence: How to develop and use AI in a responsible way* (pp. 93–105). Springer. https://doi.org/10.1007/978-3-030-30371-6_6
- Duesterwald, E., Isahagian, V., Jayaram, K. R., Kumar, R., Muthusamy, V., Oum, P., Thomas, G., & Venkateswaran, P. (2024). A conversational assistant framework for automation. In *Proceedings of the Middleware Industrial Track 2024* (pp. 1–7). Association for Computing Machinery. <https://doi.org/10.1145/3700824.3701093>

- Egbengwu, V., Garn, W., & Turner, C. J. (2025). Metaverse for manufacturing: Leveraging extended reality technology for human-centric production systems. *Sustainability*, 17(1), Article 280. <https://doi.org/10.3390/su17010280>
- Franklin, M., Ashton, H., Awad, E., & Lagnado, D. A. (2022). Causal framework of artificial autonomous agent responsibility. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 276–284). Association for Computing Machinery. <https://doi.org/10.1145/3514094.3534140>
- Goncalves, R. P., Spaniol, M. J., Rowland, N. J., & Rytter, N. G. M. (2025). Reflections on building an artificial intelligence bot to prepare students to engage in strategic conversations during foresight fieldwork. *Futures and Foresight Science*, 7(1), Article e202. <https://doi.org/10.1002/ffo2.202>
- Gong, S. (2025). Digital transformation of supply chain management in retail and e-commerce. *International Journal of Retail & Distribution Management*, 53(2), 1–19. <https://doi.org/10.1108/IJRDM-02-2023-0076>
- Gregor, S. (2006). The nature of theory in information systems. *MIS Quarterly*, 30(3), 611–642. <https://doi.org/10.2307/25148742>
- Hnatyshyn, L., Prokopyshyn, O., Maletska, O., Keleberda, T., & Pylypenko, K. (2025). Digital innovations in accounting as economic growth factors of an enterprise. *Scientific Bulletin of Mukachevo State University. Series Economics*, 12(1), 75–89. <https://doi.org/10.52566/msu-econ1.2025.75>
- Holzinger, A., Fister, I., Fister, I., Kaul, H. P., & Asseng, S. (2024). Human-centered AI in smart farming: Toward Agriculture 5.0. *IEEE Access*, 12, 62199–62214. <https://doi.org/10.1109/ACCESS.2024.3395532>
- Holzinger, A., Zatloukal, K., & Muller, H. (2025). Is human oversight to AI systems still possible? *New Biotechnology*, 85, 59–62. <https://doi.org/10.1016/j.nbt.2024.12.003>
- How, M., & Cheah, S. M. (2024). Forging the future: Strategic approaches to quantum AI integration for industry transformation. *AI*, 5(1), 290–323. <https://doi.org/10.3390/ai5010015>
- Huang, K. (2025). AI agents and business workflow. In *Agentic AI: Theories and practices* (pp. 135–166). Springer. https://doi.org/10.1007/978-3-031-90026-6_5
- Hughes, L., Dwivedi, Y. K., Malik, T., Shawosh, M., Albashrawi, M. A., Jeon, I., Dutot, V., Appanderanda, M., Crick, T., De', R., Fenwick, M., Gunaratnege, S. M., Jurcys, P., Kar, A. K., Kshetri, N., Li, K., Mutasa, S., Samothrakis, S., Wade, M., & Walton, P. (2025). AI agents and agentic systems: A multi-expert analysis. *Journal of Computer Information Systems*, 65(4), 489–517. <https://doi.org/10.1080/08874417.2025.2483832>
- Jahan, N., Ahamed, J., & Nandi, D. (2025). Enhancing e-commerce sentiment analysis with advanced BERT techniques. *International Journal of Information Engineering and Electronic Business*, 17(3), 49–61. <https://doi.org/10.5815/ijieeb.2025.03.04>
- Jensen, M. C., & Meckling, W. H. (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4), 305–360. [https://doi.org/10.1016/0304-405X\(76\)90026-X](https://doi.org/10.1016/0304-405X(76)90026-X)
- Kaur, A., & Prashar, D. (2025). Web scraping for product recommendations: A review of techniques and applications. *Journal of Computer Science*, 21(6), 1425–1439. <https://doi.org/10.3844/jcssp.2025.1425.1439>
- Kottur, R. (2024). Responsible AI development: A comprehensive framework for ethical implementation in contemporary technological systems. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 10(6), 1553–1561. <https://doi.org/10.32628/cseit241061197>

- Kwon, D., Weiss, E., Kulshrestha, T., Chawla, K., Lucas, G. M., & Gratch, J. (2024). Are LLMs effective negotiators? Systematic evaluation of the multifaceted capabilities of LLMs in negotiation dialogues. In *Findings of the Association for Computational Linguistics: EMNLP 2024* (pp. 5391–5413). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.findings-emnlp.310>
- Liu, X., Wu, L., Yu, J., & Lei, X. (2010). A holistic governance framework for e-business success. In *Proceedings of the 2010 International Conference on Management of e-Commerce and e-Government* (pp. 142–146). IEEE. <https://doi.org/10.1109/ICMcCG.2010.36>
- MacInnis, D. J. (2011). A framework for conceptual contributions in marketing. *Journal of Marketing*, 75(4), 136–154. <https://doi.org/10.1509/jmkg.75.4.136>
- Mahler, J., Matl, M., Satish, V., Danielczuk, M., DeRose, B., McKinley, S., & Goldberg, K. (2019). Learning ambidextrous robot grasping policies. *Science Robotics*, 4(26), Article eaau4984. <https://doi.org/10.1126/scirobotics.aau4984>
- Manda, V. K., Christy, V., & Rao, J. M. (2024). Ethical AI and decision-making in management leadership. In *Ethical dimensions of AI development* (pp. 197–226). IGI Global. <https://doi.org/10.4018/979-8-3693-4147-6.ch009>
- Mashatan, A. A., Lemieux, V. L., Lee, S., Szufel, P., & Roberts, Z. (2021). Usurping double-ending fraud in real estate transactions via blockchain technology. *Journal of Database Management*, 32(1), 27–48. <https://doi.org/10.4018/JDM.2021010102>
- Minarsch, D., Favorito, M., Hosseini, S. A., Turchenkov, Y., & Ward, J. (2022). Autonomous economic agent framework. In *Engineering Multi-Agent Systems* (Lecture Notes in Computer Science, Vol. 13190, pp. 237–253). Springer. https://doi.org/10.1007/978-3-030-97457-2_14
- Mitsiaki, A., Tsita, C., Dimitriou, N., & Tzovaras, D. (2025). A human-centric decision support system for zero-defect manufacturing enabled by human-in-the-loop learning. *Communications in Computer and Information Science*, 2582, 3–16. https://doi.org/10.1007/978-3-031-96199-1_1
- Munoko, I., Brown-Liburd, H. L., & Vasarhelyi, M. (2020). The ethical implications of using artificial intelligence in auditing. *Journal of Business Ethics*, 167(2), 209–234. <https://doi.org/10.1007/s10551-019-04407-1>
- Murphy, B., Feeney, O., Rosati, P., & Lynn, T. (2024). Exploring accounting and AI using topic modelling. *International Journal of Accounting Information Systems*, 55, Article 100709. <https://doi.org/10.1016/j.accinf.2024.100709>
- Nicoletti, B., & Appolloni, A. (2025). A digital twin framework for enhancing human-agent AI-machine collaboration. *Journal of Intelligent Manufacturing*. Advance online publication. <https://doi.org/10.1007/s10845-025-02637-x>
- Oprea, S. V., & Băra, A. (2025). Diverse counterfactual explanations (DiCE) role in improving sales and e-commerce strategies. *Journal of Theoretical and Applied Electronic Commerce Research*, 20(2), Article 96. <https://doi.org/10.3390/jtaer20020096>
- Orlikowski, W. J. (1992). The duality of technology: Rethinking the concept of technology in organizations. *Organization Science*, 3(3), 398–427. <https://doi.org/10.1287/orsc.3.3.398>
- Pal, G., Li, G., & Atkinson, K. M. (2018). Multi-agent big-data lambda architecture model for e-commerce analytics. *Data*, 3(4), Article 58. <https://doi.org/10.3390/data3040058>
- Patel, F. B., Jain, N., Menon, R., & Kodeboyina, S. (2020). Comparative study of privacy preserving-contact tracing on digital platforms. In *Proceedings of the 2020*

- International Conference on Communication and Information Processing* (pp. 137–141). IEEE. <https://doi.org/10.1109/ICC151257.2020.9247782>
- Petrenko, A. I. (2025). Agent-based approach to implementing artificial intelligence (AI) in service-oriented architecture (SOA). *System Research and Information Technologies*, 2025(1), 104–123. <https://doi.org/10.20535/SRIT.2308-8893.2025.1.08>
- Puhakainen, P., Candolin, C., & Kari, H. H. (2003). Using adaptive decision making based on incomplete trust in electronic commerce. In *Proceedings of the IASTED International Conference on Communication, Network, and Information Security* (pp. 191–194). ACTA Press.
- Qiu, L., & Yang, L. (2021). The research of co-evolution mechanisms between cross-border e-commerce and manufacturing cluster: An agent-based model. *E3S Web of Conferences*, 235, Article 03048. <https://doi.org/10.1051/e3sconf/202123503048>
- Queiroz, M., Anand, A., & Baird, A. (2024). Manager appraisal of artificial intelligence investments. *Journal of Management Information Systems*, 41(3), 682–707. <https://doi.org/10.1080/07421222.2024.2376383>
- Ramon-Poma, G. M., Macias-Cabrera, A. G., Yamil, A. A. P., & Jacobo, A. P. D. (2024). Impact of artificial intelligence on the accounting profession. In *Development of scientific research in the social sciences: An approach from the Latin American university* (Vol. 238, pp. 360–386). <https://doi.org/10.53486/issc2023.29>
- Rosaci, D., & Sarnè, G. M. L. (2012). A multi-agent recommender system for supporting device adaptivity in e-commerce. *Journal of Intelligent Information Systems*, 38(2), 393–418. <https://doi.org/10.1007/s10844-011-0160-9>
- Saha, G. K., Desai, M., Ghosh, A., & Saha, N. (2014). Digital signature modeling in e-business. In *Proceedings of the 11th IEEE International Conference on e-Business Engineering* (pp. 350–354). IEEE. <https://doi.org/10.1109/ICEBE.2014.67>
- Sallem, N. R. M., Hussain, N. H. C., Muhmad, S. N., Adnan, N. S., & Halmi, S. A. H. (2024). Artificial intelligence (AI) revolution in accounting and auditing field: A bibliometric analysis. *Advances in Social Sciences Research Journal*, 11(9.2), 64–78. <https://doi.org/10.14738/assrj.119.2.17402>
- Sanjamagari, H. S. G., Sougandhika, G. S., & Srinivasa Desikan, K. E. (2026). Bridging the trust gap: Leveraging explainable AI for personalized e-commerce recommendations. In *Lecture Notes in Computer Science* (Vol. 15706, pp. 225–236). Springer. https://doi.org/10.1007/978-981-96-8889-0_19
- Shankar, V. (2024). Managing the twin faces of AI: A commentary on ‘Is AI changing the world for better or worse?’ *Journal of Macromarketing*, 44(4), 892–899. <https://doi.org/10.1177/02761467241286483>
- Shili, M., & Anwar, S. (2025). Real-time electrical energy optimization in e-commerce systems based on IoT and mobile agents. *Information*, 16(7), Article 551. <https://doi.org/10.3390/info16070551>
- Sifakis, J., Li, D., Huang, H., Zhang, Y., Dang, W., Huang, R., & Yu, Y. (2025). *A reference architecture for autonomous networks: An agent-based approach*. arXiv. <https://arxiv.org/abs/2503.12871>
- Sümer, M. S., Özsoy, T., Yildirim Okay, F., & Kok, I. (2025). Smart agents as customers: A semi-autonomous digital commerce model for Industry 5.0. In *Proceedings of the 2025 International Conference on Human-Oriented Robotics and Automation*. IEEE. <https://doi.org/10.1109/ICHORA65333.2025.11017052>
- Syed, D. A., Quadri, W., Rahmani Choubbeh, N., Pinzone, M., & Gusmeroli, S. (2025). Approaching interoperability and data-related processing issues in a human-centric industrial scenario. *Communications in Computer and Information Science*, 2328, 21–34. https://doi.org/10.1007/978-3-031-78572-6_2

- Thapliyal, K., & Thapliyal, M. (2024). Chatbot-XAI: The new age artificial intelligence communication tool for e-commerce. *Studies in Computational Intelligence*, 1094, 77–100. https://doi.org/10.1007/978-3-031-55615-9_6
- Triantafyllou, S. (2023). Forward-looking and backward-looking responsibility attribution in multi-agent sequential decision making. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems: Doctoral Consortium* (pp. 2952–2954). International Foundation for Autonomous Agents and Multiagent Systems.
- Triantafyllou, S., & Radanovic, G. (2023). Towards computationally efficient responsibility attribution in decentralized partially observable MDPs. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems* (pp. 131–139). International Foundation for Autonomous Agents and Multiagent Systems.
- Trist, E. L., & Bamforth, K. W. (1951). Some social and psychological consequences of the Longwall method of coal-getting: An examination of the psychological situation and defences of a work group at the coal face. *Human Relations*, 4(1), 3–38. <https://doi.org/10.1177/001872675100400101>
- Van, N. D., Hoang, C. C., & Khoa, B. T. (2025). Bibliometric examination of artificial intelligence within the framework of e-commerce technology from 1996 to 2024. *Journal of Logistics, Informatics and Service Science*, 12(2), 138–150. <https://doi.org/10.33168/JLISS.2025.0209>
- Vannostrand, P. M., Hofmann, D. M., Ma, L., & Rundensteiner, E. A. (2024). Actionable recourse for automated decisions: Examining the effects of counterfactual explanation type and presentation on lay user understanding. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1682–1700). Association for Computing Machinery. <https://doi.org/10.1145/3630106.3658997>
- Vatankhah, S., Bamshad, V., Arici, H. E., & Duan, Y. (2024). Ethical implementation of artificial intelligence in the service industries. *The Service Industries Journal*, 44(9–10), 661–685. <https://doi.org/10.1080/02642069.2024.2359077>
- Wang, T. C., Guo, R.-S. A., Chen, C., & Li, C. (2025). Multi-stage data-driven framework for customer journey optimization and operational resilience. *Mathematics*, 13(7), Article 1145. <https://doi.org/10.3390/math13071145>
- Wirtz, B. W., Weyerer, J. C., & Kehl, I. (2022). Governance of artificial intelligence: A risk and guideline-based integrative framework. *Government Information Quarterly*, 39(4), Article 101685. <https://doi.org/10.1016/j.giq.2022.101685>
- Xu, H., Crossler, R. E., & Belanger, F. (2012). A value sensitive design investigation of privacy enhancing tools in web browsers. *Decision Support Systems*, 54(1), 424–433. <https://doi.org/10.1016/j.dss.2012.06.003>
- Xu, L., Proselkov, Y., Schoepf, S., Minarsch, D., Minaricova, M., & Brintrup, A. (2023). Implementation of autonomous supply chains for digital twinning: A multi-agent approach. *IFAC-PapersOnLine*, 56(2), 11076–11081. <https://doi.org/10.1016/j.ifacol.2023.10.812>
- Zamain, N. S. A., & Subramanian, U. (2024). The impact of artificial intelligence in the accounting profession. *Procedia Computer Science*, 238, 849–856. <https://doi.org/10.1016/j.procs.2024.06.102>
- Zhao, L., Cai, L., & Wang, J. (2025). Federated learning-based autonomous economic agent for portfolio allocation. *Mathematics*, 13(4), Article 427. <https://doi.org/10.3390/math13040427>