Available online www.ejaet.com

European Journal of Advances in Engineering and Technology, 2025, 12(6):108-115

Research Article ISSN: 2394 - 658X

Agentic Intelligence in Information Management Systems: A Framework for Autonomous Decision Workflows

Sudhir Vishnubhatla

Senior Technical Lead- Tampa, USA

ABSTRACT

The evolution of Information Management Systems (IMS) from static repositories to autonomous, cognitive ecosystems represents a major milestone in enterprise intelligence. This paper introduces an agentic framework for next-generation IMS, integrating Intelligent Document Processing (IDP), multi-agent orchestration, and autonomous decision-making. By embedding reasoning, collaboration, and learning into document-centric workflows, agentic IMS enable contextual understanding and real-time decision execution across enterprise functions. The framework employs a layered architecture that connects document ingestion, agentic orchestration, and autonomous decisioning, supported by Large Language Models (LLMs) for semantic reasoning and continuous learning. Key use cases contract management, financial document automation, and regulatory compliance demonstrate tangible benefits in efficiency, transparency, and adaptability. The discussion further explores explainability, scalability, and ethical governance as critical dimensions of trustworthy automation. Future directions include hybrid human–agent collaboration, edge-deployed intelligence, and interoperable standards for agentic ecosystems. The study concludes that agentic IMS can transform enterprises into self-optimizing, ethically governed, knowledge-driven organizations where data, context, and autonomy converge to redefine decision intelligence.

Keywords: AI Agents, Intelligent Document Processing (IDP), Information Management Systems (IMS), Autonomous Decisioning; Agentic AI, Multi-Agent Systems, Workflow Automation, Enterprise Agents, Document-to-Decision, Human-in-the-Loop.

INTRODUCTION

Information Management Systems (IMS) have long served as the digital backbone of enterprise operations, designed primarily for storing, retrieving, and classifying documents. Their evolution has historically mirrored organizational needs from simple file repositories to structured databases supporting enterprise workflows. Yet, as the scale and complexity of business information have expanded, these traditional systems have increasingly struggled to meet modern demands for scalability, contextual understanding, and real-time decision support. The exponential growth of unstructured and semi-structured content, including emails, reports, contracts, and sensor-generated logs, has rendered conventional IMS architectures insufficient for dynamic information environments.

The integration of Intelligent Document Processing (IDP) technologies marked a major milestone in this evolution. IDP systems combine optical character recognition (OCR), natural language processing (NLP), and machine learning (ML) to automate document extraction, classification, and metadata tagging. While these capabilities significantly reduced manual intervention, their primary function remained reactive, focused on interpreting documents rather than understanding broader organizational contexts. As enterprises pursue higher degrees of automation, the need for systems capable of reasoning, coordination, and adaptive learning has become increasingly evident.

This transformation has given rise to the emergence of AI-driven agents within IMS. These intelligent agents move beyond static automation by integrating perception (data acquisition and contextual awareness), reasoning (decision-making based on goals and constraints), and learning (continuous improvement through feedback). They can perform complex, interrelated tasks such as contract validation, compliance checks, and workflow orchestration, often in collaboration with other agents or human users. This represents a fundamental shift in system behavior from task execution to autonomous decision orchestration.

In this context, IMS are evolving into agentic ecosystems, where distributed AI agents interact to manage information flows, infer intent, and coordinate decisions across departments. Such systems are not merely repositories of enterprise knowledge but dynamic entities that interpret, learn, and act upon data in real time. This paradigm aligns closely with the broader movement toward multi-agent systems (MAS) and cognitive enterprise architectures, which emphasize decentralization, adaptability, and contextual intelligence.

This paper examines the architectural frameworks, enabling technologies, and challenges involved in transitioning from document-centric automation to agentic decision-making platforms. By analyzing the convergence of IDP, large language models (LLMs), and multi-agent collaboration, it highlights how IMS are being redefined as autonomous knowledge infrastructures capable of real-time reasoning and decision execution. The discussion also addresses the organizational and ethical implications of this shift, paving the way for the next generation of self-managing enterprise systems.

BACKGROUND & RELATED WORK

The evolution of Information Management Systems (IMS) toward intelligent, autonomous decision-making platforms is deeply rooted in the parallel development of Intelligent Document Processing (IDP) and Multi-Agent Systems (MAS). Early IMS frameworks were primarily concerned with document storage, indexing, and retrieval. However, as enterprise data grew exponentially, these systems needed to interpret and act upon content rather than merely store it. This demand gave rise to IDP technologies, which combined Optical Character Recognition (OCR), Natural Language Processing (NLP), and Machine Learning (ML) to extract structured information from unstructured sources such as invoices, contracts, and forms.

While IDP has been highly effective for automating repetitive document tasks, its scope has remained largely procedural. Systems could classify, tag, and extract entities with impressive precision, but they lacked the ability to reason about the extracted information or make context-aware decisions. For instance, an IDP pipeline might identify and digitize a contract's terms but would not autonomously assess its compliance with internal policies or external regulations. The next logical evolution, therefore, involves embedding reasoning and decision capabilities within IMS, giving rise to agentic document intelligence.

The concept of Multi-Agent Systems (MAS) provides the theoretical foundation for this transformation. MAS research focuses on distributed systems composed of autonomous entities or "agents" that perceive their environment, reason about it, and act to achieve defined goals. Li et al. (2024) explored how the integration of Large Language Models (LLMs) enhances reasoning and contextual adaptability in such multi-agent workflows. Their findings indicate that agents equipped with LLM-based cognitive capabilities can dynamically allocate tasks, interpret ambiguous data, and coordinate decisions across distributed systems.

Zhang et al. (2025) offered a comprehensive analysis of multi-agent collaboration frameworks, emphasizing coordination strategies, role assignment, and inter-agent communication. Their research highlighted that efficient cooperation among agents is vital to achieving scalability and reliability in complex enterprise environments. Similarly, Sun et al. (2025) examined the ongoing challenges of coordination, autonomy, and interpretability within agent-based ecosystems, underscoring the trade-offs between self-directed behavior and centralized control mechanisms.

Within IMS, the application of these ideas is still emerging. Current systems incorporating agentic components primarily use AI-driven assistants for document routing, data validation, or contextual summarization. However, true multi-agent reasoning integration where agents interpret, deliberate, and act collectively within document-centric workflows remains at a nascent stage. This gap represents both a research and implementation opportunity. Developing architectures that enable communication between document agents, decision agents, and governance agents could redefine IMS as a network of collaborative decision-makers rather than isolated automation modules.

Another significant area of study involves governance and explainability. As IMS evolve into multi-agent ecosystems, ensuring accountability and transparency becomes crucial. Existing literature (Dignum, 2019; Russell, 2021) emphasizes the importance of embedding ethical and interpretability frameworks within autonomous systems. In the context of enterprise information management, this translates to mechanisms that allow traceability of agentic decisions, auditability of automated actions, and human oversight in high-stakes operations.

In summary, the convergence of IDP, MAS, and LLM-based reasoning marks the foundation of the next generation of Information Management Systems. While current research demonstrates promising advances in individual areas, the integration of multi-agent intelligence into enterprise-scale IMS remains an open research frontier. Addressing this integration while maintaining scalability, interpretability, and compliance will be key to realizing truly cognitive, self-managing enterprise ecosystems.

AGENTIC ARCHITECTURE & WORKFLOW PATTERNS

The emergence of agentic Information Management Systems (IMS) represents a major step forward in the automation and cognitive transformation of enterprise operations. These systems move beyond static, rule-based automation to dynamic environments where autonomous agents collaborate, reason, and act in real time. A modern

AI-augmented IMS operates through a multi-layered architecture that integrates document understanding, cognitive reasoning, and autonomous decision-making within a unified framework.

At the highest level, the architecture can be divided into three primary tiers:

1. Document Ingestion & Processing
2. Agentic Orchestration
3. Autonomous Decisioning

Each tier builds upon the capabilities of the preceding one, creating a continuous data-to-decision pipeline that blends human insight with machine autonomy.

1. Document Ingestion & Processing Layer

At the foundation lies the document ingestion and processing layer, responsible for converting raw enterprise documents into structured, machine-readable information. This process typically involves Intelligent Document Processing (IDP) components such as Optical Character Recognition (OCR), Natural Language Processing (NLP), and Entity Recognition models.

In this layer, heterogeneous inputs — contracts, invoices, reports, or emails — are parsed, classified, and semantically enriched. IDP pipelines apply contextual tagging to identify key metadata such as customer identifiers, transaction values, or compliance clauses. Figure 2 illustrates this process, showing how raw documents transition through recognition, classification, and extraction phases to produce normalized datasets ready for higher-level reasoning.

While traditional systems stop at extraction, modern IMS architectures treat this layer as a knowledge foundation. The extracted entities and relationships become inputs to agentic reasoning modules, enabling cross-referencing with enterprise knowledge graphs, databases, and external APIs. This allows the system to infer meaning rather than merely process text, marking the first step toward intelligence in information management.

2. Agentic Orchestration Layer

The second layer, agentic orchestration, introduces multi-agent coordination for contextual reasoning and task management. Here, intelligent agents function as autonomous software entities each specializing in distinct cognitive or operational domains. Examples include:

- A Classification Agent that routes documents to appropriate workflows.
- A Compliance Agent that evaluates extracted content against policy or regulatory constraints.
- A Decision Agent that synthesizes contextual signals to recommend or execute actions.

These agents communicate through shared message buses or reasoning frameworks, exchanging data and insights dynamically. Figure 3 depicts this orchestration, illustrating how multiple agents operate in a federated structure, interacting through protocols that ensure coordination, fault tolerance, and adaptive decision-making.

A key advantage of agentic orchestration is scalability and flexibility. Unlike static workflows, agent-based systems can reconfigure themselves dynamically. When a new document type, regulation, or data source is introduced, the orchestration layer assigns new roles or re-trains existing agents to handle the change eliminating the need for manual reprogramming. The system's adaptability makes it particularly valuable in sectors where policies or compliance frameworks evolve rapidly, such as finance, healthcare, or government operations.

3. Autonomous Decisioning Layer

The final tier, autonomous decisioning, represents the operational intelligence of the IMS. At this stage, agents not only analyze but also act initiating business processes, approving transactions, or generating alerts based on confidence thresholds and learned patterns.

This layer relies heavily on policy-driven governance and reinforcement learning to maintain a balance between automation and human oversight. Actions are guided by decision policies encoded within the system, ensuring compliance while allowing flexibility. Over time, feedback from users and outcomes continuously refines these models, leading to improved accuracy and reduced intervention.

For example, in a document-heavy financial workflow, the system might autonomously verify loan applications, cross-check credit data, and flag anomalies for human review. The result is a closed loop learning system that continually enhances its performance through observation and correction.

Together, these three layers form a cohesive, adaptive, and explainable architecture that defines the next generation of Information Management Systems. The shift from static workflows to agentic orchestration introduces continuous learning and reasoning as core architectural principles. Each agent acts as both a decision node and a knowledge contributor, enabling the IMS to evolve alongside organizational needs.

The figures collectively represent this progression:

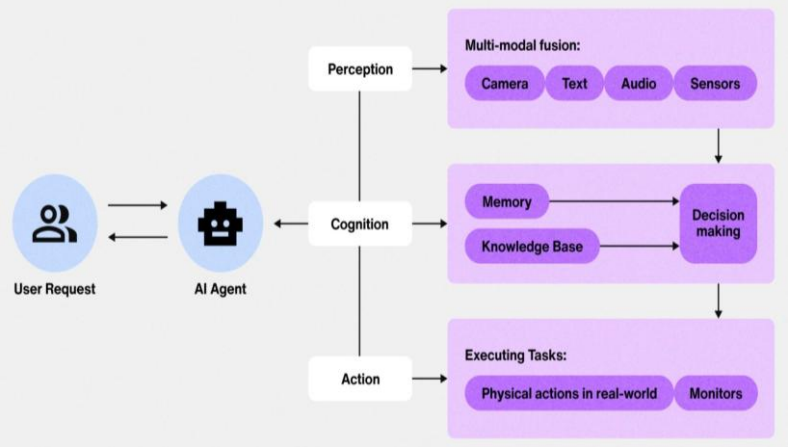

Figure 1: illustrates the overall paradigm shift from traditional automation to agentic workflows.

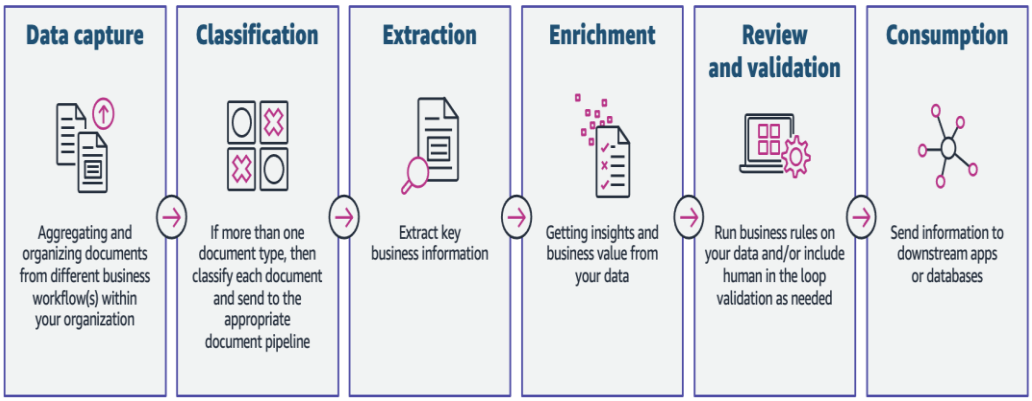

Figure 2: Details the IDP foundation that feeds intelligent data into higher reasoning layers.

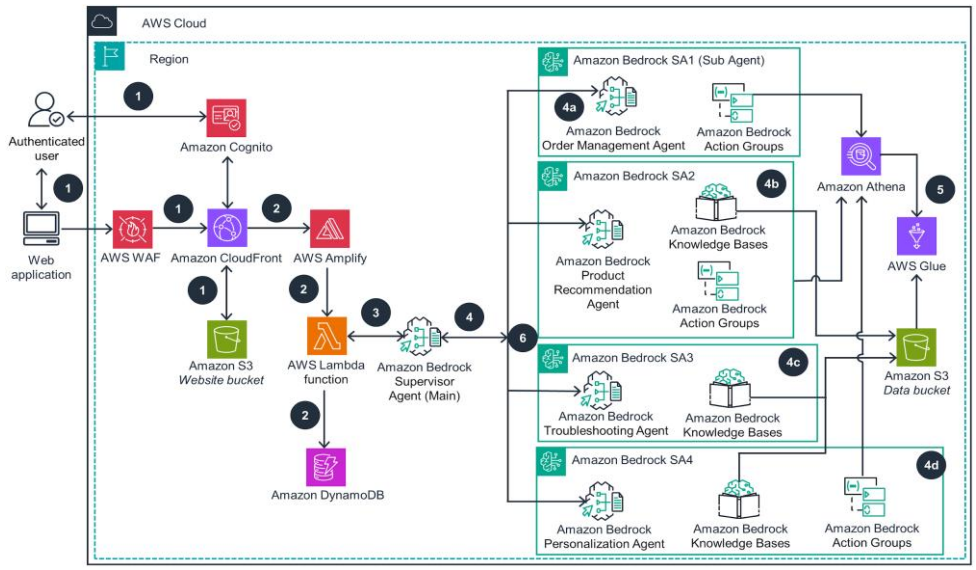

Figure 3: Demonstrates the orchestration of agents for collaborative, autonomous decision-making.

This layered model not only enhances efficiency and accuracy but also lays the groundwork for self-managing enterprise ecosystems, where IMS function as digital collaborators interpreting, reasoning, and acting in alignment with human and organizational objectives.

USE-CASES IN IMS

The integration of AI agents into Information Management Systems (IMS) is fundamentally reshaping how enterprises handle documents, workflows, and decision-making. By embedding cognitive and autonomous capabilities within traditional document ecosystems, these systems are moving from passive information storage to

active operational intelligence. Three critical application areas contract management, financial document automation, and regulatory compliance demonstrate how agentic IMS architectures drive measurable business impact.

1. Contract Management

Contract management has traditionally been one of the most labor-intensive components of enterprise information handling. Large organizations manage thousands of contracts simultaneously, covering vendors, clients, service providers, and internal departments. Manual review processes are prone to delays, inconsistencies, and compliance gaps.

AI agents embedded in IMS streamline this process by automating the entire contract lifecycle from document ingestion to renewal. Intelligent agents equipped with Natural Language Understanding (NLU) and contextual reasoning can extract key clauses such as payment terms, liability limits, and termination conditions. Once extracted, these clauses are cross-referenced against corporate standards or legal frameworks to identify deviations or risks.

Agents also monitor critical deadlines and obligations by integrating with enterprise scheduling systems. For instance, an agent may automatically alert the legal team of an upcoming contract renewal or escalate a contract with non-compliant clauses for human review. Over time, feedback from human interventions trains these agents to refine their evaluation logic, enabling adaptive contract governance. The result is a self-regulating contract management ecosystem that minimizes risk, ensures compliance, and significantly reduces administrative overhead.

2. Financial Document Automation

In financial operations, the accuracy and speed of document processing directly influence decision-making and risk control. AI agents in IMS enhance financial workflows by automating data extraction, validation, and anomaly detection across invoices, purchase orders, and expense reports.

According to Jeong et al. (2025), integrating agentic reasoning into financial document pipelines enables systems to identify subtle irregularities that traditional rule-based systems often miss. For instance, agents can flag inconsistencies between purchase order amounts and invoice totals, detect duplicate expense submissions, or recognize suspicious transaction patterns that may indicate fraud.

Furthermore, quantum-inspired and statistical learning techniques are being incorporated into these financial agents to handle dynamic, real-time data streams. This allows for continuous risk scoring and trend analysis rather than post-hoc audits. When connected with enterprise resource planning (ERP) systems, agents can autonomously trigger corrective actions such as holding payments pending verification or alerting compliance officers of anomalies.

By integrating reasoning capabilities into financial document workflows, IMS evolve into proactive decision-support systems that combine accuracy, transparency, and efficiency in financial management.

3. Regulatory Compliance and Audit Readiness

Regulatory compliance represents one of the most challenging domains for modern enterprises, requiring constant vigilance across diverse legal frameworks and data governance standards. AI agents enhance compliance management by continuously scanning, interpreting, and validating documentation to ensure adherence to policies, standards, and regulations.

Within an agentic IMS, Compliance Agents operate autonomously to track document trails, verify version integrity, and ensure policy conformance. They monitor evolving regulatory updates such as GDPR amendments or new financial disclosure mandates and cross-reference these with internal policies. This dynamic monitoring ensures that all enterprise documentation remains audit-ready and aligned with current regulatory expectations.

In more advanced implementations, compliance agents collaborate with Governance Agents that maintain digital ledgers or blockchain-based audit trails. This guarantees traceability for every document interaction and modification, enhancing transparency and accountability. These capabilities are particularly valuable in industries such as banking, healthcare, and insurance, where regulatory oversight is both continuous and stringent.

By embedding compliance reasoning directly into IMS workflows, organizations can achieve a continuous compliance posture reducing risk exposure, minimizing manual audits, and enhancing stakeholder confidence.

CHALLENGES & TRADE-OFFS

Despite their growing adoption and transformative potential, agentic Information Management Systems (IMS) encounter a set of complex challenges that affect scalability, transparency, governance, and security. As enterprises move from document-centric automation toward autonomous, reasoning-driven ecosystems, balancing computational performance with ethical and regulatory obligations becomes increasingly critical.

1. Scalability and System Complexity

One of the foremost challenges in deploying agentic IMS lies in scalability. As organizations integrate multiple AI agents each performing specialized functions such as classification, compliance verification, or anomaly detection the complexity of orchestration grows exponentially. High data velocity, driven by real-time document inflows and continuous learning cycles, can strain system resources and degrade response times.

In multi-agent settings, communication overhead and synchronization latency become bottlenecks. Ensuring consistent state management, workload distribution, and fault tolerance across hundreds or thousands of interacting agents demands advanced coordination frameworks and distributed computing infrastructures. Enterprises must carefully architect modular, load-balanced systems that scale elastically without compromising decision accuracy or stability.

2. Explainability and Transparency

Explainability remains a cornerstone requirement for enterprise-grade IMS, especially in regulated sectors like finance, healthcare, and public administration. As agentic systems become more autonomous, their decisions ranging from document classification outcomes to contract risk assessments must remain interpretable and auditable.

Complex reasoning pipelines powered by large language models (LLMs) and deep learning often behave as black boxes, producing results without easily understandable justifications. This lack of transparency poses significant barriers for audit compliance and stakeholder trust. For instance, a contract may be rejected or flagged by an AI agent without a clear explanation of the criteria used, leading to legal or operational disputes.

Addressing this requires the integration of Explainable AI (XAI) techniques within agentic architectures. Visualization dashboards, reasoning trace logs, and rule-based overlays can provide human reviewers with interpretable insights into how and why a decision was made. The objective is not only to satisfy regulatory standards but to foster human-AI collaboration where explanations reinforce confidence in autonomous outcomes.

3. Ethical Governance and Bias Mitigation

As highlighted by TechRadar (2025), the expansion of agentic systems introduces new challenges in ethical governance and bias control. AI agents trained on enterprise or third-party data may inadvertently learn and propagate biases present in historical records. For example, decision agents assessing vendor contracts or credit applications could exhibit unintended bias based on regional or demographic factors.

To mitigate this, enterprises must establish governance frameworks that enforce fairness, accountability, and human oversight. Periodic model audits, ethical review boards, and algorithmic transparency protocols are crucial to preventing both intentional misuse and unintentional bias propagation. Governance agents dedicated to AI components responsible for monitoring and enforcing ethical standards are emerging as a best practice within large-scale IMS deployments.

Ethical design must also extend to human oversight. Even in highly autonomous systems, human-in-the-loop checkpoints ensure that critical or high-impact decisions remain under human review, maintaining accountability and control.

4. Data Security and Privacy Protection

Given that IMS manages vast volumes of sensitive corporate and personal data, data security is paramount. The distributed and interconnected nature of multi-agent ecosystems increases the attack surface for cyber threats, including unauthorized access, model manipulation, and data leakage.

Advanced encryption, secure agent communication protocols, and zero-trust architectures are necessary to protect document content and metadata across layers of ingestion, reasoning, and decisioning. Additionally, privacy-preserving AI techniques such as differential privacy and federated learning can safeguard data confidentiality while allowing continuous model improvement. These measures help maintain compliance with global standards like GDPR, CCPA, and industry-specific security frameworks.

5. Balancing Speed, Interpretability, and Security

The central trade-off for organizations adopting agentic IMS lies in balancing performance with governance. Systems designed for high-speed automation risk compromising explainability or security, while those emphasizing transparency and compliance may suffer from latency or computational overhead.

Enterprises must therefore adopt adaptive deployment strategies, where agents dynamically adjust their behavior based on context prioritizing speed for routine document processing while enforcing stricter explainability and audit protocols for high-risk or compliance-sensitive tasks. The ability to calibrate these parameters in real time defines the maturity and trustworthiness of next-generation IMS platforms.

FUTURE DIRECTIONS

The evolution of AI agents within Information Management Systems (IMS) is entering a new phase characterized by autonomy, adaptability, and ethical intelligence. As organizations seek to transition from static automation to self-managing digital ecosystems, the next generation of IMS will be defined by advancements in edge deployment, hybrid collaboration, explainable intelligence, and governance standardization. Together, these developments will shape the trajectory of intelligent enterprise systems over the coming decade.

1. Edge Deployment and Distributed Intelligence

One major frontier lies in the deployment of agentic IMS at the network edge. As enterprises handle growing volumes of time-sensitive data, processing information closer to its source will become essential for real-time responsiveness and scalability. Edge-based agents embedded within branch systems, IoT endpoints, or localized

servers will be capable of executing document processing, classification, and compliance checks independently, without relying on centralized infrastructure.

This approach reduces latency and improves resilience in distributed organizations, particularly in sectors such as finance, logistics, and healthcare. Moreover, coupling edge deployment with federated learning enables agents to learn from decentralized data while preserving privacy, ensuring that IMS remain efficient and secure across geographically dispersed environments.

2. Hybrid Human–Agent Collaboration

While automation continues to expand, hybrid human–agent collaboration will remain critical to enterprise decision-making. Instead of replacing human expertise, future IMS will augment it by allowing seamless interaction between human users and autonomous agents.

In this model, agents handle repetitive, data-driven tasks such as classification, extraction, and anomaly detection while humans oversee strategic reasoning, conflict resolution, and ethical evaluation. Human feedback loops will help train and calibrate agent behavior, resulting in continuously improving decision systems. This symbiotic collaboration not only ensures operational reliability but also fosters user trust by keeping human judgment central to critical processes.

3. Explainable and Ethical AI (XAI) Integration

As AI agents gain autonomy, Explainable AI (XAI) will be indispensable for maintaining accountability, fairness, and regulatory compliance. Future IMS architectures will embed interpretability directly into their reasoning frameworks, allowing stakeholders to understand the logic behind agent decisions.

Explainability tools such as visual trace graphs, causal inference mappings, and natural-language justifications will enable auditors and end-users to evaluate agent decisions transparently. These capabilities will not only strengthen compliance with governance mandates but also enhance the trustworthiness of autonomous systems. In parallel, embedding ethical AI principles including bias detection, consent management, and equitable treatment will ensure that agentic IMS operate responsibly across cultural and legal boundaries.

4. Interoperability and Benchmarking Frameworks

As agentic IMS proliferates across industries, standardization and interoperability will become crucial. Currently, many enterprise systems operate as isolated silos with proprietary communication protocols, limiting collaboration and scalability. Future research will focus on developing open interoperability standards that allow agents from different vendors or domains to communicate seamlessly within multi-enterprise ecosystems.

In addition, benchmarking frameworks will be essential for evaluating agent performance, safety, and alignment. Standardized metrics for autonomy, interpretability, and energy efficiency will help organizations assess the maturity and reliability of agentic systems. These frameworks will also facilitate industry-wide certification processes, like existing ISO and IEEE standards for information systems.

5. Toward Ethical, Context-Aware Ecosystems

Ultimately, the future of agentic IMS lies in the creation of intelligent, context-aware ecosystems capable of adaptive, ethical decision-making. These systems will integrate reasoning with situational awareness, understanding not just what actions to take, but when and why to take them. By combining contextual intelligence with moral and organizational constraints, future IMS will act as responsible digital collaborators that align operational efficiency with human values.

As organizations embrace this vision, they will redefine enterprise information management from passive document handling to autonomous, ethical decision orchestration. The convergence of edge intelligence, hybrid collaboration, explainable AI, and interoperability will position agentic IMS as foundational pillars of the next-generation intelligent enterprise.

CONCLUSION

The emergence of AI agents within Information Management Systems (IMS) marks a transformative evolution from static data repositories to dynamic, self-adaptive ecosystems. These systems no longer serve merely as document management tools but as cognitive infrastructures capable of understanding, reasoning, and executing tasks with minimal human intervention. By embedding agentic intelligence into IMS, enterprises unlock the ability to perform continuous decision-making, automate complex workflows, and enhance the precision and agility of organizational operations.

Agentic IMS redefines the relationship between data and decision-making. Through the integration of Intelligent Document Processing (IDP), multi-agent orchestration, and autonomous reasoning, they bridge the gap between document-centric automation and knowledge-driven action. This shift establishes IMS as active participants in enterprise processes systems that not only store and retrieve information but also interpret context, assess risk, and initiate decisions aligned with strategic goals.

However, the path to widespread adoption demands attention to key pillars: scalability, explainability, and ethical governance. As agentic systems grow in autonomy and scope, transparent reasoning mechanisms and robust oversight become indispensable to ensure accountability and trust. Designing architectures that harmonize

performance with interpretability will determine the long-term viability of these systems in regulated and high-stakes environments.

Looking ahead, agentic IMS will form the core of the intelligent enterprise, supporting adaptive, context-aware decision-making at scale. They will serve as intermediaries between human expertise and digital intelligence augmenting human judgment rather than replacing it. By blending autonomy with oversight and analytics with ethics, AI-driven IMS will not only streamline operations but also shape a new paradigm of responsible, intelligent automation for the data-driven era.

REFERENCES

[1]. Dignum, V. (2019). Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way. Springer. https://doi.org/10.1007/978-3-030-30371-6

[2]. Russell, S. (2022). Artificial Intelligence and the Problem of Control. Perspectives on digital humanism, 19, 1-322.

[3]. Armbrust, M., et al. (2010). A View of Cloud Computing. Communications of the ACM, 53(4), 50–58. https://dl.acm.org/doi/10.1145/1721654.1721672

[4]. Shoham, Y., & Leyton-Brown, K. (2008). Multiagent systems: Algorithmic, game-theoretic, and logical foundations. Cambridge University Press.

[5]. Wooldridge, M. (2021). The Road to Conscious Machines: The Story of AI. Pelican Books.

[6]. Gunning, D., & Aha, D. (2019). DARPA's Explainable Artificial Intelligence (XAI) Program. AI Magazine, 40(2), 44–58. https://onlinelibrary.wiley.com/doi/epdf/10.1609/aimag.v40i2.2850

[7]. Van der Aalst, W. (2016). Process Mining: Data Science in Action. Springer. https://link.springer.com/book/10.1007/978-3-662-49851-4

[8]. Chen, J., & Zhao, Y. (2023). Human–Agent Collaboration Models for Cognitive Enterprises. Knowledge-Based Systems, 284, 111064. https://www.sciencedirect.com/science/article/abs/pii/S0950705123008146?via%3Dihub

[9]. Rahwan, I., Cebrian, M., Obradovich, N., et al. (2019). Machine Behaviour. Nature, 568(7753), 477–486. https://www.nature.com/articles/s41586-019-1138-y

[10]. Crawford, K. (2021). The atlas of AI: Power, politics, and the planetary costs of artificial intelligence. https://yalebooks.yale.edu/book/9780300264630/atlas-of-ai/