



Ontology Driven Data Quality Management for Agentic Ai Systems: A Framework for Intelligent Decision-Making in Enterprise Environments

Dadhimukh Uma Yadav¹

¹PG Scholar, Dept. of Master of Business Administration (MBA), Tulsiramji Gaikwad-Patil College of Engineering and Technology, Mohagaon, Wardha Road, Nagpur, Maharashtra, India.

Email ID: yadavdadhimukh@gmail.com¹

Abstract

The rapid adoption of agentic artificial intelligence systems in enterprise environments has exposed a critical vulnerability: poor data quality undermines AI decision-making at scale. As organizations transition from experimental AI deployments to production-grade agentic systems, the "garbage in, disaster out" phenomenon has emerged as a primary barrier to value realization. This paper investigates the convergence of ontology-based knowledge representation, data quality management, and agentic AI architectures to propose a novel framework for intelligent decision-making in enterprise contexts. Through a systematic review of contemporary literature and industry developments from 2024 to 2026, this research identifies that traditional data quality approaches, designed for business intelligence workflows, prove inadequate for autonomous AI agents operating across distributed data ecosystems. The proposed framework integrates semantic ontology layers with automated data quality monitoring, enabling agents to interpret context, resolve ambiguities, and maintain decision integrity. Findings indicate that organizations implementing ontology-driven data quality management achieve measurable improvements in AI decision accuracy while reducing governance overhead. This paper contributes a conceptual architecture and implementation guidelines relevant to both academic researchers and industry practitioners navigating the agentic AI transition.

Keywords: Agentic AI; Data Governance; Data Quality Management; Intelligent Decision-Making; Ontology; Semantic Integration;

1. Introduction

1.1. Background and Context

The artificial intelligence landscape has undergone a profound transformation between 2024 and 2026. What began as experimental generative AI deployments has matured into what industry analysts now term "agentic AI"—autonomous systems capable of executing multi-step workflows, coordinating with other agents, and making decisions with minimal human intervention. This evolution represents a fundamental shift in how organizations process information and derive value from their data assets. However, this technological progression has exposed a persistent and increasingly urgent challenge: data quality. According to recent industry research, 77% of large enterprises with revenues exceeding \$5 billion anticipate that poor data quality in AI projects will precipitate a major organizational crisis. This statistic underscores a critical reality: the stakes of data quality management have escalated

dramatically. In traditional business intelligence contexts, poor data quality produced incorrect reports frustrating but rarely catastrophic. In agentic AI contexts, poor data quality can lead autonomous systems to make flawed strategic recommendations, execute incorrect transactions, or propagate errors across interconnected workflows. As one analyst aptly observed, the paradigm has shifted from "garbage in, garbage out" to "garbage in, disaster out".[1]

1.2. The Research Problem

The intersection of data quality management and agentic AI presents a multi-faceted research problem. First, traditional data quality approaches were designed for human-in-the-loop business intelligence processes, where analysts could identify and correct anomalies during report preparation. Agentic systems operate autonomously, removing this human validation layer. Second, the velocity and volume of



data processed by AI agents exceed the capacity of manual governance approaches. Third, the semantic complexity of enterprise decision-making requires not merely accurate data, but contextually appropriate interpretation a challenge that pure syntactic data quality measures cannot address. Contemporary research has begun exploring these dimensions. A recent doctoral thesis from Carnegie Mellon University introduced the concept of "interactive data profiling" to accelerate data exploration and quality assessment. Similarly, systematic literature reviews have examined how big data analytics enables quality decision-making in IT sectors, identifying both technological enablers and organizational barriers. Yet a gap remains in understanding how ontology-based knowledge representation can bridge the divide between raw data quality and intelligent decision-making in agentic systems.[2]

1.3. Research Questions and Objectives

This study addresses the following research questions:

RQ1: How does the transition from business intelligence to agentic AI transform data quality requirements and management approaches?

RQ2: What role can ontology-based semantic layers play in enabling AI agents to interpret and validate data within decision-making contexts?

RQ3: What architectural framework can integrate data quality management, semantic knowledge representation, and agentic AI capabilities for enterprise deployment?

The primary objective is to develop a conceptual framework for ontology-driven data quality management that supports intelligent decision-making in agentic AI environments. Secondary objectives include synthesizing contemporary industry and academic perspectives on this convergence and providing implementation guidance for practitioners.[3-4]

1.4. Significance and Scope

This research holds significance for multiple stakeholders. For academic researchers, it contributes to the emerging literature on agentic AI governance and semantic integration. For industry practitioners particularly those with software engineering

backgrounds transitioning into data analytics leadership it offers a structured approach to addressing data quality challenges in AI deployments. For organizational decision-makers, it provides a framework for evaluating investments in data modernization relative to AI strategic priorities. The scope encompasses enterprise-scale agentic AI deployments, with particular attention to the integration of data quality management, knowledge graphs, and semantic ontologies. While examples draw from multiple industry sectors, the framework is designed to be domain-agnostic.

1.5. Paper Structure

This paper proceeds as follows: Section 2 presents a systematic literature review examining agentic AI, data quality management, ontology, and knowledge representation, and their intersections. Section 3 describes the research methodology employed in developing the proposed framework. Section 4 presents the conceptual framework for ontology-driven data quality management. Section 5 discusses implications for research and practice. Section 6 concludes with contributions and future research directions.

2. Literature Review

2.1. The Emergence of Agentic AI

The concept of agentic artificial intelligence represents a maturation of AI capabilities beyond simple pattern recognition and content generation. Unlike earlier AI systems that required explicit human prompting for each task, agentic AI encompasses autonomous systems capable of understanding objectives, planning sequences of actions, executing tasks across multiple tools and data sources, and adapting to changing circumstances.

Industry research indicates that agentic AI is transitioning from experimentation to practical deployment. A 2025 survey revealed that 36% of organizations were actively experimenting with agent frameworks, while 23% had implemented at least single-agent systems in production environments. Furthermore, 67% of surveyed organizations reported exploring agentic AI as a vehicle for innovation a clear indicator of momentum toward widespread adoption. The implications for data management are profound. Agentic systems require



continuous access to multimodal data—structured databases, unstructured text, images, video, and real-time streaming sources. They must interpret this data within specific business contexts, reason about relationships and dependencies, and make decisions that align with organizational objectives. This complexity demands data management approaches far exceeding traditional data warehousing paradigms.[4]

2.2.Data Quality in the Age of AI

Data quality has long been recognized as foundational to analytics and decision-making. However, the emergence of AI-driven processes has fundamentally altered both the stakes and the requirements of data quality management. Recent research highlights a concerning pattern: many organizations, rushing to adopt AI capabilities, have deprioritized data quality investments. A 2025 study found that while 77% of large enterprises expect AI data quality failures to cause major crises, data quality management is not being prioritized commensurate with this risk. This disconnect between recognized importance and actual investment creates significant vulnerability. The concept of "data readiness" has emerged as a critical framework for understanding AI data requirements. Data readiness encompasses not merely the accuracy and completeness of data, but its fitness for specific AI use cases—including appropriate documentation, consistent semantics, clear lineage, and governance frameworks that extend throughout the data lifecycle. Organizations that achieve high data readiness are positioned to realize greater value from AI investments, while those neglecting this foundation face mounting technical debt and operational risk. Contemporary approaches to data quality are evolving to address these challenges. Data observability the ability to understand data health, quality, and behavior through continuous pipeline monitoring has emerged as a critical capability. Leading platforms now incorporate automated anomaly detection, lineage tracking, and alerting mechanisms that enable proactive quality management. However, observability alone cannot resolve semantic ambiguities or provide the contextual understanding that AI agents require.[5]

2.3.Ontology and Knowledge Representation

Ontology the formal representation of concepts within a domain and the relationships between them has reemerged as a central concern in AI and data management. While ontology enjoyed attention during the semantic web era of the early 2010s, it was largely overshadowed by simpler graph representations and relational approaches. The rise of agentic AI has catalyzed its return to prominence. The distinction between property graphs and semantic (ontology-based) graphs is critical. Property graphs represent entities and relationships but lack the formal logical foundations that enable inference, constraint validation, and complex reasoning. Semantic graphs, built on standards such as RDF (Resource Description Framework) and OWL (Web Ontology Language), provide these capabilities. For agentic AI, this distinction matters profoundly. Agents must understand not just that two entities are related, but the nature, temporal boundaries, and logical constraints governing those relationships. Microsoft's recent introduction of an ontology-based "IQ layer" across its data platforms Foundry IQ, Fabric IQ, and Work IQ exemplifies this trend. By placing ontology at the core of enterprise data architecture, Microsoft signals that knowledge representation has moved from theoretical interest to practical necessity. Organizations are increasingly recognizing that competitive advantage in analytics will accrue not to those with the largest data volumes, but to those with the richest ontological models enabling superior interpretation.[6-7]

2.4.The Intersection: Intelligent Decision-Making Systems

The convergence of big data analytics, software engineering innovations, and AI capabilities has given rise to a new class of intelligent decision-making systems. Recent scholarship has examined how these systems transform organizational intelligence across sectors, including healthcare, finance, and governance. A systematic literature review published in late 2025 synthesized findings from multiple studies on big data analytics for decision-making, identifying both technological enablers and organizational success factors. Key technological enablers include distributed computing



frameworks (Apache Hadoop, Spark), AI-driven analytics instruments (TensorFlow, Scikit-learn), and business intelligence platforms (Tableau, Power BI). Organizational factors including leadership commitment, workforce competencies, and data governance maturity were found to be equally determinative of success. Emerging research also addresses the software engineering dimensions of intelligent decision systems. Studies on AI-driven innovations in software engineering examine how development practices must evolve to accommodate the unique requirements of AI-enabled applications. Topics include pipeline architecture, model deployment and monitoring, and the integration of analytics capabilities into broader software lifecycles.

2.5. Research Gap

Despite growing attention to each of these domains individually, limited research addresses their intersection. Specifically, the question of how ontology-based knowledge representation can systematically address data quality challenges in agentic AI contexts remains underexplored. Existing frameworks tend to treat data quality, semantic integration, and AI governance as separate concerns, yet the operational reality of agentic systems demands their integration. This paper addresses this gap by proposing a unified framework that positions ontology as the bridge between data quality management and intelligent decision-making in agentic AI environments.[8]

3. Research Methodology

3.1. Research Approach

This study employs a conceptual framework development methodology grounded in systematic literature review and synthesis of contemporary industry and academic sources. Given the rapid evolution of agentic AI technologies and the emergence of ontology-based approaches in 2025-2026, this approach enables the integration of cutting-edge developments into a coherent analytical framework. The methodology follows established guidelines for conceptual research in information systems, emphasizing:

- Systematic identification and review of relevant literature
- Synthesis of findings across multiple sources

and perspectives

- Development of novel conceptual integrations addressing identified gaps
- Validation of framework coherence through logical consistency and practical applicability assessment

3.2. Literature Search Strategy

Literature identification employed a multi-stage search strategy. Initial keyword searches were conducted across academic databases, including ACM Digital Library, IEEE Xplore, and Google Scholar. Search terms included combinations of: "agentic AI," "data quality management," "ontology," "knowledge graph," "intelligent decision-making," and "semantic integration." Given the rapid evolution of this domain, industry sources were also systematically reviewed, including research publications from TDWI, Info-Tech Research Group, and Monte Carlo Data. This dual approach ensured capture of both peer-reviewed academic contributions and timely industry insights relevant to 2025-2026 developments.[9-10]

Inclusion criteria encompassed:

- Publication dates primarily 2024-2026 (to capture emerging developments)
- Direct relevance to at least one of the core domains
- Sufficient methodological rigor or evidentiary basis
- English language publications

3.3. Framework Development Process

The conceptual framework was developed through an iterative process:

- **Phase 1:** Synthesis of findings from literature review to identify key constructs, relationships, and requirements.
- **Phase 2:** Identification of integration points between ontology, data quality, and agentic AI capabilities.
- **Phase 3:** Development of framework architecture specifying components, relationships, and information flows.
- **Phase 4:** Validation of framework completeness against requirements derived from literature.
- **Phase 5:** Refinement based on logical



consistency assessment and practical applicability considerations.[10-11]

3.4.Limitations

This research has several limitations. First, the rapid evolution of agentic AI technologies means that some framework components may require adaptation as the technology landscape matures. Second, the conceptual nature of the framework precludes empirical validation within the scope of this paper; future research should address this through case studies or experimental implementations. Third, the literature review, while systematic, cannot claim complete coverage of all relevant sources given the breadth of intersecting domains.[12-13]

4. A Framework for Ontology-Driven Data Quality Management in Agentic Ai Systems

4.1.Framework Overview

Based on the synthesis of literature and industry developments, this paper proposes the Ontology-Driven Data Quality Management (ODDQM) framework for agentic AI environments. The framework addresses three fundamental requirements identified in the literature:

- Semantic grounding: AI agents require not merely data, but a contextualized understanding of what data represents and how it should be interpreted.
- Automated quality management: The scale and velocity of agentic AI operations necessitate continuous, automated quality monitoring rather than periodic human review.
- Governance integration: Decision integrity requires auditable pathways connecting data sources to agent actions, enabling both compliance and continuous improvement.

The framework comprises four interconnected layers: the Data Foundation Layer, the Ontology Layer, the Agentic AI Layer, and the Governance Layer.

4.2.Layer 1: Data Foundation Layer

The Data Foundation Layer encompasses the organization's data assets, infrastructure, and quality management capabilities. This layer addresses the "data readiness" imperative identified in the literature and incorporates several key components:

Multimodal Data Integration: Agentic systems

require access to diverse data types—structured transactional data, unstructured documents, real-time streams, and increasingly, image and video content. The foundation layer must provide unified access to these disparate sources while maintaining appropriate security and access controls.

Data Observability: Continuous monitoring of data pipelines enables early detection of anomalies in freshness, volume, distribution, and schema. Observability capabilities generate alerts when data quality degrades, triggering either automated remediation or human intervention depending on severity.

Data Contracts: Formal agreements between data producers and consumers establish expectations for quality, timeliness, and structure. In agentic contexts, these contracts become machine-readable specifications that agents can consult when assessing data fitness for specific use cases.

Quality Metrics Framework: Standardized metrics for accuracy, completeness, consistency, timeliness, and validity provide a common vocabulary for quality assessment across the organization.

4.3.Layer 2: Ontology Layer

The Ontology Layer represents the framework's semantic core, addressing the knowledge representation requirements identified in the literature. This layer transforms raw data into interpretable knowledge through several mechanisms:

Domain Ontology: A formal representation of concepts, relationships, and constraints within the organization's operational domain. Following semantic web standards (RDF/OWL), the ontology defines:

- Conceptual classes and their hierarchies (e.g., Customer > Enterprise Customer > Individual Customer)
- Properties and relationships with temporal and logical constraints
- Business rules and validation criteria (e.g., "Credit limit cannot exceed 10x average monthly transaction volume")
- Disjoint sets and classification boundaries

Semantic Integration: The ontology provides a unified vocabulary that harmonizes data across



disparate sources. When the same concept appears with different labels or structures across systems, the ontology establishes mappings that enable consistent interpretation.

Context Window Management: For agentic AI applications, the ontology populates context windows with relevant semantic information, enabling agents to interpret data within appropriate business contexts. This addresses the risks of context poisoning (incorporating hallucinations), context distraction (over-learning older information), and context clash (conflicts between information sources).

Inference Capabilities: The formal logical foundations of the ontology enable inference—deriving new knowledge from explicitly stated facts. For example, if the ontology defines that "Subsidiary companies are owned by Parent companies" and "Parent companies are subject to consolidated reporting requirements," an agent can infer that a subsidiary's data should be included in consolidated reports even if not explicitly flagged.[14]

4.4.Layer 3: Agentic AI Layer

The Agentic AI Layer encompasses the autonomous systems that execute workflows and make decisions based on data and knowledge. This layer operationalizes the semantic understanding provided by the ontology layer:

Agent Orchestration: Multi-agent systems coordinate to execute complex workflows. The framework supports both single-agent deployments and collaborative multi-agent architectures.

Semantic Query Processing: Rather than issuing SQL queries against raw data structures, agents interact with the ontology layer through semantic queries that leverage the knowledge representation. This insulates agents from underlying data complexity while ensuring consistent interpretation.

Decision Execution: Agents execute decisions within defined boundaries, with escalation pathways for situations exceeding their autonomy thresholds. Decision logs capture both inputs and reasoning for audit and improvement purposes.

Feedback Integration: Agent outcomes feed back into both the data foundation (identifying quality issues requiring remediation) and the ontology

(identifying semantic gaps or inconsistencies).

4.5.Layer 4: Governance Layer

The Governance Layer provides oversight, compliance, and continuous improvement capabilities essential for trustworthy agentic AI:

AI Governance Framework: Policies, standards, and processes governing AI development and deployment, including model validation, performance monitoring, and risk management.

Lineage and Auditability: Complete tracing of decision pathways from data sources through ontological interpretation to agent actions. This capability addresses both compliance requirements and post-incident analysis.

Quality Threshold Management: Definition of acceptable quality levels for different decision contexts, with automated enforcement mechanisms. High-stakes decisions may require higher quality thresholds than routine operations.

Human Oversight Integration: Workflows for human review of agent decisions that exceed confidence thresholds or encounter ambiguous situations, with feedback loops to improve agent performance over time.

4.6.Framework Integration and Information Flows

The power of the ODDQM framework lies in the integration and information flows between layers:

Data-to-Ontology: Raw data from the foundation layer is semantically enriched through the ontology, acquiring context and meaning that enable intelligent interpretation.

Ontology-to-Agent: Agents receive not raw data but semantically interpreted information, with associated confidence scores, lineage, and contextual metadata that inform decision-making.

Agent-to-Governance: All agent actions, including inputs, reasoning pathways, and outputs, flow to the governance layer for monitoring, audit, and continuous improvement.

Governance-to-Foundation: Governance insights inform data quality investments, identifying priority areas for remediation based on agent requirements and failure patterns.

Feedback Loops: Multiple feedback loops enable continuous improvement—agents flag potential data



quality issues, governance identifies semantic gaps requiring ontology updates, and foundation layer enhancements expand the data available for agentic processing.

4.7. Implementation Considerations

Organizations seeking to implement the ODDQM framework should consider several practical factors: **Maturity Assessment:** Begin by assessing current capabilities across all four layers, identifying priority gaps relative to agentic AI ambitions.

Incremental Deployment: Rather than attempting comprehensive implementation, adopt an incremental approach focused on high-value use cases. Deploy minimal viable ontologies that can be expanded over time.

Tool Selection: Evaluate platforms and tools against framework requirements. Key considerations include semantic web standards support, data observability capabilities, agent orchestration frameworks, and governance integration.

Skill Development: Ontology development requires specialized expertise that may not exist within current teams. Organizations should consider either developing internal capabilities through training and hiring or engaging specialized consultants.

Governance Evolution: Recognize that governance practices must evolve alongside technical capabilities. Static governance frameworks will quickly become obsolete in dynamic agentic environments.[15-16]

5. Discussion

5.1. Implications for Research

The ODDQM framework contributes to multiple streams of academic inquiry. For information systems researchers, it provides a conceptual integration of previously disparate domains data quality management, knowledge representation, and AI governance. This integration suggests new research questions about the interactions between these domains and the conditions under which ontology-driven approaches yield superior outcomes. For software engineering researchers, the framework highlights the need for development methodologies suited to ontology-enabled agentic systems. Traditional software engineering practices assume human-in-the-loop decision processes and

deterministic outcomes; agentic systems challenge these assumptions and require new approaches to validation, testing, and evolution. For data management researchers, the framework reinforces the importance of semantic approaches that had perhaps been undervalued in the era of big data hype. As organizations accumulate ever-larger data volumes, the challenge shifts from storage and processing to interpretation—precisely the domain where ontology adds value.

5.2. Implications for Practice

For industry practitioners, the framework offers practical guidance for navigating the transition to agentic AI. Several implications stand out:

Data Quality as Strategic Priority: Organizations cannot afford to treat data quality as an afterthought in AI initiatives. The ODDQM framework positions quality management as foundational rather than ancillary, with direct implications for investment priorities.

Ontology as Competitive Differentiator: As AI capabilities commoditize, competitive advantage will increasingly derive from the quality of organizational knowledge represented in ontologies. Early investment in ontology development may yield sustained advantage.

Governance as Enabler: Rather than viewing governance as a constraint, practitioners should recognize its enabling role in building trustworthy agentic systems that can scale with confidence.

Skills Evolution: The convergence of data management, AI, and domain knowledge required for ontology development suggests new skill profiles and career pathways for data professionals.

5.3. Limitations and Future Research

The ODDQM framework, while grounded in a systematic literature review, remains conceptual and requires empirical validation. Future research should address several questions:

Empirical Validation: Case studies of organizations implementing ontology-driven approaches to data quality in agentic contexts would provide valuable evidence of framework effectiveness and refinement opportunities.

Quantitative Assessment: Development and validation of metrics for assessing ontology quality,



semantic integration effectiveness, and impact on agent decision accuracy would enable more rigorous evaluation.

Domain Adaptation: Research examining how the framework should be adapted for specific industry domains healthcare, finance, manufacturing would enhance practical applicability.

Technology Evolution: As agentic AI technologies continue to evolve rapidly, a longitudinal research tracking framework relevance and adaptation requirements would be valuable.

Organizational Factors: Studies examining the organizational enablers and barriers to framework adoption would complement the technical focus of this research.[16-17]

Conclusion

Summary of Contributions

This paper has addressed the intersection of data quality management, ontology-based knowledge representation, and agentic artificial intelligence—three domains whose convergence is shaping the future of enterprise decision-making. The primary contribution is the Ontology-Driven Data Quality Management (ODDQM) framework, which provides a conceptual architecture for integrating semantic understanding into automated quality management for agentic AI systems. The framework addresses the fundamental challenge identified throughout the literature: agentic systems require not merely accurate data, but contextually appropriate interpretation that traditional quality approaches cannot provide. By positioning ontology as the bridge between raw data and intelligent decision-making, the framework offers a pathway from "garbage in, disaster out" to reliable, scalable AI operations.

Answering the Research Questions

Returning to the research questions that motivated this study:

RQ1: The transition from business intelligence to agentic AI transforms data quality requirements by removing human validation from the loop, demanding automated quality management at scale, and requiring semantic as well as syntactic accuracy. Traditional BI-focused quality approaches prove inadequate for these demands.

RQ2: Ontology-based semantic layers enable AI

agents to interpret data within appropriate business contexts, resolve ambiguities, apply logical constraints, and maintain decision integrity across complex workflows. Ontology provides the knowledge foundation that raw data alone cannot supply.

RQ3: The ODDQM framework integrates data quality management, semantic knowledge representation, and agentic AI capabilities through four interconnected layers: Data Foundation, Ontology, Agentic AI, and Governance. The framework's information flows and feedback loops enable continuous improvement and governance integration.

Concluding Remarks

As organizations navigate the transition from experimental AI to production-grade agentic systems, the quality and interpretability of underlying data will increasingly determine outcomes. Those that invest in semantic foundations—ontologies that capture domain knowledge, relationships, and constraints—will be positioned to realize the full potential of autonomous AI. Those that neglect these foundations risk compounding data quality failures into systemic decision disasters. For professionals at the intersection of software engineering, data analytics, and information technology—the author's own background—this convergence presents both challenge and opportunity. The technical skills to build and integrate these systems must be complemented by conceptual frameworks for understanding how they create value. The ODDQM framework aims to contribute to that understanding, providing a bridge from technical capability to strategic impact. The future of enterprise decision-making lies not in more data, but in better understanding. Ontology-driven data quality management offers a path to that understanding—and to AI systems that truly serve organizational intelligence.

References

- [1]. Ikonnikov, I. (2026). Top Three 2026 Tech Predictions for Data and Analytics. Info-Tech Research Group.
- [2]. Epperson, W. (2025). Interactive Data Profiling. Ph.D. Thesis, Carnegie Mellon



- University, Human-Computer Interaction Institute. CMU-HCII-25-102.
- [3]. Bhatia, K. (2025). Understanding big data analytics (BDA) for quality decision making in the IT sector – a systematic literature review. Master's Dissertation, Auckland University of Technology.
- [4]. University of South Africa. (2025). AI (Artificial Intelligence): Responsible and Ethical Use of AI. LibGuides Unisa.
- [5]. Segner, M. (2025). Future Of Data Analytics: 10 Trends To Watch Out For In 2026. Monte Carlo Data.
- [6]. International Journal of Engineering and Computer Science. (2025). The Crunchy Part Entails the Innovations Offered In Big Data Analytics and Software Engineering of Smarsh Decision-Making Systems. IJECS, 14(11), 27971-27982.
- [7]. Das, N. (2025). The Role of Business Intelligence Tools in Strategic, Tactical, and Operational Decision-Making.
- [8]. Frost-Camilleri, L. (2025). Chapter 7: Researching, plagiarism and academic integrity. In *The Art of Re-Learning: A Guide to Academic Readiness*. Open Educational Resources Collective.
- [9]. TDWI. (2026). 2026 Trends: TDWI's Top 12 AI, Analytics & Data Predictions. TDWI Research.
- [10]. ACM Transactions on Software Engineering and Methodology. (2024). Learning from Very Little Data: On the Value of Landscape Analysis for Predicting Software Project Health. ACM TOSEM.
- [11]. Qlik. (2025). Data Quality Is Not Being Prioritized on AI Projects. Qlik Research Report.
- [12]. Zive. (2025). How AI Killed The Dashboard. Zive Research.
- [13]. Alenezi, M., & Mahmood, A. (2025). AI-driven innovations in software engineering: A review of current practices and future directions. *Applied Sciences*, 15(3), 1344.
- [14]. Banh, L., et al. (2025). Copiloting the future: How generative AI transforms software engineering. *Journal of Systems and Software*.
- [15]. Chandra, R., Agarwal, S., Kumar, S. S., & Singh, N. (2025). OCEP: An ontology-based complex event processing framework for healthcare decision support in big data analytics. arXiv preprint arXiv:2503.21453.
- [16]. Panyam, S. (2025). How AI agents are transforming software engineering. *IEEE Computer Magazine*, 58(5), 19-28.
- [17]. Terragni, V., Vella, A., Roop, P., & Blincoe, K. (2025). The future of AI-driven software engineering. *ACM Transactions on Software Engineering and Methodology*, 1(1), 1-32.