International Journal of Emerging Trends in Computer Science and Information Technology
ISSN: 3050-9246 | https://doi.org/10.63282/3050-9246.IJETCSIT-V5I4P117
Eureka Vision Publication | Volume 5, Issue 4, 151-155, 2024

Original Article

Data Stewardship: How AI Agents Form the Pillars for Effective Data and AI Governance

Susmit Sen
Independent Researcher and Industry Expert.

Abstract - The exponential growth of enterprise data and the rapid adoption of artificial intelligence (AI) have fundamentally altered the landscape of data governance. Traditional, manual approaches to data stewardship are increasingly inadequate for managing the complexity, velocity, and scale of modern data ecosystems. In response, agentic AIautonomous systems capable of interpreting intent, executing workflows, and adapting to dynamic conditions has emerged as a transformative solution. This manuscript provides a comprehensive analysis of how AI agents are redefining data stewardship, establishing the foundational pillars for effective data and AI governance. It examines the deployment of specialized AI agents across critical governance domains, including data quality management, metadata curation, master data management, and data retention. Furthermore, this paper addresses the imperative of regulatory compliance and responsible AI frameworks, demonstrating how multi-agent architectures enable real-time policy enforcement and risk mitigation. Through a synthesis of recent technological advancements and enterprise implementation models, this study underscores that agentic data governance is not merely an operational enhancement, but a critical evolution necessary for organizations to maintain data integrity, security, and trust in the AI era.

Keywords - Agentic AI, Data Governance, Data Quality, Data Stewardship, Large Language Models, Master Data Management, Metadata Management, Responsible AI.

1. Introduction

The modern enterprise operates within a highly complex, decentralized data ecosystem. As organizations accelerate their digital transformation initiatives, the volume and variety of data they generate, consume, and store have reached unprecedented levels. Concurrently, the integration of generative AI and machine learning models into core business processes has elevated the strategic importance of data, making it the foundational asset for competitive advantage and operational efficiency [1]. However, this data-driven paradigm is accompanied by significant risks. Poor data quality, fragmented metadata, and inconsistent governance policies frequently undermine analytics initiatives and compromise the reliability of AI models. According to industry research, poor data quality costs organizations an average of $12.9 million annually and contributes to the failure of a substantial percentage of AI projects [2].

Historically, data governance relied heavily on manual processes executed by human data stewards. These professionals were tasked with defining data standards, monitoring compliance, resolving quality issues, and maintaining business glossaries. While effective in smaller, structured environments, this manual approach cannot scale to meet the demands of modern, real-time data architectures [3]. The static rules and periodic audits characteristic of traditional governance frameworks are ill-equipped to handle continuous data flows, dynamic schema changes, and evolving regulatory requirements, such as those mandated by the European Union Artificial Intelligence Act (EU AI Act) and the National Institute of Standards and Technology (NIST) AI Risk Management Framework [4].

To overcome these structural limitations, the discipline of data stewardship is undergoing a profound evolution, driven by the advent of agentic AI. Agentic data governance represents a shift from reactive oversight to intelligent, autonomous control. By deploying specialized AI agents directly into the data lifecycle, organizations can automate complex stewardship tasks, enforce policies in real time, and dynamically adapt to changing conditions [5].

The objective of this manuscript is to systematically evaluate the role of AI agents in establishing the pillars of effective data and AI governance. Section II examines the conceptual framework of agentic data governance. Section III explores the application of AI agents across core data stewardship domains. Section IV addresses the intersection of data governance and responsible AI. Section V presents multi-agent architectural models, and Section VI concludes the manuscript.

2. The Framework of Agentic Data Governance

Agentic data governance leverages autonomous AI agents to enforce, monitor, and adapt data policies across complex enterprise environments. Unlike traditional automation, which relies on rigid, linear scripts, agentic systems possess reasoning capabilities, situational awareness, and the capacity to make independent decisions based on predefined policies and real-time context [6].

2.1. Core Principles of Agentic Systems

The efficacy of agentic data governance is rooted in several core principles that distinguish it from legacy approaches. First, AI agents operate as autonomous governance actors. Powered by Large Language Models (LLMs) and vector databases, these agents can interpret natural language intent, analyze metadata, and evaluate risk signals to determine appropriate actions, such as granting access, masking sensitive fields, or escalating anomalies to human overseers [7].

Second, agentic governance is fundamentally event-driven. Agents continuously monitor the data lifecycle and respond to specific triggers, such as schema modifications, pipeline updates, or model deployments. This continuous observability ensures that governance policies are applied at the point of data creation or transformation, rather than retroactively [8].

Third, these systems exhibit continuous learning and policy adaptation. By analyzing feedback loops, user overrides, and historical incident data, AI agents can refine risk scoring models and detect policy drift, proactively suggesting updates to governance frameworks that no longer align with operational realities [9].

2.2. The Human-in-the-Loop Imperative

While AI agents significantly augment the capacity of data stewards, they do not eliminate the need for human oversight. Effective agentic governance architectures incorporate robust human-in-the-loop (HITL) mechanisms. Complex, ambiguous, or high-impact decisions particularly those involving regulatory compliance or sensitive master dataare automatically routed to human domain experts for review [10]. This collaborative model preserves critical domain knowledge while leveraging AI to eliminate routine, labor-intensive tasks, thereby democratizing data access without compromising security or integrity.

3. AI Agents in Core Stewardship Domains

The practical implementation of agentic data governance involves deploying specialized AI agents to manage distinct pillars of data stewardship. Table I summarizes the primary agent roles and their respective functions within the enterprise data ecosystem.

Table 1. AI Agent Roles in Data Stewardship Domains

Agent Type	Primary Function	Key Capabilities
Data Quality Agent	Ensures accuracy, consistency, and reliability of enterprise data.	Anomaly detection, automated remediation, complex rule generation, root cause analysis via lineage tracking.
Metadata Agent	Curates technical and business metadata to provide context.	Automated schema extraction, continuous catalog updates, semantic linking, self-healing lineage.
Master Data Agent	Manages the lifecycle of critical data elements (CDEs).	Smart matching, deduplication, automated enrichment, cross-system data standardization.
Retention Agent	Enforces compliance with data lifecycle and privacy policies.	Automated PII classification, policy-driven archiving, secure deletion, usage pattern analysis.

3.1. Data Quality Management

Data quality is the bedrock of trustworthy analytics and AI. The Data Quality Agent transforms traditional profiling and remediation by utilizing machine learning algorithms to continuously scan structured and unstructured datasets for anomalies, missing values, and inconsistencies [11]. Beyond basic validation, advanced agents can cluster similar data exceptions, autonomously execute complex remediations, and utilize process mining to identify the root causes of data degradation.

Furthermore, these agents can translate natural language business requirements into executable data quality rules, significantly accelerating the deployment of governance controls. In a financial institution context, for example, a deployed Data Quality Agent can scan customer records from multiple source systems, identify address format anomalies, cluster similar exceptions, and automatically apply corrections while flagging edge cases for human review [12].

3.2. Metadata Curation and Lineage

Metadata provides the essential context required for data discovery and understanding. The Metadata Management Agent automates the extraction of technical metadata from new sources and employs natural language processing (NLP) to enrich data catalogs with business descriptions [13]. Crucially, these agents enable self-healing data lineage by continuously monitoring data pipelines and automatically repairing broken links or missing tags caused by schema drift.

By linking technical assets to business glossaries through semantic embeddings, the Metadata Agent ensures that data remains accessible and comprehensible to non-technical users. This capability is particularly significant as organizations increasingly rely on large language models trained on internal data, where provenance and lineage documentation are regulatory requirements [14].

3.3. Master Data Management

Master Data Management (MDM) involves governing the critical data elements that define an organization's core entities, such as customers, products, and suppliers. The Master Data Agent leverages AI to automate the creation, enrichment, and deduplication of these records [15]. By recognizing patterns across disparate source systems, the agent can intelligently merge duplicate entries and standardize schemas, maintaining a single, authoritative source of truth.

This automated reconciliation reduces the manual burden on data stewards and accelerates the integration of new data domains. Compliance checks remain essential for sensitive and business-critical master data; accordingly, the most complex decisions retain a human oversight layer, guided by the domain data strategy [16].

3.4. Data Retention and Privacy Compliance

As global data privacy regulations become increasingly stringent, managing data retention and secure disposal is a critical compliance imperative. The Data Retention Agent autonomously analyzes metadata and content to classify sensitive information, such as Personally Identifiable Information (PII) or Protected Health Information (PHI) [17]. It then enforces retention schedules by automatically triggering archiving, anonymization, or deletion protocols when data reaches the end of its legal lifecycle.

In the healthcare sector, for instance, a deployed Data Retention Agent can identify patient records subject to HIPAA retention requirements, automatically archive them after the mandated period, and securely delete them upon expiry, while continuously monitoring access patterns for potential compliance violations [18]. This proactive approach minimizes the risk of regulatory penalties and optimizes storage infrastructure.

4. Aligning Data Governance with Responsible AI

The proliferation of generative AI has intrinsically linked data governance with AI governance. The quality, provenance, and security of training data directly dictate the safety, fairness, and compliance of the resulting AI models [19]. Organizations that fail to establish robust data governance foundations before deploying AI systems face compounding risks, including model bias, regulatory non-compliance, and erosion of stakeholder trust.

4.1. Regulatory Frameworks and Data Provenance

Emerging regulatory frameworks, notably the EU AI Act, mandate rigorous data governance practices for high-risk AI systems. Organizations must maintain comprehensive documentation regarding data provenance, lineage, and the methodologies used for dataset preparation [20]. AI agents facilitate compliance by automatically generating immutable audit trails that track the origin and transformation of data throughout the model development lifecycle.

This automated lineage tracking provides the transparency required by auditors and regulatory bodies. Furthermore, the NIST AI Risk Management Framework (AI RMF 1.0) emphasizes lifecycle thinking and ongoing measurement as governance principles, which align directly with the continuous monitoring and adaptation capabilities of agentic systems [21].

4.2. Mitigating Bias and Ensuring Fairness

Responsible AI principles require organizations to actively identify and mitigate biases within their datasets. Specialized Ethics and Fairness Agents can be deployed to continuously monitor training data for demographic imbalances or historical prejudices [22]. By flagging skewed distributions and enforcing diversity constraints before model training commences, these agents help ensure that AI outputs remain equitable and aligned with organizational values.

The integration of governance agents into the model development pipeline creates a closed-loop system in which data quality, lineage, and fairness checks are performed continuously, rather than as one-time pre-deployment activities. This approach is increasingly recognized as a best practice by leading AI governance frameworks and enterprise technology providers [23].

5. Multi-Agent Architectures for Enterprise Scalability

To achieve comprehensive governance at scale, organizations are adopting multi-agent orchestration architectures. In these environments, domain-specific agents operate independently but coordinate through shared signals and semantic layers [24]. An orchestration layer manages the sequencing of agent actions, ensuring that governance decisions are consistent and non-conflicting across domains.

For example, when a schema change is detected in a production data pipeline, the orchestration layer triggers a coordinated sequence of events. The Metadata Agent updates the data catalog and refreshes lineage graphs. The Data Quality

Agent validates the new schema structure against existing quality rules. The Retention Agent checks whether newly introduced fields contain sensitive data requiring classification. Finally, the Access Control Agent adjusts permissions based on the updated sensitivity profile [25]. This interconnected architecture ensures that security, quality, and compliance measures are synchronized, preventing the emergence of governance silos.

Table II presents a comparative overview of traditional, automated, and agentic governance models, illustrating the progressive enhancement of governance capabilities across each paradigm.

Table 2. Comparative Governance Model Capabilities

Capability	Traditional	Automated	Agentic
Policy Enforcement	Manual, periodic	Rule-based, scheduled	Real-time, context-aware
Data Quality	Batch profiling	Automated scanning	Continuous, self-healing
Lineage Tracking	Manual documentation	Partial automation	End-to-end, auto-updated
Compliance Audit	Periodic review	Automated reporting	Continuous, immutable trail
Scalability	Limited by headcount	Moderate	Highly scalable, adaptive

6. Conclusion

The integration of agentic AI into data stewardship represents a fundamental paradigm shift in enterprise data governance. As demonstrated in this manuscript, autonomous AI agents transcend the limitations of manual oversight, offering unprecedented capabilities in data quality remediation, metadata curation, master data reconciliation, and regulatory compliance. By embedding intelligent, context-aware decision-making directly into the data lifecycle, organizations can ensure the continuous integrity and security of their data assets.

However, the transition to agentic data governance requires careful orchestration. Enterprises must prioritize the development of robust semantic layers, establish clear human-in-the-loop escalation protocols, and align their AI agents with established responsible AI frameworks, including the NIST AI RMF and the EU AI Act. The emerging field of Explainable AI (XAI) holds particular promise in this regard, as it provides the transparency necessary for regulatory acceptance and human-machine collaboration in high-stakes governance environments.

Ultimately, the deployment of the Digital Data Steward is not merely a mechanism for operational efficiency; it is an existential imperative for organizations seeking to build trustworthy, compliant, and scalable AI systems in an increasingly complex digital economy. Future research must focus on federated governance architectures that enable cross-organizational policy alignment without compromising proprietary data, as well as the development of standardized benchmarks for evaluating AI agent performance in data governance tasks.

Acknowledgment

The author acknowledges the contributions of the data governance and artificial intelligence research communities whose published findings form the foundation of this review. Special recognition is extended to the researchers at the World Economic Forum, the National Institute of Standards and Technology (NIST), and the European Union AI Office for their pioneering work in advancing responsible AI and data governance standards.

References

[1] IBM, “Data Governance and AI: Building a Foundation for Enterprise Intelligence,” IBM Think Insights, 2022. [Online]. Available: https://www.ibm.com/think/topics/data-governance.
[2] Alation, “The State of Data Culture Report,” Alation Inc., 2021. [Online]. Available: https://www.alation.com/state-of-data-culture/.
[3] M. Makhlouf, “On the Applicability of Machine Learning Fairness Notions,” IEEE Access, vol. 10, pp. 99651–99671, 2022, doi: https://doi.org/10.1109/ACCESS.2022.3206014.
[4] European Parliament and Council of the European Union, “Regulation (EU) 2024/1689 Artificial Intelligence Act,” Official Journal of the European Union, Jun. 2024. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689.
[5] P. Kaewkamol, “A Framework for Agentic AI in Enterprise Data Governance,” in Proc. IEEE IEEM, 2022, pp. 1–6.
[6] World Economic Forum, “Data Governance in the Fourth Industrial Revolution,” WEF White Paper, Jan. 2021. [Online]. Available: https://www.weforum.org/reports/data-governance-in-the-fourth-industrial-revolution.
[7] S. Chakraborty, “Transforming Enterprise Data Governance with Generative AI,” European Journal of Computer Science and Information Technology, vol. 10, no. 4, pp. 1–12, 2022.
[8] Ataccama, “The State of Data Quality 2022,” Ataccama Corporation, 2022. [Online]. Available: https://www.ataccama.com/resources/state-of-data-quality-2022.

[9] Acceldata, “Data Observability and Governance: A Practitioner’s Guide,” Acceldata Inc., 2022. [Online]. Available: https://www.acceldata.io/resources.
[10] Immuta, “The State of Data Engineering 2022,” Immuta Inc., 2022. [Online]. Available: https://www.immuta.com/resources/state-of-data-engineering-2022.
[11] O. Azeroual, “Data Quality Issues in Research Information Systems and Approaches for Data Cleaning,” Data, vol. 7, no. 9, p. 131, Sep. 2022, doi: https://doi.org/10.3390/data7090131.
[12] T. C. Redman, “If Your Data Is Bad, Your Machine Learning Tools Are Useless,” Harvard Business Review, Apr. 2018. [Online]. Available: https://hbr.org/2018/04/if-your-data-is-bad-your-machine-learning-tools-are-useless.
[13] Dataversity, “Data Management Trends: A Foundation for Efficiency,” Dataversity Education LLC, 2022. [Online]. Available: https://www.dataversity.net.
[14] T. Gebru et al., “Datasheets for Datasets,” Communications of the ACM, vol. 64, no. 12, pp. 86–92, Dec. 2021, doi: https://doi.org/10.1145/3458723.
[15] D. Loshin, “Master Data Management,” Morgan Kaufmann / OMG Press, 2009, ISBN: 978-0-12-374225-4.
[16] Tamr, “AI-Native Master Data Management: Empowering Responsible Data Stewardship,” Tamr Inc., 2022. [Online]. Available: https://www.tamr.com/resources.
[17] M. Janssen, H. van der Voort, and A. Wahyudi, “Factors influencing big data decision-making quality,” Government Information Quarterly, vol. 34, no. 3, pp. 418–426, 2020, doi: https://doi.org/10.1016/j.giq.2017.05.007.
[18] A. Cavoukian, “Privacy by Design: The 7 Foundational Principles,” Information and Privacy Commissioner of Ontario, 2009. [Online]. Available: https://www.ipc.on.ca/wp-content/uploads/Resources/7foundationalprinciples.pdf.
[19] V. Eubanks, “Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor,” St. Martin’s Press, New York, NY, 2018, ISBN: 978-1-250-07431-7.
[20] European Commission, “Explainability of Artificial Intelligence: Delivering on the Promise,” Publications Office of the European Union, 2021. [Online]. Available: https://ec.europa.eu/digital-single-market/en/news/explainability-artificial-intelligence.
[21] NIST, “Artificial Intelligence Risk Management Framework (AI RMF 1.0),” National Institute of Standards and Technology, Jan. 2023. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf.
[22] AI Now Institute, “AI Now Report 2023,” New York University, 2023. [Online]. Available: https://ainowinstitute.org/reports.
[23] Acceldata, “The Evolution of AI-Driven Data Governance,” Acceldata Inc., 2022. [Online]. Available: https://www.acceldata.io/blog/ai-driven-data-governance-evolution.
[24] M. Wooldridge, “An Introduction to MultiAgent Systems,” 2nd ed., John Wiley & Sons, 2009, ISBN: 978-0-470-51946-2.
[25] S. Russell and P. Norvig, “Artificial Intelligence: A Modern Approach,” 4th ed., Pearson, 2020, ISBN: 978-0-13-468599-1.