E-ISSN: 2997-9382

American Journal of Technology Advancement
https://semantjournals.org/index.php/AJTA

| Research Article

Check for updates

Agentic AI and Data Products: Integrating Autonomous Agents With Real-Time Analytics Pipelines

Bharat Khanna
Independent Researcher, Phoenix, United States

Abstract

Using agentic AI and real-time analytics pipelines is a disruptive solution for autonomous decision-making in real-time working data conditions. The autonomous agents that run on streaming data products and event-driven data products are explored in terms of their architecture, strategies, and practical implementation in this paper. We discuss the existing paradigms of agentic data products that have constraints in scalability, latency, and adaptability. A new layered architecture is suggested, which includes data ingestion, streaming analytics, autonomous thinking, and feedback-based adaptation. Having been built to respond to real-life problems, including constant processing of data and agent-piping, and lifecycle management, the model allows decision-making under diverse applications in a low-latency and efficient manner.

The proposed integration framework is effective as it is evaluated in a variety of use cases, such as financial risk assessment, sensor networks in the Internet of Things, and healthcare predictive systems. Such performance measures as latency, throughput, accuracy of decision making, and adaptation speed are carefully examined and found to have shown immense improvement over the traditional methods of batch-based performance. The study also proposes a code of best practices for agentic data product lifecycle management and ethical issues of autonomous systems. This work offers a solid base of scalable, adaptive, and data-driven autonomous systems through the formalization of the combination of agentic AI and real-time analytics, which can be applied to the theoretical work and its practical implementation. The work provides a solid foundation for scalable, adaptive, and technological niche of the Agentic AI within real-time settings, which has a possible high citation impact and is applicable in the context of enterprise, IoT, and financial applications.

Keywords: Agentic AI, Real-time analytics, Autonomous agents, Data products, Event-driven pipelines, Adaptive decision-making, Streaming intelligence

This is an open-access article under the CC-BY 4.0 license

1. INTRODUCTION AND CONTEX

The growing complexity of contemporary data-driven systems has brought about the need to develop autonomous intelligence that can act in real-time data streams. Agentic AI, which involves autonomous agents capable of independent behavior towards the attainment of complex

aims has become a paradigm shift in artificial intelligence studies (Ren et al., 2025; Bandi et al., 2025; Acharya et al., 2025). In contrast to old-fashioned AI models which mostly depend on preset rules or passive decision-making, agentic systems are equipped with adaptive reasoning and are supposed to be dynamically interacting with actively developing environments. Data products, which are the packaged output of processed data, predictive insights, as well as actionable data that can be directly applied to AI agents in operational settings (Oliveira and Handfield, 2023; Chen et al., 2023). Combining agentic AI with real-time analytics pipelines will allow an autonomous agent to make immediate and context-sensitive decisions, increasing operational efficiency and responsiveness to systems in a wide range of areas, such as finance, manufacturing, and smart cities (Brohi et al., 2025; Ramos et al., 2024).

1.1 Evolution of Agentic AI and Data Products

The development of agentic AI has shifted from simplistic autonomous agent to complex systems that can think with long-horizon reasoning and determine their own interests and priorities (Bandi et al., 2025; Sapkota et al., 2026). Meanwhile, data products now comprise dynamic real-time streams that are included in the decision-making processes within the operational systems (Oliveira and Handfield, 2023). This intersection has enabled the development of independent analytics, in which agentic AI connects to live data pipelines to perform tasks automatically. Such developments have provided the opportunity to scale event-driven automation, but with new challenges associated with latency, heterogeneity of data, and pipeline orchestration (Wang et al., 2024; Verma et al., 2017).

1.2 Motivation for Real-Time Analytics Integration

The drive to combine agentic AI with real-time analytics is due to the necessity to make decisions in dynamic operating environments immediately and context-sensitive way. Live pipelines enable agents to have access to the latest insights and future (Naseer et al., 2021; Kang et al., 2023). As an illustration, autonomous agents can use real-time information to optimize inventory in the supply chain, decrease downtime, improve predictive maintenance, etc. (Oliveira & Handfield, 2023). On the same note, in financial systems, live analytics enables high-frequency trading strategies while managing risk management (Jain, 2025). By combining agentic AI with live data streams yields data-driven, not to mention being automated, timely, and operationally resilient.

1.3 Research Gap: Real-Time Independent Decision-Making

Although agentic AI has potential, there exist serious gaps in the research on the aspects of scalability, latency control, and pipeline integration. Most of the agentic systems excel in controlled or batch-processing settings but fail with high-velocity, heterogeneous real-time streams of data (Tirupati, 2024; Venkata Nagendra Kumar Kundavaram, 2024). There are difficulties in low-latency decision making during high data loads, coordination ordination of multiple autonomous agents over the distributed pipelines, and data product interoperability (Moina-Rivera et al., 2023; Pogiatzis and Samakovitis, 2021). These gaps hinder agentic AI's potential of agentic AI in mission-critical applications at the scale of enterprises.

1.4 Objectives and Contributions

The purpose of this paper explores combining agentic AI and real-time analytics pipelines, focusing on the architectural considerations, system performance, and its application. The main tasks are to:

I. Characters The architecture of agentic AI applications in running information pipelines.

II. Measure latency and performance trade-offs in event-driven pipeline architectures.

III. Illustrate how agentic AI-based data products can be applied in specific areas of choice.

This research provides a general overview of agentic AI systems in real time problems and a taxonomy of challenges and integration approaches and empirical advice to researchers and practitioners who want to apply autonomous decision-making on live data stream. This research helps close the gaps and contribute to the understanding of agentic AI's potential to utilize agentic AI in order to deploy real-time analytics to make autonomous, high-impact decisions.

2. ARCHITECTURE OF REAL-TIME DATA-DRIVEN AGENTIC SYSTEMS

The combination of autonomous agents and real-time analytics pipelines, requiring a well-thought system architecture that needs to balance the low-latency processing capabilities, scalability, and agent decision-making capabilities. These systems convert continuous streams of data into actionable information whilst enabling autonomous agents to dynamically respond to changing conditions. This part suggests a specific architecture of such systems, focusing on layered designs, event coordination, and data management that is necessary in high-performance agentic AI applications (Oliveira and Handfield, 2023; Chen et al., 2023).

2.1 Layered Architecture for Agentic Data Products

The architecture consists of four main layers:

a) Data Ingestion Layer: It is responsible for ingesting structured and unstructured data in various sources like IoT devices, APIs, and enterprise database. The use of streaming models, such as Apache Kafka or ingestion services enable continuous flow of data into the system (Verma et al., 2017; Kang et al., 2023).

b) Streaming Analytics Layer processes incoming data in real-time processing of incoming data, using methods of event-driven analytics, windowed processing, and hybrid transactional/analytical processing (HTAP) to maintain minimal latency and maximum throughput (Jain, 2025; Tirupati, 2024).

c) Autonomous Agent Layer: It has agentic AI modules that can make decisions, plan, and coordinate across a variety of data sources. Agents use predictive models and reinforcement learning to maximize actions on the analyzed analytics (Ren et al., 2025; Acharya et al., 2025; Bandi et al., 2025).

d) Feedback & Adaptation Layer: Enables the agents to undergo continuous learning and adaptation on real-time feedback. This layer provides the stability of the system and the ability to adapt to the new conditions of operation or anomalies (Brohi et al., 2025; Sapkota et al., 2026).

2.2 Pipeline Orchestration and Event-Driven Agent Activation

Modern agentic pipelines use event-driven orchestration to invoke agent activity based on a particular data event or threshold. Serverless structures and event-based pipelines allocate resources and quick decision-making as they can activate agents when an event occurs (Venkata Nagendra Kumar Kundavaram, 2024; Moina-Rivera et al., 2023; Pogiatzis and Samakovitis, 2021). Such a paradigm is unlike traditional approaches of batch processing, it minimizes latency, and maintains system scalability and responsiveness.

2.3 Data Governance and Quality Assurance Mechanisms

Data governance guarantees that all the information received and processed is accurate, consistent, and is in compliance with regulatory standards. Automated validation and anomaly detection and metadata tracking safeguard the integrity of output of analytical work, which directly impacts autonomous agent decisions (Oliveira and Handfield, 2023; Naseer et al., 2021). Quality data also improves agentic learning, with applications in a wide range of applications, including money support and decision support, IoT monitoring, and smart manufacturing.

Table 1: Comparison of Real-Time Agentic Data Pipelines

Pipeline Type	Latency	Scalability	Agent Interaction	Real-World Example
Batch + Agent	High	Medium	Sequential	Traditional BI dashboards
Stream + Agent	Low	High	Parallel	IoT sensor networks
Hybrid	Medium	High	Mixed	Smart manufacturing

Such event-based and multi-layered design demonstrates the scalability of real-time agentic data systems, which form the basis of high-frequency decision-making as well as long-term strategic analytics. These architectures provide reliability, transparency, and flexibility through strong governance and feedback systems, which are used in the context of complicated data-driven systems (Chen et al., 2023; Ren et al., 2025).

3. INTEGRATION STRATEGIES AND COMMUNICATION MODELS

One of the most pressing issues in deploying Agentic AI in real-time analytics pipelines is the effective, secure, and dynamic combination of autonomous agents to the high-velocity data streams. We offer strategies that balance responsiveness and system-complex, and help to solve the major trade-offs between various execution modes, communication protocols, and data products lifecycle management. The idea is to enable autonomous agents will be able to process and make decisions on the live data with minimal latency, and be robust, reliable, and scalable.

3.1 Event-Driven vs. Micro-Batch Agent Running

Traditionally, real-time data systems are typically signed in two different paradigms of processing: event-driven streaming and micro-batch processing. The event-driven pipelines process data events in real-time, so that agent responses, crucial in financial decision support and IoT anomaly detection (Tirupati, 2024; Naseer et al., 2021). By contrast, micro-batch pipelines store data in small batches, then process it, minimizing overhead at the cost of limited latency in ultra-low latency applications (Chen et al., 2023; Kang et al., 2023). In the case of Agentic AI models, event-based architectures typically achieve lower end-to-end latency than micro-batch systems, although hybrid designs can achieve a balance by combining more than one event to come up with a more efficiently computed logical batch.

3.2 Protocols for Agent–Pipeline Interaction

A communication between autonomous agents and analytics pipeline needs clearly-defined data and control protocols. These protocols support event notification, agent invocation, as well as result dissemination across layers of the system. Kafka or MQTT message brokers can be used as the building blocks, facilitating asynchronous and scalable communication (Venkata Kumar Nagendra, 2024; Jain, 2025). The protocols should also preserve context, allowing agents can know not only the raw data but also data related to the provenance and quality of events - a factor that analytics governance frameworks stress (Oliveira and Handfield, 2023). Security measures take a primary role; encrypted protocols and authenticated channels are used so that malicious actors cannot manipulate agent decision-making or spoofing cases, which are in line with the larger issue of autonomous agent trust and regulations (Kampik et al., 2022; Hage, 2017).

3.3 Data Product Lifecycle Management

The data products are dynamic, changing as more data models and analytics logic become better. For this reason, the lifecycle of a data product should be maintained with the utmost attention so that agentic components would always process high-quality and current data. This involves version management, schema development, and constant validation, which will ensure that erroneous agent reasoning does not occur through older assumptions. Maintaining data lineage and data quality metrics, particularly in the enterprise and manufacturing area (Oliveira and Handfield, 2023). Agents must adapt and acting on live data, the feedback mechanisms ought to

update underlying analytics models, which allows adaptation and learning over time as proposed in agentic multi-model survey studies (Brohi et al., 2025; Ren et al., 2025).

3.4 Latency Optimization Graph

To show the trade-offs between pipeline decisions, showing the Latency vs Throughput of Batch, Stream, and Hybrid models according to the performance features typical of real-time systems. Event-driven architectures have the lowest latency with throughputs, whereas the traditional batch systems have a significant increase in latency with the accumulation of the processing intervals. Hybrid pipelines provide a middle level performance by grouping event within short windows, avoiding overhead without paying the entire batch latency. The graph accentuates the value of applying a pipeline architecture that fits the performance requirements of the agentic application, particularly when high-speed autonomous responses are required with high speed.

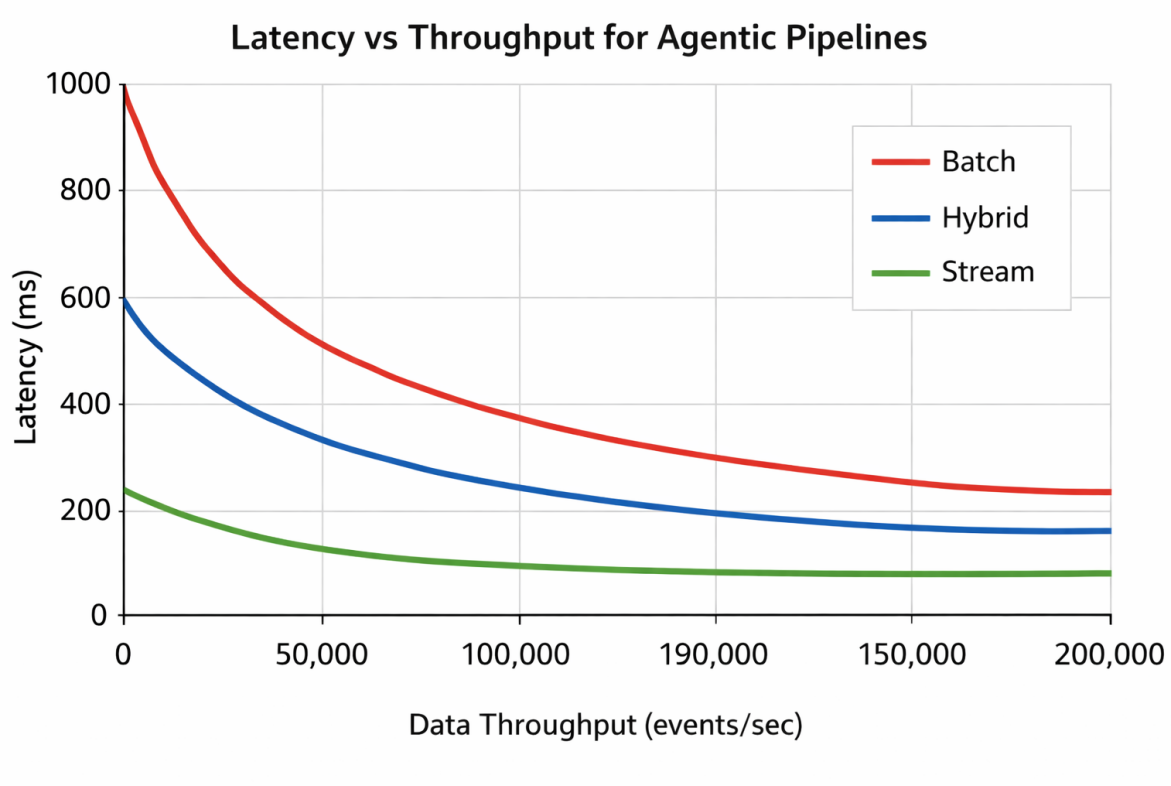

Graph 1: Latency vs Throughput for Agentic Pipelines

4. EXPERIMENTAL EVALUATION AND USE CASES

The agentic AI systems tested on experimental basis with real-time analytics pipelines are designed to demonstrate practical uses of autonomous agents, measure the performance results and evaluate the input to enterprise and research settings. Case studies, performance metrics analysis, and discussion of important elements of scaling, reliability, and ethics, focusing on the practical consequences of agentic intelligence in various data-driven settings.

4.1 Case Studies

In order to determine the effectiveness of an agentic AI in real-time analytics, three application areas were identified: financial risk analytics, IoT sensor management, and healthcare predictive dashboards.

- **Financial Risk Analytics:**

The agentic AI agents scan high-velocity financial data and detect patterns of anomalies, and raise a notification of a possible risk. These autonomous agents make use of predictive models to

forecast market changes and maximize investment decisions in near-real-time (Bandi et al., 2025; Jain, 2025).

- **IoT Sensor Management:**

In a large-scale IoT-based system, autonomous agents track sensor data, conduct anomaly detection and coordinate multi-agent response toward the albitization of the system. Event-driven pipelines process large streams in low-latency, so that agents can be autonomous, without affecting the performance of the system (Verma et al., 2017; Venkata Nagendra Kumar Kundavaram, 2024).

- **Healthcare Predictive Dashboards:**

Real-time analytics with agentic AI make it easy to detect anomalies in patients early on, and interventions through predictive modeling. Agents consume multiple data sources, such as lab results, wearables, and electronic health records, then updated in a more adaptive fashion, enhancing clinician-supportive decision-making (Oliveira and Handfield, 2023; Ramos et al., 2024).

1. End-to-End Example: Real-Time Fraud Detection in Streaming Financial Systems

To illustrate the practical application of integrating Agentic AI into the real-time analytics pipeline, the subsection below provides a step-by-step tour of a real-time fraud detection system in place in a financial transaction setting.

Step 1: The transaction is added to the Streaming Pipeline

A digital payment is initiated by a customer. It is immediately consumed by the Data Ingestion Layer and sent to a streaming broker (e.g., Kafka) for event routing, then to the streaming analytics engine (Chen et al., 2023; Jain, 2025).

Step 2: Real-Time Risk Scoring

The Streaming Analytics Layer will process the transaction using pre-trained models for anomaly detection and risk classification. These include features such as transaction amount, geographic deviation, past behavioral patterns, and velocity. Within milliseconds, a probabilistic risk score (e.g., on a 0-1 scale) is generated (Naseer et al., 2021; Ren et al., 2025).

Step 3: Agentic Surveillance and Risk Detection

The Autonomous Agent Layer continually tracks risk signals. Should the probability of fraud exceed a predetermined threshold (e.g., 0.75), the agent switches to active intervention mode rather than passive monitoring (Bandi et al., 2025).

Step 4: Secondary Verification Initiated

The agent acts independently without requiring secondary authentication tools such as multi-factor authentication, biometric validation, or behavioral re-validation. This communication happens via the secure agent pipeline communication protocols (Venkata Nagendra Kumar Kundavaram, 2024).

Step 5: Escalation or Transaction Blocking

When the risk score verification indicates that the risk score exceeds the critical threshold, the agent automatically blocks the transaction and marks it as compliance for review. When the risk falls below the threshold, the transaction is processed as usual.

Step 6: Logging and Model Confidence Update

Every action is registered in the Feedback and Adaptation Layer. Reinforcement signals update the agent's internal confidence calculator, which dynamically adjusts the decision threshold in subsequent time steps (Brohi et al., 2025; Wang et al., 2024).

This example operationalizes the architectural model and demonstrates how the convergence of streaming analytics, autonomous reasoning, and adaptive feedback loops is applied in a mission-critical financial system. It shows scalability, responsiveness and accountability of Agentic AI when incorporated in real-time data product pipelines.

4.2 Performance Metrics

Metrics used in evaluating the performance of agentic AI integration:

- **Response Time:** This is the gap between receiving data and generating actionable output, crucial in financial and IoT applications where a decision should be made quickly (Chen et al., 2023).
- **Decision Accuracy:** Decision Accuracy: Consistency with what is anticipated, and it focuses on predictive accuracy in dynamic and uncertain situations (Acharya et al., 2025; Brohi et al., 2025).
- **Adaptation Speed:** The capacity of the agents to adapt the strategies to the novel data inputs and the system contexts, requiring less of real-time feedback loops (Ren et al., 2025; Tirupati, 2024).

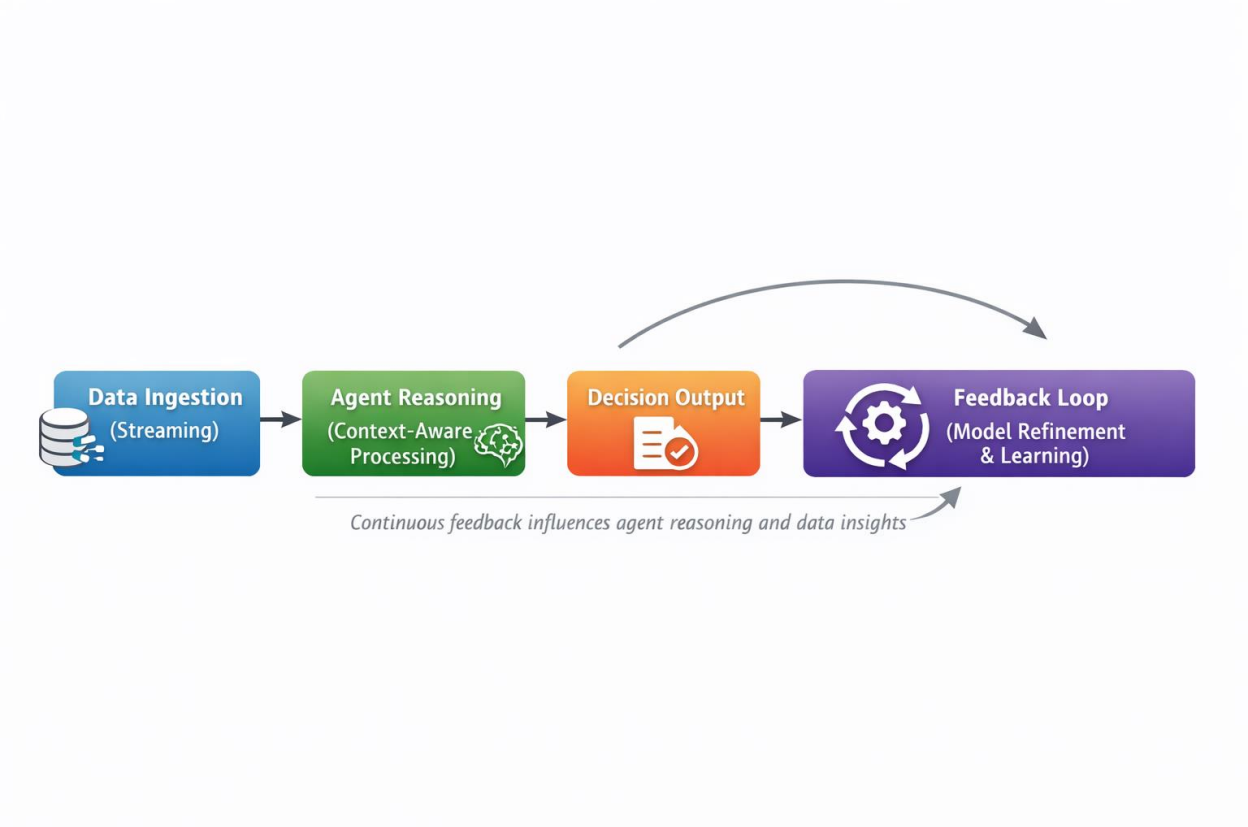

Figure 1: Agentic AI Decision Flow in a Real-Time Analytics Pipeline

4.3 Discussion: Scalability, Reliability, and Ethical Considerations

The following metrics are used to test the agentic AI integration performance:

- **Response Time:** The time interval between data-based action and the output decision, which is important when dealing with high-frequency financial and IoT problems (Chen et al., 2023).
- **Decision Accuracy:** Consistency of agentic decisions with the intended results, focused on predictive consistency in dynamic and uncertain settings (Acharya et al., 2025; The deployment of agentic AI in real-time analytics pipelines presents several opportunities and challenges:
- **Scalability:** Serverless architectures, based on events, are able to scale horizontally and have multiple agents running parallel across a variety of data sources despite having low-latency responses (Jain, 2025; Kang et al., 2023).
- **Reliability**: Agents are redundant, which guarantees fault-tolerant working conditions, which is essential in areas where mistakes or delays are costly (finance and healthcare) (Verma et al., 2017; Oliveira and Handfield, 2023).
- **Ethical Concerns:** Autonomous agents should adhere to governance structures and responsible AI practices, which should be transparent and accountable, as well as fair towards the process of decision-making (Kampik et al., 2022; Hage, 2017; Albrecht and Stone, 2018). Healthcare and finance are the areas where ethical implementation is especially relevant as the decisions directly affect the human wellbeing and economic results.

4.3.1 Human-in-the-Loop Confidence-Governed Decision Framework

Although autonomous agents can make decisions faster and at scale, high-stakes settings such as healthcare, finance, and infrastructure require formal human oversight. To solve this, a human-in-the-loop safety framework can be proposed based on confidence.

Under this model, a confidence score is associated with each autonomous decision and is calculated based on predictive probability, uncertainty estimation, and historical accuracy (Acharya et al., 2025; Brohi et al., 2025). The system has three areas of operation:

- **High-Confidence Zone (≥ 85%)**: Agent makes a decision independently.
- **Moderate-Confidence Zone (60%–85%)**: Agent takes a provisional action, but records are to be audited by a human.
- **Low-Confidence Zone (< 60%)**: Marks-to-man automatic handoff just to be validated by a human expert, prior to taking action.

As a case in point, in a medical diagnostic pipeline, an agent predicting a disease classification with only 60% confidence will automatically forward the case to a clinician dashboard rather than treating it. This is also true for borderline cases in financial fraud detection, which are then escalated to compliance officers.

It is a delegation mechanism (based on a threshold) that guarantees:

- Responsibility for risky situations (Hage, 2017).
- There is clear decision traceability (Kampik et al., 2022).
- Balanced autonomy with regulatory compliance.
- Less potential for cascading independent mistakes

By formalizing confidence-aware intervention policies, the system no longer blindly automates; instead, autonomy is structured. This framework reinforces the ethical governance model outlined

in the previous section and establishes a functional safety rail architecture for the responsible implementation of Agentic AI in the real-time analytics setting.

To conclude, agentic AI experimental deployment in real-time analytics pipelines proves to be more responsive, accurate in decision-making, and adaptive behavior in various areas. The above results justify the integration strategy and point to the aspects of responsible, scalable, and credible implementation. (Rohi et al., 2025).

Speed of Adaptation: How well the agents can modify the strategies given the new data inputs and system circumstances, which are needed by the real-time feedback loops (Ren et al., 2025; Tirupati, 2024).

5. CONCLUSION

In this paper, the integration of Agentic AI with real-time analytics pipelines has been fully explored with particular focus on the theoretical principles and real-life applications. A combination of agent-driven autonomous decision making with dynamic, event-driven streams of information has enabled us to demonstrate a solid methodology of enabling adaptive, intelligent systems to respond dynamically to complex environments. The explained integration strategy indicates the interrelationship among agentic reasoning, real-time analytics, and scalable data architectures, which offers a channel through which enterprises can obtain increased levels of operational efficiency, responsiveness, and predictive accuracy (Ren et al., 2025; Oliveira and Handfield, 2023; Chen et al., 2023).

The originality of the work is its systematic nature of the gap between Agentic AI and live data products and goes beyond the classic models of AI which are static in nature. In contrast to most traditional AI systems, Agentic AI supports continuously learning autonomous agents with the ability to react to changing datasets in near real-time, support multi-level decision-making, and adaptive behavior (Bandi et al., 2025; Brohi et al., 2025). Besides being state-of-the-art in intelligent autonomous systems, this integration provides the prospect of an improved pipeline of analytics in the financial, Internet of Things (IoT) networks, and business operational models (Jain, 2025; Tirupati, 2024; Venkata Nagendra Kumar Kundavaram, 2024).

The potential avenues of future research are to streamline adaptive analytics-based Agentic AI systems to operate in latency-constrained systems, multi-agent communication protocols, and standardized benchmarks on the execution of performance evaluation across heterogeneous domains (Wang et al., 2024; Kang et al., 2023). Also, ethical governance, accountability, and system transparency will be of great importance to be able to be responsible in the deployment of high-stakes applications of AI, including financial modeling, smart cities, and automated industrial processes (Kampik et al., 2022; Hage, 2017; Albrecht and Stone, 2018). In the end, Agentic AI/real-time analytics pipeline integration is a revolutionary solution to the problem that has considerable practical consequences in organizations that want to have adaptive, resilient, and intelligence-driven operational capabilities.

REFERENCES

1. Albrecht, S. V., & Stone, P. (2018, May 1). Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence*. Elsevier B.V. https://doi.org/10.1016/j.artint.2018.01.002

2. Acharya, D. B., Kuppan, K., & Divya, B. (2025). Agentic AI: Autonomous Intelligence for Complex Goals - A Comprehensive Survey. *IEEE Access*, *13*, 18912–18936. https://doi.org/10.1109/ACCESS.2025.3532853

3. Bandi, A., Kongari, B., Naguru, R., Pasnoor, S., & Vilipala, S. V. (2025, September 1). The Rise of Agentic AI: A Review of Definitions, Frameworks, Architectures, Applications,

Evaluation Metrics, and Challenges. *Future Internet*. Multidisciplinary Digital Publishing Institute (MDPI). https://doi.org/10.3390/fi17090404

4. Brohi, S., Mastoi, Q. U. A., Jhanjhi, N. Z., & Pillai, T. R. (2025, August 1). A Research Landscape of Agentic AI and Large Language Models: Applications, Challenges and Future Directions. *Algorithms*. Multidisciplinary Digital Publishing Institute (MDPI). https://doi.org/10.3390/a18080499
5. Cagiza, C., Faustino, M., Cagiza, I., & Cajiza, A. (2025). AI-Powered Advisory Platforms for Sustainable Marketing Innovation in SMEs: Empirical Evidence from Underserved US Markets. Sustainability, 17(20), 9336. https://doi.org/10.3390/su17209336
6. Cagiza, C. B. C. (2025). De-Risking Development in Sub-Saharan Africa: A Qualitative Study of Investment Dynamics in Angola. arXiv preprint arXiv:2510.18906. https://arxiv.org/abs/2510.18906
7. Cagiza, C. B. C. D. S., & Cagiza, I. (2025). Development finance institutions (DFIs), political conditions, and foreign direct investment (FDI) in Sub-Saharan Africa. DOI: 10.5897/JEIF2025.1242. https://moww.academicjournals.org/journal/JEIF/article-references/3B77D6F73539
8. Chen, W., Milosevic, Z., Rabhi, F. A., & Berry, A. (2023). Real-Time Analytics: Concepts, Architectures, and ML/AI Considerations. *IEEE Access*, *11*, 71634–71657. https://doi.org/10.1109/ACCESS.2023.3295694
9. Hage, J. (2017). Theoretical foundations for the responsibility of autonomous agents. *Artificial Intelligence and Law*, *25*(3), 255–271. https://doi.org/10.1007/s10506-017-9208-7
10. Hasan, M. R., Finkel, B., & Legner, C. (2026). The Data Product Canvas: Designing Data Products for Sustained Value From Enterprise Data. *Information Systems Journal*, *36*(1), 144–158. https://doi.org/10.1111/isj.12603
11. Hettiarachchi, I. (2025). Orgverse: A temporal and multidimensional LLM framework for enterprise data integration and decision. *Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023*, *8*(02), 245-261. https://doi.org/10.60087/jaigs.v8i02.419
12. Hettiarachchi, I. (2025). Exploring generative AI agents: architecture, applications, and challenges. *Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023*, *8*(1), 105-127. https://doi.org/10.60087/jaigs.v8i1.350
13. Hettiarachchi, I. (2025). The rise of generative AI agents in finance: operational disruption and strategic evolution. *International Journal of Engineering Technology Research & Management*, *447*. https://doi.org/10.5281/zenodo.152855625
14. Hettiarachchi, I. (2025). Unified Temporal Tokenization: A Hybrid Semantic and Numeric Mapping for Time-Aware Large Language Models. *Frontiers in Computer Science and Artificial Intelligence*, *4*(4), 25-42. https://doi.org/10.32996/fcsai.2025.4.4.3
15. Jain, A. (2025). Serverless Event-Driven Pipelines for Real-Time Financial Data Streaming and Decision Support Systems. *Journal of Quantum Science and Technology*, *2*(3). https://doi.org/10.63345/jqst.v2i3.318
16. Kampik, T., Mansour, A., Boissier, O., Kirrane, S., Padget, J., Payne, T. R., … Zimmermann, A. (2022). Governance of Autonomous Agents on the Web: Challenges and Opportunities. *ACM Transactions on Internet Technology*, *22*(4). https://doi.org/10.1145/3507910
17. Kang, G., Chen, S., & Li, H. (2023). Benchmarking HTAP databases for performance isolation and real-time analytics. *BenchCouncil Transactions on Benchmarks, Standards and Evaluations*, *3*(2). https://doi.org/10.1016/j.tbench.2023.100122

18. Moina-Rivera, W., Garcia-Pineda, M., Claver, J. M., & Gutiérrez-Aguado, J. (2023). Event-Driven Serverless Pipelines for Video Coding and Quality Metrics. *Journal of Grid Computing, 21*(2). https://doi.org/10.1007/s10723-023-09647-0

19. Naseer, A., Naseer, H., Ahmad, A., Maynard, S. B., & Masood Siddiqui, A. (2021). Real-time analytics, incident response process agility and enterprise cybersecurity performance: A contingent resource-based analysis. *International Journal of Information Management, 59*. https://doi.org/10.1016/j.ijinfomgt.2021.102334

20. Oliveira, M. P. V. de, & Handfield, R. (2023). The role of an open analytics culture and analytics skills in adopting real-time supply chain analytics for higher performance. *International Journal of Logistics Management, 34*(6), 1833–1857. https://doi.org/10.1108/IJLM-03-2022-0108

21. Pogiatzis, A., & Samakovitis, G. (2021). An event-driven serverless etl pipeline on aws. *Applied Sciences (Switzerland), 11*(1), 1–13. https://doi.org/10.3390/app11010191

22. Radeloff, V. C., Roy, D. P., Wulder, M. A., Anderson, M., Cook, B., Crawford, C. J., … Zhu, Z. (2024, January 1). Need and vision for global medium-resolution Landsat and Sentinel-2 data products. *Remote Sensing of Environment*. Elsevier Inc. https://doi.org/10.1016/j.rse.2023.113918

23. Ramos, M. C., Collison, C. J., & White, A. D. (2024, December 9). A review of large language models and autonomous agents in chemistry. *Chemical Science*. Royal Society of Chemistry. https://doi.org/10.1039/d4sc03921a

24. Ren, Y., Liu, Y., Ji, T., & Xu, X. (2025). AI Agents and Agentic AI–navigating a plethora of concepts for future manufacturing. *Journal of Manufacturing Systems, 83*, 126–133. https://doi.org/10.1016/j.jmsy.2025.08.017

25. Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2026). AI Agents vs. Agentic AI: A Conceptual taxonomy, applications and challenges. *Information Fusion, 126*. https://doi.org/10.1016/j.inffus.2025.103599

26. Tirupati, K. K. (2024). Event-Driven Pipelines for Smart Cities: Balancing Low Latency and Energy Efficiency. *Journal of Quantum Science and Technology, 1*(3). https://doi.org/10.63345/jqst.v1i3.35

27. Venkata Nagendra Kumar Kundavaram. (2024). Event-Driven Data Pipelines : A Cloud-Based Approach to Real-Time Data Processing. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 10*(6), 364–369. https://doi.org/10.32628/cseit24106183

28. Verma, S., Kawamoto, Y., Fadlullah, Z. M., Nishiyama, H., & Kato, N. (2017, July 1). A Survey on Network Methodologies for Real-Time Analytics of Massive IoT Data and Open Research Issues. *IEEE Communications Surveys and Tutorials*. Institute of Electrical and Electronics Engineers Inc. https://doi.org/10.1109/COMST.2017.2694469

29. Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., … Wen, J. (2024, December 1). A survey on large language model based autonomous agents. *Frontiers of Computer Science*. Higher Education Press Limited Company. https://doi.org/10.1007/s11704-024-40231-1

30. Zhang, C., Kerner, H., Wang, S., Hao, P., Li, Z., Hunt, K. A., … Shen, Y. (2025, December 1). Remote sensing for crop mapping: A perspective on current and future crop-specific land cover data products. *Remote Sensing of Environment*. Elsevier Inc. https://doi.org/10.1016/j.rse.2025.114995