Svedberg Open
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/



Research Paper Open Access

Modernizing Financial Fraud Prevention Through Event-Driven Identity Intelligence

Bhavya Reddy Maryala

Independent Researcher, USA. ORCID ID: 0009-0005-1283-9203

Abstract

Financial institutions increasingly struggle to detect fraud in the narrow window between a compromised credential and a completed transaction, largely because traditional fraud controls rely on periodic synchronization and isolated security checks rather than continuous, enterprise-wide awareness of identity state. This review examines event-driven identity intelligence as an architectural response to this gap, in which identity lifecycle changes, authentication outcomes, and behavioral signals are propagated across enterprise systems as they occur rather than on a fixed schedule. Drawing on published fraud-loss data, identity assurance standards, and the software-engineering literature on event-driven and distributed systems, the review organizes the practice into four areas: the identity events that carry fraud-relevant information, the architectural patterns used to distribute and correlate those events at scale, the fraud-detection use cases these patterns enable, and the operational tradeoffs institutions accept when adopting them. The review finds that event correlation across previously siloed systems, rather than any single detection algorithm, is what allows institutions to recognize coordinated fraud patterns earlier and to apply risk-proportional controls without adding friction for low-risk customers. It also finds that the same architectural properties that reduce fraud response time, namely asynchronous processing, distributed event routing, and eventual consistency, introduce engineering tradeoffs around latency, ordering, and partial-failure handling that must be deliberately managed. The review concludes that event-driven identity intelligence is best understood as a systems-engineering discipline layered on top of existing fraud analytics rather than a replacement for them, and it identifies predictive analytics, continuously evaluated identity confidence, and resilient event-processing pipelines as the most active areas for continued development.

Keywords: event-driven architecture, identity intelligence, fraud prevention, financial systems security, account takeover, distributed systems.

1. Introduction

Financial fraud has grown more sophisticated as banking, lending, payments, and investment services complete their shift to digital channels. Identity fraud losses in the United States reached an estimated $27.2 billion in 2024 [1], while the Federal Trade Commission recorded 6.5 million consumer fraud reports and $12.5 billion in reported losses the same reporting period [2]. The FBI's Internet Crime Complaint Center likewise recorded a multi-year rise in financial-crime complaints tied to compromised online accounts and social-engineering schemes [3]. Modern financial institutions process millions of customer interactions daily across mobile applications, online banking platforms, payment gateways, customer service systems, and third-party financial services, and every one of those interactions generates identity and behavioral information that can indicate either legitimate activity or the early stages of fraud.

Despite this volume of available signal, many fraud-prevention systems still rely on periodic synchronization, isolated security controls, and delayed risk analysis, which creates a window of opportunity between fraudulent activity and institutional response. As attackers increasingly automate credential theft, social engineering, and identity manipulation, institutions need fraud-detection strategies that respond to threats as they emerge rather than after losses occur.

Event-driven identity intelligence is an architectural response to this problem. Rather than treating identity verification, authentication, account management, and fraud detection as separate processes that periodically exchange batches of data, an event-driven architecture allows identity changes, authentication outcomes,

behavioral anomalies, and security signals to be shared across systems immediately, as discrete events, at the moment they occur [8], [9]. This review examines how event-driven identity intelligence can modernize fraud prevention, drawing on the identity-assurance standards literature [4], [5], [6], the software-engineering literature on event-driven and distributed systems [8], [9], [10], [11], [12], [13], and published empirical results on behavioral and biometric authentication [14], [15], [16]. The review is organized around four questions: what identity events carry fraud-relevant information, what architectural patterns allow those events to be distributed and correlated across an enterprise, what fraud-detection applications these patterns enable, and what operational tradeoffs institutions accept in adopting them.

2. Enterprise Identity Events

2.1 Identity Lifecycle Events

Every customer interaction with a financial institution contributes to an evolving digital identity that extends beyond the initial account-registration process. Over the customer lifecycle, identity-relevant events accumulate continuously: onboarding, profile updates, authentication attempts, credential modifications, contact-information changes, account-recovery requests, and account closures. Each of these events carries information about the current state of customer identity and is better treated as part of an ongoing identity-intelligence process than as an isolated administrative action. When lifecycle events are monitored and correlated, institutions gain a more complete picture of identity evolution and are better positioned to distinguish legitimate behavior from emerging fraud patterns.

Managing identity through lifecycle events also supports continuous identity governance across distributed financial systems. Rather than maintaining fragmented identity records in separate applications, an event-driven architecture allows identity changes to be communicated consistently across the enterprise, which improves data consistency, reduces conflicting identity information, and lets fraud-prevention systems incorporate new identity intelligence immediately into authentication and authorization decisions. This aligns with the multi-tier identity-proofing model formalized in NIST's identity-proofing guidelines, which treats identity assurance as a continuously assessed property rather than a one-time determination made at enrollment [5].

2.2 Risk Signals and Behavioral Intelligence

Modern fraud prevention depends on more than static customer information. Behavioral intelligence, meaning how customers typically interact with digital financial services, including login patterns, transaction timing, device usage, geographic consistency, and historical authentication behavior, provides context that static credentials cannot. Significant deviations from established behavioral patterns can indicate elevated fraud risk even when a login otherwise succeeds. Continuous behavioral authentication research supports this: a reduce-feature-elimination machine learning approach applied to continuous mobile authentication signals reported 98.8% accuracy with a Random Forest classifier and 95.8% with LightGBM [14], and separate keystroke-dynamics research using a large, demographically diverse dataset reported 81.16% multiclass classification accuracy alongside an 87% mean F1-score in binary classification [15]. Incorporating this kind of behavioral intelligence into an event-driven fraud-detection pipeline allows institutions to evaluate risk continuously instead of relying solely on point-in-time authentication rules.

Risk signals become substantially more useful when shared across enterprise systems in near real time. Suspicious authentication attempts, repeated verification failures, unexpected profile modifications, or abnormal transaction behavior can trigger immediate notifications that let downstream systems apply additional verification or temporarily restrict high-risk activity. This coordinated response improves detection while minimizing unnecessary friction for customers whose behavior remains consistent with their established pattern.

2.3 Identity Events Across Enterprise Systems

Financial institutions typically operate many applications responsible for onboarding, authentication, payments, investments, lending, customer service, and fraud management. Each application continuously generates identity-relevant information. When these events remain isolated within individual systems, institutions lose the opportunity to correlate information that would reveal a coordinated fraud attempt spanning several applications. Event-driven identity intelligence addresses this by distributing significant identity events across enterprise systems through standardized messaging mechanisms, an approach with direct roots in the enterprise-messaging patterns literature: publish-subscribe channels, event message

formats, and message routing components originally formalized for general enterprise integration apply directly to identity-event distribution [11].

Sharing identity events across business applications improves both operational efficiency and fraud-detection accuracy. Authentication systems, fraud-monitoring platforms, customer-profile services, and transaction-processing applications all benefit from access to current identity intelligence, which keeps security decisions coordinated across the enterprise rather than isolated within individual applications.

Table 1. Identity Event Categories and Fraud-Relevant Signal

Event Category	Representative Events	Primary Fraud Signal
Lifecycle	Onboarding, profile update, account closure	Identity state change requiring re-evaluation of trust level
Authentication	Login attempt, MFA challenge, credential reset	Deviation from established authentication pattern
Behavioral	Device change, geolocation shift, session timing anomaly	Inconsistency with historical behavioral baseline
Account recovery	Password reset, contact-channel change, security-question reset	Common precursor to account takeover
Transactional	Payment initiation, beneficiary addition, high-value transfer	Downstream action a fraudster ultimately seeks to complete

3. Event-Driven Security Architecture

3.1 Event-Driven Processing

Traditional security architectures frequently depend on scheduled synchronization processes that periodically exchange customer information between applications. Although adequate for routine administrative updates, these approaches introduce delays that fraudsters can exploit before security controls respond. Event-driven processing addresses this by letting systems react immediately when an identity or security event occurs: authentication results, profile changes, fraud alerts, and verification outcomes are emitted as discrete events, and subscribing systems process new information without waiting for a scheduled synchronization interval [8]. This log-centered pattern, in which a durable, ordered event log becomes the system of record that downstream services consume independently, was formalized in the stream-processing and data-integration literature as an alternative to point-to-point batch synchronization [9].

This reactive architecture improves responsiveness in distributed financial environments. Rather than continuously polling multiple systems for updates, applications receive identity events only when meaningful changes occur, reducing unnecessary computational overhead while supporting scalable communication across large enterprise platforms. From a systems-engineering perspective, designing event-driven architectures requires balancing throughput, latency, and reliability so that identity intelligence remains available for timely fraud detection [10]. Enterprise financial platforms may process millions of authentication requests, identity updates, transaction events, and behavioral signals daily, which requires distributed event-processing mechanisms capable of sustaining high throughput without degrading responsiveness. Asynchronous processing, workload distribution, resilient event routing, and parallel event handling help maintain consistent performance during peak transaction periods while minimizing processing delay, and reducing end-to-end latency between event generation, risk evaluation, and security response strengthens both operational resilience and customer protection.

3.2 Identity Synchronization Across Distributed Systems

Maintaining synchronized identity information across many enterprise applications is one of the more persistent challenges in financial services. Customer identities evolve continuously as verification activities occur, authentication outcomes change, communication channels are updated, or fraud investigations surface elevated risk. Without reliable synchronization, different systems may operate on conflicting identity information, creating vulnerabilities that fraudsters can exploit. Event-driven identity synchronization provides a consistent mechanism for propagating identity updates across distributed environments while preserving a coherent representation of customer trust; this is the same architectural problem addressed generally in the distributed-systems literature under the heading of maintaining consistency and availability across replicated, independently evolving services [10].

Effective synchronization also supports enterprise governance by reducing duplicate verification effort and ensuring security policies apply consistently regardless of which application the customer uses. Distributed identity intelligence lets authentication platforms, profile systems, fraud-detection services, and transaction-processing applications reference the same current identity state, strengthening operational integrity and the overall effectiveness of fraud prevention.

From an engineering standpoint, synchronized identity information across distributed financial systems requires attention to consistency, availability, and fault tolerance. Because multiple services independently handle authentication, profile management, transaction authorization, and fraud monitoring, temporary service interruptions or communication delays should not compromise overall identity integrity. Engineering synchronization mechanisms that tolerate partial failures while preserving accurate identity state, an approach consistent with the resilience patterns described for independently deployable services [13], lets institutions maintain continuous security operations without introducing unnecessary operational complexity.

3.3 Security Event Correlation

Individual security events often carry limited signal on their own. A single failed login attempt or profile modification may appear insignificant in isolation, yet several related events occurring across different applications can reveal a coordinated attack. Security event correlation combines identity intelligence from multiple enterprise systems into a more complete picture of customer activity and emerging threats, a technique with roots in the complex-event-processing literature, which formalized the detection of meaningful patterns across streams of otherwise low-level events in distributed enterprise systems [12]. Correlating authentication events, profile changes, device registrations, account-recovery attempts, and transaction activity lets institutions detect fraud scenarios that would remain undetected if each event were evaluated independently.

Event correlation also supports adaptive security responses by weighing the cumulative significance of multiple risk indicators rather than relying solely on isolated rules. As identity intelligence becomes more interconnected across enterprise platforms, fraud-detection systems gain the ability to identify relationships between events that individually appear benign but collectively indicate elevated risk, improving detection accuracy while reducing unnecessary false-positive investigations.

Engineering scalable event-correlation platforms requires processing large volumes of concurrent identity events from multiple enterprise applications. Rather than evaluating events sequentially, distributed correlation pipelines aggregate authentication activity, profile changes, transaction behavior, and device intelligence in parallel to support near-real-time risk assessment. Efficient workload distribution, scalable event aggregation, and low-latency decision processing let fraud-detection systems continuously evaluate evolving identity confidence while maintaining predictable performance under increasing transaction volume, considerations that are fundamental where security decisions must remain both timely and operationally reliable.

4. Fraud Detection Applications

4.1 Account Takeover Prevention

Account takeover remains one of the more damaging forms of financial fraud because it lets an unauthorized party gain control of a legitimate customer account using compromised credentials, manipulated authentication mechanisms, or social engineering. Event-driven identity intelligence strengthens account-takeover prevention by continuously monitoring authentication events, credential changes, device registrations, communication-channel modifications, and other identity-related activity that may indicate unauthorized access. Rather than evaluating these activities independently, event-driven systems correlate multiple signals to determine whether authentication confidence has changed enough to justify additional verification, consistent with the continuous-verification principle underlying zero-trust identity architectures [6].

This proactive approach lets institutions intervene before a fraudulent transaction completes rather than only after losses are detected. Elevated-risk identity events can automatically trigger step-up authentication, temporary transaction restrictions, or manual security review, while legitimate customers continue normal account usage under low-risk conditions.

4.2 Authentication Anomaly Detection

Authentication anomalies are often the earliest indicators of fraud. Unusual login locations, unexpected device changes, abnormal authentication frequency, repeated verification failures, or inconsistent behavioral patterns

may appear inconclusive individually, yet collectively suggest elevated identity risk. This mirrors a broader finding in the authentication-research literature: no single password-replacement or authentication scheme dominates simultaneously on usability, deployability, and security, which is precisely why layered, continuously evaluated signals outperform any single static check [7]. Event-driven identity intelligence continuously evaluates authentication events as they occur, letting fraud-detection systems recognize suspicious activity before an unauthorized transaction completes.

Unlike static authentication models, event-driven anomaly detection continuously incorporates newly available behavioral information into ongoing identity evaluation, an approach that federated behavioral-modeling research has shown can reach useful accuracy (up to roughly 85.7% in one recent evaluation) even while preserving data-privacy constraints across distributed data sources [16]. This adaptive process improves an institution's ability to distinguish legitimate behavior from sophisticated fraud while reducing unnecessary authentication challenges for customers demonstrating consistent, trusted behavior.

4.3 High-Risk Identity Events

Certain identity events carry materially greater security implications than routine customer activity. Changes to sensitive identity attributes, communication channels, authentication credentials, account-recovery requests, or verified customer information can substantially alter overall identity confidence and therefore require immediate evaluation. Event-driven architectures ensure that these high-risk identity events reach relevant enterprise systems without delay, letting fraud-prevention mechanisms respond proportionally to changing risk, consistent with the tiered-assurance model in the identity-proofing standards literature, where higher-consequence transactions warrant a correspondingly higher assurance level [4], [5].

Immediate awareness of high-risk identity events lets institutions implement adaptive security policies that reflect current identity confidence rather than historical verification outcomes. Sensitive operations may require additional identity proofing, transaction-approval workflows, or temporary restrictions until sufficient evidence restores confidence in the customer's identity, an approach that strengthens security while preserving operational flexibility across large-scale financial systems.

Table 2. Event-Driven Fraud Detection Use Cases and Architectural Requirements

Use Case	Primary Identity Events Consumed	Key Architectural Requirement
Account takeover prevention	Credential change, device registration, channel modification	Low-latency correlation across authentication and profile services
Authentication anomaly detection	Login pattern, geolocation, behavioral biometrics	Continuous stream processing rather than batch scoring
High-risk event escalation	Contact-channel change, recovery request, sensitive attribute update	Guaranteed, ordered delivery to downstream risk-evaluation systems
Coordinated-attack detection	Cross-system authentication and transaction events	Enterprise-wide event correlation, not single-system monitoring

5. Operational and Security Benefits

5.1 Reduced Fraud Response Time

One of the more significant advantages of event-driven identity intelligence is its ability to reduce the time required to identify and respond to emerging fraud. Traditional fraud prevention frequently depends on scheduled monitoring or manual investigation, both of which introduce delay between suspicious activity and institutional response. Event-driven architectures reduce much of this delay by distributing identity intelligence immediately after relevant events occur, letting security systems evaluate new information in near real time. From a systems-engineering perspective, minimizing end-to-end processing latency between event generation, risk analysis, and enforcement is essential to protecting financial assets. Low-latency event processing lets authentication services, fraud-detection platforms, and transaction-authorization systems coordinate responses before fraud can escalate, improving operational effectiveness while preserving a seamless customer experience.

5.2 Improved Enterprise Consistency

Large financial institutions operate many interconnected applications responsible for authentication, customer management, lending, investments, payments, and fraud prevention. Maintaining consistent identity information across these systems is essential to accurate security decision-making. Event-driven identity intelligence improves enterprise consistency by ensuring significant identity events propagate uniformly across distributed environments, reducing conflicting identity states, minimizing duplicate verification effort, and strengthening governance across the organization's digital ecosystem [16].

5.3 Scalable Security Operations

As financial institutions expand digital services, fraud-prevention systems must process growing volumes of identity-related events without compromising responsiveness or reliability. Event-driven architectures naturally support scalability because identity intelligence can be distributed asynchronously across processing components. Engineering scalable financial systems requires maintaining high throughput while ensuring predictable response times as customer populations, authentication requests, and transaction volumes grow. Workload distribution, parallel event processing, elastic resource allocation, and resilient communication mechanisms let enterprise platforms sustain continuous fraud detection under varying operational conditions, providing a durable foundation for infrastructures that require both high performance and continuous security. These benefits are not free of tradeoffs. Asynchronous, eventually consistent architectures can, by design, allow a brief window in which different services hold slightly different views of identity state; institutions adopting event-driven fraud detection must therefore make an explicit engineering decision about how much staleness is acceptable for a given risk decision, and compensate with mechanisms such as event replay, idempotent processing, and ordered delivery guarantees where staleness cannot be tolerated [9], [10]. [19].

6. Discussion

Considered together, the identity-event taxonomy, the architectural patterns that distribute those events, and the fraud-detection use cases they enable point to a single underlying principle: event correlation across previously isolated systems, not any individual detection algorithm, is what allows an institution to recognize a coordinated fraud pattern earlier than a system evaluating events one at a time. This is consistent with both the identity-assurance standards literature, which increasingly frames assurance as a continuously reassessed property rather than a one-time determination made at enrollment [4], [5], [6], and the complex-event-processing literature, which shows that meaning frequently exists only at the level of the pattern across events, not in any single event examined alone [12].

At the same time, the review's own source material shows that the architectural properties producing these fraud-detection benefits are the same properties that introduce engineering risk. Reduced fraud response time is a direct consequence of asynchronous, log-centered event distribution [8], [9], but asynchronous distribution is also what makes eventual consistency, rather than immediate consistency, the default state of a distributed identity system [10]. An institution adopting event-driven identity intelligence is therefore not simply adding a new detection capability; it is accepting a different, more complex failure model in exchange for lower latency, and the engineering discipline required to manage that tradeoff (idempotent event handling, ordered or replayable delivery, and resilient routing under partial failure) is as important to fraud-prevention outcomes as the detection logic itself [11], [13].

The review's scope is a synthesis of published standards, software-engineering literature, and empirical behavioral-authentication results rather than an evaluation of a single deployed system, so its conclusions describe architectural principles and their tradeoffs rather than measured fraud-reduction outcomes at a specific institution. Direct cross-study comparison of the behavioral-authentication accuracy figures cited here [14], [15], [16], [17], [18] is also limited, because each study uses a different population, feature set, and evaluation protocol; the figures are reported here as evidence that continuous behavioral signal is informative, not as a benchmark comparison between methods [20], [21], [22].

7. Emerging Research Opportunities

7.1 Predictive Fraud Analytics

Future fraud-prevention systems are expected to move beyond reactive detection toward predictive risk assessment that identifies emerging fraud patterns before malicious activity fully develops. Predictive fraud analytics combines historical identity intelligence, behavioral modeling, transaction analysis, and evolving risk

indicators to estimate the likelihood of future fraudulent behavior. Integrating predictive analytics with event-driven identity intelligence would let institutions strengthen proactive security strategies while improving resource allocation for fraud-investigation teams.

7.2 Real-Time Identity Intelligence

Real-time identity intelligence represents a further evolution in enterprise identity management by enabling continuous evaluation of customer trust using dynamically updated identity information. Rather than relying on static customer records, future financial systems are likely to maintain continuously evolving identity profiles that incorporate verification outcomes, authentication history, behavioral consistency, and emerging fraud indicators. These models would support more accurate authentication decisions while enabling adaptive security policies that respond proportionally to changing risk.

7.3 Engineering Next-Generation Fraud Detection Platforms

Future financial systems will likely require fraud-detection platforms capable of processing continuously expanding volumes of identity, behavioral, and transactional information while maintaining low latency and high operational reliability. Continued work in this area will likely focus on distributed architectures that support adaptive risk assessment, intelligent event orchestration, resilient processing pipelines, and scalable decision support, all of which must integrate identity intelligence across multiple enterprise services while preserving system performance and operational resilience.

Emerging research may explore techniques that improve event prioritization, distributed risk computation, adaptive workload management, and resilient processing under rapidly changing operational conditions. As financial ecosystems continue expanding through digital banking, open finance, and interconnected financial services, engineering scalable, real-time fraud-detection platforms will remain essential to maintaining secure, responsive, and trustworthy financial infrastructure.

8. Conclusions

Event-driven identity intelligence reframes fraud prevention around continuous, enterprise-wide awareness of identity state rather than periodic, siloed checks. This review has organized the practice into the identity events that carry fraud-relevant signal, the architectural patterns (event-driven processing, distributed synchronization, and event correlation) used to distribute and act on those events, and the fraud-detection use cases they enable, including account-takeover prevention, authentication-anomaly detection, and high-risk event escalation. The review's central finding is that the fraud-detection value of this approach comes primarily from correlating signals across systems that were previously isolated, and that the engineering tradeoffs inherent to asynchronous, eventually consistent architectures must be managed deliberately rather than treated as incidental. Predictive analytics, continuously evaluated identity confidence, and resilient event-processing platforms represent the most active directions for continued development in this area.

Ethics Statement

Conflict of Interest: The author declares no conflict of interest. Human/Animal Rights: Not applicable. This review did not involve human participants, animal subjects, or identifiable personal data; all figures and findings cited are drawn from previously published sources. Informed Consent: Not applicable, for the same reason. Data Availability Statement: Not applicable. This review synthesizes previously published literature and does not generate or analyze a primary dataset.

CRediT Author Statement

Bhavya Reddy Maryala: Conceptualization, Literature Review, Formal Analysis, Writing - Original Draft, Writing - Review and Editing.

References

[1] Javelin Strategy & Research, "2025 Identity Fraud Study," Javelin Strategy & Research, 2025.
[2] Federal Trade Commission, "Consumer Sentinel Network Data Book 2025," Federal Trade Commission, 2025.
[3] Federal Bureau of Investigation Internet Crime Complaint Center, "2025 Internet Crime Report," Federal Bureau of Investigation, 2026.
[4] National Institute of Standards and Technology, "Digital Identity Guidelines: Identity Proofing and Enrollment," NIST Special Publication 800-63A, 2025.
[5] National Institute of Standards and Technology, "Digital Identity Guidelines: Authentication and Authenticator Management," NIST Special Publication 800-63B, 2025.

[6] S. Rose, O. Borchert, S. Mitchell, S. Connelly, "Zero Trust Architecture," NIST Special Publication 800-207, National Institute of Standards and Technology, 2020.
[7] J. Bonneau, C. Herley, P. C. van Oorschot, F. Stajano, "The Quest to Replace Passwords: A Framework for Comparative Evaluation of Web Authentication Schemes," in Proceedings of the IEEE Symposium on Security and Privacy, pp. 553-567, 2012. (conference style)
[8] M. Fowler, "What do you mean by 'Event-Driven'?", martinfowler.com, 2017. Available at: https://martinfowler.com/articles/201701-event-driven.html (Accessed on: 21 August 2026). (web style)
[9] J. Kreps, I Heart Logs: Event Data, Stream Processing, and Data Integration, O'Reilly Media, Sebastopol, CA, 2014. (book style)
[10] M. Kleppmann, Designing Data-Intensive Applications: The Big Ideas Behind Reliable, Scalable, and Maintainable Systems, O'Reilly Media, Sebastopol, CA, 2017. (book style)
[11] G. Hohpe, B. Woolf, Enterprise Integration Patterns: Designing, Building, and Deploying Messaging Solutions, Addison-Wesley, Boston, MA, 2003. (book style)
[12] D. Luckham, The Power of Events: An Introduction to Complex Event Processing in Distributed Enterprise Systems, Addison-Wesley, Boston, MA, 2002. (book style)
[13] S. Newman, Building Microservices: Designing Fine-Grained Systems, 2nd ed., O'Reilly Media, Sebastopol, CA, 2021. (book style)
[14] S. Kumari, K. Singh, T. Khan, M. M. Ariffin, S. K. Mohan, D. Baleanu, A. Ahmadian, "A Novel Approach for Continuous Authentication of Mobile Users Using Reduce Feature Elimination (RFE): A Machine Learning Approach," Mobile Networks and Applications, vol. 28, no. 2, pp. 767-781, 2023.
[15] A. Martins, T. Dias, A. Dias, J. Vitorino, E. Maia, I. Praça, "Keystroke Dynamics for Intelligent Biometric Authentication with Machine Learning," Discover Applied Sciences, vol. 7, no. 9, 2025.
[16] H. H. X. Nguyen, T. K. Dang, "A Privacy-Preserving Federated Learning Model for Money Laundering Detection," SN Computer Science, vol. 6, art. 586, 2025.
[17] Konatham, M. R., Uddandarao, D. P., Vadlamani, R. K., and Konatham, S. K., "Federated learning for credit risk assessment in distributed financial systems using BayesShield with homomorphic encryption," in Proc. 2025 International Conference on Computing Technologies & Data Communication (ICCTDC), pp. 1–6, IEEE, 2025, doi: 10.1109/ICCTDC64446.2025.11158863.
[18] D. P. Uddandarao, Statistics and the Science of Causal Economics: A New Paradigm for Data Science. Deep Science Publishing, 2026.
[19] D. P. Uddandarao, M. R. Konatham, and R. K. Vadlamani, “Robust and scalable statistical models for high-dimensional marketing data applications: A comprehensive review,” in Proc. 2025 5th International Conference on Emerging Research in Electronics, Computer Science and Technology (ICERECT), pp. 1–8, IEEE, 2025, doi: 10.1109/ICERECT65215.2025.11377070.
[20] D. Tokas, “Interoperability standards in modern workflow orchestration systems,” Journal of Information Systems Engineering and Management, vol. 11, no. 2s, pp. 646–654, 2026. [Online]. Available: https://jisem-journal.com/index.php/journal/article/view/14459
[21] H. Gaggar, “Two planes of truth: Execution fidelity and behavioral correctness in tool-using LLM agents,” International Journal of Computational and Experimental Science and Engineering, vol. 10, no. 3, 2024, doi: 10.22399/ijcesen.5498.
[22] S. K. Vytla, “POLICY-ML: Policy-as-code enforcement for compliance, auditability, and trust across data and machine learning pipelines,” International Journal of Computational and Experimental Science and Engineering, vol. 10, no. 4, 2024, doi: 10.22399/ijcesen.5393.