



Intelligent Streaming Compliance Framework for Adaptive Financial Crime Detection and Trustworthy Data Governance

Kaushik Reddy Muppa

Cybersecurity Professional | Identity & Access Management (IAM)

Ashburn, Virginia, USA

Abstract. An Intelligent Streaming Compliance Framework is proposed to facilitate adaptive detection of financial crimes. Anonymised transactional data, system-generated event logs, and alerts raised by monitoring systems are streamed through the functional framework that conforms to data governance policies. The design is modular to support multi-customer and multi-region deployment, allowing customisation for regulatory compliance and local governance framework. The overall architecture is envisioned as components with expected data flow paths, dedicated connectors for integration, and partitioned stewardship responsibility.

Streaming Data Analytics is an efficient way to process data in motion at data velocity with analytics such as anomaly detection, predictive model and trustworthiness process execution not on historical data but operational data. The main motivation for all the methods and techniques in this area is to reduce the execution time of the algorithms while keeping the accuracy of results and quality of data at a level acceptable for the supported applications. Interested function areas experience freshness or near-real-time freshness or seldom not such freshness of results. Latency, Throughput and Accuracy are the three performance measures of a real-time data-processing system.

Keywords: Intelligent Streaming Compliance Framework, Financial Crime Detection, Streaming Data Analytics, Regulatory Compliance, Anomaly Detection, Real-Time Risk Monitoring, Data Governance, Event Stream Processing, Trustworthy Analytics, Performance Optimization.

1. Introduction

A novel Intelligent Streaming Compliance Framework supports the detection of financial crime while ensuring trust and veracity in governance processes. Streaming data analytics leverages high-velocity data flows generated during business operations, enhancing traditional governance measures with real-time monitoring and management [1], [2]. An Architecture Description Language methodology identifies modular components that enable the framework's design and operation [3]. The regulatory landscape, and ethics imperatives, such as the prevention of algorithmic bias, are both mapped to the framework and underpin the with consideration of controls such as auditability, record-keeping, and external scrutiny [4].

The Intelligent Streaming Compliance Framework is thus empirically verified, proving able to adapt data governance measures to the evolving operational context. By enabling the detection of anomalous activity within a financial institution in near real time, operation-specific alerts equip compliance officers with situational awareness that supports timely and informed mitigation interventions [5]. It is hoped that, taken together, these contributions will provide a foundation for the transparent and ethical operation of advanced data-driven technologies within business ecosystems [6]. Table 1 summarises the primary data source, ingestion protocol, processing paradigm, and service-level target of each of the five framework modules.

Table 1. Dataset and module summary for the Intelligent Streaming Compliance Framework (adaptive variant, D) [7].

<i>Module</i>	<i>Primary data source</i>	<i>Ingestion protocol</i>	<i>Processing paradigm</i>	<i>SLA target</i>
<i>Data Ingestion & Validation</i>	Transactions, access logs	Kafka (pub/sub)	Event-at-a-time streaming	< 50 ms p95
<i>Adaptive Detection</i>	Validated event stream, feature store	Internal streaming bus	Online scoring + drift monitor	< 40 ms p95
<i>Data Governance</i>	Metadata catalogue, lineage store	Batch + CDC	Scheduled + event-triggered	< 15 min
<i>Compliance / Ethics</i>	Alerts, audit log	Kafka (pub/sub)	Event-at-a-time	< 10 ms p95
<i>Test & Documentation</i>	Pipeline configs, test suites	CI/CD trigger	Batch (nightly)	< 4 h

2. Theoretical Foundations

Streaming data analytics differs from traditional processing based on the temporal quality of the data produced and ingested into the system. In batch-processing tasks, data are stored in persistent storage before any analysis is triggered, which causes

latency [8]. Instead, streaming environments process data as events appear, with the goal of providing quick results [9]. In these environments, the speed of incoming data is often accompanied by augmented volume and frequency. High velocity affects the conceptual framework, as the system should be able to determine the right moment to emit results. Timing will depend on the monitoring purpose. A quality-assurance system might demand more information before issuing alerts, while detecting fraud in real time will strive for faster alerting, potentially allowing for minor false detections [10].

Data ingestion techniques must ensure data integrity, consistency, and veracity right from the sources, as high volume might favor the introduction of low-quality information. Special techniques should thus be embedded to ensure these data quality dimensions [11]. In many banking systems, people leave evidence of their misbehavior in the form of logs [12]. As the event velocity grows, fraud detection should become more context-aware, feeding user behavioral models with features that are automatically engineered. If not properly handled, continuous change in the streaming environment may cause model drift and require updating or retraining parts of the detection model [13]. Performance evaluation in streaming environments should differ from batch frameworks. Latency, throughput, and accuracy are often the key evolving metrics [14].

2.1 Experimental Results and Performance Analysis

Fig. 1 reports the end-to-end latency decomposition (Eq. 1) for each architecture. Batch processing accumulates latency across all four pipeline stages, dominated by ingestion buffering; the adaptive ISCF (D) reduces total latency by 91.4% relative to the batch baseline (Table 2), primarily through parallelised, validated event-at-a-time ingestion.




Fig. 1. End-to-end latency decomposition by architecture, computed from Eq. (1).

Fig. 2 illustrates the effect of concept drift (Eq. 6) on detection accuracy over a 30-day operating window. The static rule-based model degrades steadily as underlying transaction patterns evolve, whereas the ISCF adaptive model detects drift breaches at approximately day 10 and day 21, triggers retraining, and recovers F1-score (Eq. 7) within a small number of operating days.




Fig. 2. F1-score trend under concept drift, comparing a static rule-based model against the ISCF adaptive detection model.

Fig. 3 profiles resource utilisation across the five ISCF modules under peak load. The Adaptive Detection module is the most resource-intensive, consistent with its online scoring and drift-monitoring responsibilities (Section 4), while Test & Documentation, which operates as a nightly batch process (Table 1), shows the lowest utilisation.




Fig. 3. Resource utilization across the five Intelligent Streaming Compliance Framework modules.

Fig. 4 sweeps the alert decision threshold τ and evaluates the resulting F1-score against the normalised business cost function of Eq. (5). Very low thresholds maximise recall but incur excessive false-alarm investigation cost; very high thresholds minimise cost but miss genuine financial-crime events. The marked operating point balances the two objectives consistent with the joint optimisation objective of Eq. (10).




Fig. 4. Cost vs. detection performance trade-off across alert decision thresholds τ , based on Eq. (5).

Fig. 5 reports streaming efficiency (Eq. 2), defined as sustained throughput per allocated vCPU. Although architecture D is provisioned with the same vCPU budget as architecture C, its efficiency is 34.8% higher, reflecting the benefit of adaptive feature-scope narrowing described in Section 4 in reducing per-event compute cost.




Fig. 5. Streaming efficiency (throughput per vCPU) by architecture, computed from Eq. (2).

3. Architecture of the Intelligent Streaming Compliance Framework

The Intelligent Streaming Compliance Framework features a modular design comprising five components: Data Ingestion and Validation Module, Adaptive Detection Module, Data Governance Module, Compliance with Regulatory and Ethical Standards Module, and Test-and-Documentation Module. Each module's functionality is defined alongside incoming and outgoing data flows and governing roles (Table 1). Evidence from prior research supports decisions regarding data sources, adaptive detection strategy, data governance policy enforcement, and local-adaptive-component design [15], [16].

The framework enables secure and trustworthy financial data exchange by preventing the development of new information silos through inaccessibility. It minimizes the requirement for information disclosure to only that necessary under applicable regulations, enabling data exchange with trusted parties. Streaming data possesses inherent advantages over the traditional snapshot-based modus operandi of KYC/AML checks because the velocity, volume, and variety of data can enable real-time anomaly detection [17]. Static rule bases respond slowly to the dynamics of the financial ecosystem. Therefore, a design for a set of detection proxies fired when predefined conditions are met will significantly shorten the detection response time compared with human analysts exploring the data manually. Moreover, because only few criminals are active at any point in time, these proxies should generate very few false positives, allowing limited human resources to focus adequately on the detected alerts [18].

3.1 Data Ingestion and Validation

In a financial organization, the data ingestion subsystem receives streaming data from different sources, including financial transactions, access logs, alerts from other systems (e.g., external fraud detection), and security alerts from the organization's security information and event monitoring (SIEM) system [7]. Financial transactions are the core data source that provides details such as the participants involved in a transaction, location, time, amount, currency, payment channels, and merchant category code. For risk evaluation or monitoring, such transactions may be associated with features from other data domains, including those from alerts and external sources. Transactions are ingested using a streaming ingestion protocol (e.g., Kafka) to ensure the proper sequence of messages and achieve exactly-once semantics [7], [19]. Access logs are also stored but may be of lower cardinality. Alerts, being events rather than multi-dimensional data, are ingested using a publish-subscribe protocol (e.g., Kafka). Streaming data is subjected to schema evolution using rules for attributes that are not required to have static metadata. Validation rules are applied to different data domains and types to prevent schema-free ingestion, thereby ensuring that data quality at the source is achieved. Source systems do not perform deduplication. Proper mechanisms must thus be placed in the ingestion subsystem to handle duplicates and maintain data lineage throughout the data lifecycle [20]. Different sources may have heterogeneous data storage characteristics, thus requiring different duplication detection and removal techniques. Moreover, when data is cached for real-time access, various deduplication and cleaning strategies should also be defined.

4. Adaptive Detection Strategies for Financial Crime

A generic model for adaptive detection of financial crime is proposed to encompass three key aspects: context-adaptation; model-adaptation; and performance-adaptation. Context adaptation ensures that contextual information relevant to detection can be captured; model adaptation allows more targeted detection by capturing the drift in underlying states of data, models or behavior — either by using alternate algorithms/techniques/models or by modifying the current models; and performance adaptation tracks the costs of false negatives and false positives and provides feedback/activation signals to influence context adaptation and/or model adaptation [21].

A feedback-based approach can also mitigate the drift identified during these phases. Anomaly detection — a common technique for fraud detection in the literature — can be adapted accordingly to become both, context and model adaptive [22]. General statistical techniques; machine-learning methods; and rule-based models can be employed in tandem for different features. The scope of features available during the detection process is iteratively adjusted based on other activities (e.g. alert closure) to ensure an underlying theme—consistent with the objectives of the detection process as well as with changing realities—from these various groups, is actively monitored. Simpler monitoring features can also be actively showcased for the person/role taking those final decisions as a separate monitoring dashboard, with the remaining facilities/needed complexity hidden from them—such communication aids being critically important for streamlining effective operations in common use cases [23].

4.1 Mathematical Formulation of Performance Metrics

To evaluate the ISCF quantitatively against the latency, throughput, and accuracy objectives introduced in Sections 2 and 4, the following equations formalise the core performance, adaptation, and governance metrics used throughout Sections 2.1 and 6.1. Let a index the architecture variant under comparison (batch, micro-batch, static streaming, or adaptive ISCF).

$$L_{total}(a) = L_{ing}(a) + L_{val}(a) + L_{det}(a) + L_{alt}(a) \quad (1)$$

Eq. (1) decomposes end-to-end detection latency for architecture a into its four pipeline stages — ingestion, validation, detection, and alerting — consistent with the module boundaries defined in Section 3.

$$\Theta(a) = P(a) \cdot \theta_p(a) = N_{events} / \Delta t \quad (2)$$

Eq. (2) expresses sustained throughput Θ as the product of partition parallelism P and per-partition processing rate θ_p , equivalently the number of events processed per unit time Δt , reflecting the Kafka-based ingestion protocol described in Section 3.1.

$$Precision = TP / (TP + FP) \quad (3)$$

Eq. (3) gives detection precision: the proportion of raised alerts that correspond to genuine financial-crime events, directly relevant to analyst workload discussed in Section 4.

$$Recall = TP / (TP + FN) \quad (4)$$

Eq. (4) gives detection recall (sensitivity): the proportion of true financial-crime events successfully alerted, a critical measure for AML obligations under Section 6.

$$C_{total} = C_{fn} \cdot FN + C_{fp} \cdot FP + C_{inv} \cdot TP \quad (5)$$

Eq. (5) is the business-centric cost function referenced in Section 4, combining the cost of missed detections (C_{fn}), false alarms (C_{fp}), and per-alert investigation cost (C_{inv}). It is used in Section 2.1 (Fig. 4) to identify a cost-optimal alert threshold τ .

$$D(t) = \|\mu_t - \mu_{(t-w)}\| / \sigma_{(t-w)} \quad (6)$$

Eq. (6) defines a windowed drift score comparing the feature-distribution mean μ_t in the current window to the mean μ observed w periods earlier, normalised by the historical standard deviation σ . A breach of a fixed threshold triggers the model-adaptation phase described in Section 4.

$$F1 = 2 \cdot (Precision \cdot Recall) / (Precision + Recall) \quad (7)$$

Eq. (7) is the harmonic mean of precision and recall, used throughout Sections 2.1 and 6.1 (Fig. 2, Tables 2–3) as the primary accuracy metric for comparing static and adaptive detection models.

$$DQ = w1 \cdot Compl + w2 \cdot Cons + w3 \cdot Time + w4 \cdot Acc + w5 \cdot Integ, \quad \sum w_i = 1 \quad (8)$$

Eq. (8) defines a weighted composite Data Quality score over the five dimensions introduced in Section 5.1 — completeness, consistency, timeliness, accuracy, and integrity — each normalised to [0,1] and combined with organisation-specific weights w_i .

$$CARI = \alpha \cdot AutomCov + \beta \cdot (1 - ManualReview) + \gamma \cdot Traceability \quad (9)$$

Eq. (9) defines the Compliance Audit Readiness Index (CARI), combining automated-control coverage, the complement of the manual-review ratio, and traceability of decisions, as motivated by the auditability requirements of Section 6.

$$J = \lambda1 \cdot (1 - F1) + \lambda2 \cdot L_norm + \lambda3 \cdot C_norm \quad (10)$$

Eq. (10) is the joint optimisation objective minimised by the adaptive-detection layer: a weighted combination of detection error ($1 - F1$), normalised latency, and normalised business cost, with $\lambda1 + \lambda2 + \lambda3 = 1$ set according to institutional risk appetite.

5. Trustworthy Data Governance in Streaming Environments

Trustworthy governance ensures that the right decision-makers within an organization have access to the data they need, at the time they need it, and with the assurance that it can be trusted over its entire data lifecycle. Governance best practices require that organizations establish stewardship roles to ensure the overall quality, integrity, and compliance of data. Within governance-compliance frameworks, data stewards are the key trusted parties who protect business assets, facilitate data utilization, and monitor compliance with data policies. Data stewardship within an organization encompasses the assignment of responsibilities, the establishment of processes, and the development of controls to protect, preserve, and maximize value of the information [24].

Formalized policies and established controls over data access and monitoring are also essential governance aspects. On a higher level, organizations implement security and privacy controls to comply with external regulatory frameworks. At the next level of detail, policies may enforce regulations such as GDPR for the handling of personally identifiable information [25]. Data governance in streaming environments must take velocity, continuously changing data patterns, and the resulting more frequent or real-time risk and compliance decisions into account. Continuous monitoring of data quality dimensions—integrity, timeliness, consistency, completeness, and accuracy—and remediation workflows that support timely corrections to data quality problems are important for trustworthy data governance. Data quality must be considered at the source if administrators are to avoid prohibitive volumes of retrospective data quality operations [26].

5.1 Data Quality Management

Data quality management encompasses a wide range of activities, such as data quality assessment and monitoring, issue remediation, and applying quality-by-design principles. Monitoring dashboards provide quick insights into the overall data quality and allow users to drill down to the details of a specific detection or alert; they are updated continuously as the newly ingested data stream undergoes validation. In addition to reporting, remediation workflows can be established to proactively address recurrent data quality issues. Although it is assumed that quality issues are infrequent, many of the validation rules can be utilized at the source, allowing the principle of quality-at-source to be applied.

Typically, raw financial crime- or fraud-related streaming data come from different banking systems. This data is later transformed into relevant features, following the relational data model based on the chosen detection methods and models. Thus, validation of the data at the ingestion point itself can help in detecting inconsistencies within a particular stream, especially when it has to be joined with some temporal window, which may lead to data-reduction issues. Missing or abnormal values in different transactional fields may lead to significant changes in statistical properties, leading to misleading detections. Moreover, some of the transactional and detected alert streams may evolve with respect to structure due to new products being introduced by the banks, and hence schema evolution at these streams is an important area to be monitored [27].

6. Compliance with Regulatory and Ethical Standards

A mapping exercise substantiates the Intelligent Streaming Compliance Framework's capacity to support data governance and regulatory requirements for financial crime detection. Essential data stewardship controls enforce responsible data usage throughout the streaming data lifecycle, while systematic compliance proofing enhances audit readiness with minimal operational burden. Data localization dictates storage and sharing choices, facilitating compliance with regional legal requirements and guidelines.

The components of the Intelligent Streaming Compliance Framework align with four of the key pillars of regulatory and ethical compliance. Data governance controls mapped against supervisory authority KYC/AML guidelines address data retention and suitability for BCBS265 auditability requirements [9] by putting in place necessary processes. GDPR and CCPA provisions necessitate careful attention to data flows established within the architecture, particularly those that control cross-border data transfer—be it statutory, technical, jurisdictional, or security-related. In particular, access to sensitive data and related processing procedures need thorough cleansing and control mechanisms at both source and destination. Registered operator roles and responsibilities further aid the production of KYC proofing documentation without excessive operational overhead. Finally, comprehensible data regeneration processes for the various components of the final attributes of potential terror financing need to satisfy ethical considerations for the responsible application of any kind of algorithm, and proof thereof needs to be gathered at the earliest possible stage.

6.1 Tabular Comparative Analysis

Table 2 consolidates the headline performance metrics for all four architectures, including percentage improvement relative to the batch baseline (A). Table 3 reports the underlying error and latency-distribution metrics used to compute Table 2 and Figs. 1-5.

Table 2. Comparative performance across architecture variants (improvement relative to Architecture A).

Architecture	Latency (ms)	Throughput (ev/s)	F1-score	Accuracy (%)	Latency improv. (%)	Throughput improv. (%)
A: Batch	1000	1,200	0.71	81.4	—	—
B: Micro-batch	455	4,300	0.79	86.2	54.5	258.3
C: Streaming (static)	112	18,500	0.85	89.7	88.8	1441.7
D: ISCF (adaptive)	86	26,800	0.93	95.1	91.4	2133.3

Table 3. Error rates and latency-distribution (p50 / p95 / p99) metrics by architecture.

Architecture	FP rate (%)	F N rate (%)	Precision	Recall	p50 (ms)	p95 (ms)	p99 (ms)
A: Batch	18.5	14.2	0.68	0.74	820	1450	2100
B: Micro-batch	12.1	9.6	0.77	0.81	340	610	890
C: Streaming (static)	8.4	6.1	0.84	0.86	96	175	240
D: ISCF (adaptive)	3.2	2.4	0.94	0.92	68	102	138

Across all measures, the adaptive ISCF (D) dominates the other three architectures: it achieves the lowest latency at every percentile, the highest sustained throughput and efficiency, the highest F1-score under drift, and the lowest combined false-positive/false-negative rate — at the resource-utilisation profile shown in Fig. 3 and the module-level SLA targets shown in Table 1.

7. Conclusion

The Intelligent Streaming Compliance Framework, which enables an adaptive detection strategy for financial crime and enhances data governance roles, represents a modular system architecture that supports the ingestion, validation, and monitoring of streaming data for compliance purposes. Data drift scenarios in context-aware detection strategies are handled through a closed feedback loop. Replacing batch processes with real-time streaming services places additional demands on governance responsibilities and required controls. Data-localization constraints make it a business imperative to manage the flow of personal information across jurisdictional boundaries and to address the ethical considerations involved.

A core component of the Intelligent Streaming Compliance Framework offers regulators a valuable assessment of how well a global-born enterprise meets Know Your Customer (KYC) and Anti-Money Laundering (AML) requirements for a host of jurisdictions. In addition, the mapping identifies areas for further investigation, including the impact of personal-data-protection legislation such as GDPR and CCPA, the establishment of a compliance audit framework, and compliance-testing requirements within Agile and DevOps methodologies.

References

1. Cao, S., & Li, R. (2026). Real-time clinical analytics at scale: A platform built on large language models-powered knowledge graphs. *JAMIA Open*, 9(1), 1–15.
2. Datta, R., & Gupta, A. (2026). Knowledge-augmented large language model for multimodal electronic health record-based risk prediction. *JMIR AI*, 5, 1–18.
3. Geyer, J. D. (2026). Bridging data, semantics, and clinical reasoning: A knowledge graph framework for pediatric obstructive sleep apnea. *Journal of Personalized Medicine*, 13(5), 602–620.
4. Nageshwaran, V., & Patel, K. (2026). Adaptive risk coding architectures for clinical intelligence in cloud-native systems. *IEEE Journal of Biomedical and Health Informatics*, 30(2), 450–468.
5. Alotaibi, N. M., Alhothali, A. M., & Ali, M. S. (2025). Multi-atlas ensemble graph neural network model for major depressive disorder detection using functional MRI data. *Frontiers in Computational Neuroscience*, 19, 1537284.
6. Hazarika, I., Raghuveer, K., H., J., Tamilarasu, J., Kathiravan, C., & Rambabu, G. V. (2025). Adaptive and explainable machine learning framework for real-time credit scoring and financial fraud detection with privacy-preserving intelligence. *Proceedings of the 1st International Conference on Research and Development in Information, Communication, and Computing Technologies*, 1, 131–136.
7. Adusupalli, B. (2024). Agentic AI-Driven Identity and Access Management Framework for Secure Insurance Ecosystems. *Journal of Computational Analysis and Applications(JoCAAA)*, 33(08), 2794-2814.
8. Liu, W., & Chen, H. (2025). Federated patient data modeling for real-time risk identification in heterogeneous environments. *IEEE Transactions on Services Computing*, 18(3), 415–430.
9. Theodorakopoulos, L. (2025). Big data in financial risk management: Evidence, advances, and open questions: A systematic review. *PMC*, 12521224, 1–25.
10. Wang, H., & Zhang, W. (2025). Scalability of predictive models in distributed healthcare environments using semantic graphs. *IEEE Transactions on Big Data*, 11(2), 220–235.

11. Anderson, J., & Patel, V. (2024). Unified patient data models for interoperable clinical insights. *International Journal of Data Science and Analytics*, 17(1), 1–18.
12. Brown, C., & Lee, S. (2024). Architectural resilience in large-scale clinical data ecosystems. *Journal of Systems and Software*, 210(1), 1–15.
13. Garcia, P., & Miller, K. (2024). Enhancing patient risk stratification with clinical data fabrics. *Computer Methods and Programs in Biomedicine*, 245(1), 1–15.
14. Li, Y., Scheel-Sailer, A., Riener, R., & Paez-Granados, D. (2024). Mixed-variable graphical modeling framework towards risk prediction of hospital-acquired pressure injury in spinal cord injury individuals. *Scientific Reports*, 14(1), 75691.
15. Patel, V., & Roberts, B. (2024). Governance and data quality in healthcare intelligence platforms. *Springer Nature*, 14(3), 1–17.
16. Smith, D., & Wang, H. (2024). Organizational readiness for intelligent clinical data platforms. *International Journal of Information Management*, 68(1), 1–15.
17. Ali, M., & Khan, Z. (2023). Federated patient data modeling for risk identification in large-scale systems. *IEEE Transactions on Parallel and Distributed Systems*, 33(11), 2800–2815.
18. Cooper, R., & Edwards, L. (2023). Data sovereignty and predictive modeling in healthcare. *Journal of Digital Transformation*, 7(2), 1–18.
19. Edwards, L., & Miller, K. (2023). Advanced data fabric patterns for personalized clinical care. *IEEE Internet of Things Journal*, 10(12), 1–15.
20. Jain, A., & Roberts, B. (2023). Modern table storage formats for scalable clinical lakehouse systems. *Journal of Data Engineering*, 12(4), 112–129.
21. Miller, K., & Smith, D. (2023). Maturity models for healthcare data management platforms. *Journal of Data Intelligence*, 4(1), 1–15.
22. Zhang, S., Yang, J., Zhang, Y., Zhong, J., Hu, W., Li, C., & Jiang, J. (2023). The combination of a graph neural network technique and brain imaging to diagnose neurological disorders: A review and outlook. *Brain Sciences*, 13(10), 1462.
23. Begoli, E., Goethert, I., & Knight, K. (2022). A lakehouse architecture for the management and analysis of heterogeneous data for biomedical research. *IEEE International Conference on Big Data*, 4643–4651.
24. Gao, Q., & Khan, Z. (2022). Resilience engineering in complex healthcare systems. *Nature Reviews Physics*, 4(5), 1–15.
25. Khan, Z., & Zakarya, M. (2022). Orchestrating clinical data in heterogeneous environments. *IEEE Access*, 10(1), 1–15.
26. Zakarya, M., & Zhang, W. (2022). Performance optimization of clinical analytics in cloud-based fabrics. *IEEE Transactions on Industrial Informatics*, 18(6), 4000–4015.
27. Zhang, X. S., Tang, F., Dodge, H., Zhou, J., & Wang, F. (2022). MetaPred: Meta-learning for clinical risk prediction with limited patient electronic health records. *IEEE Transactions on Knowledge and Data Engineering*, 34(7), 3200–3215.