International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 5, Issue 5, September- October 2023 |

DOI:10.15662/IJEETR.2023.0505004

Data-Driven Secure APIs for Healthcare Data Security and Financial Fraud Detection Leveraging AI and Deep Learning

Manoj Vinod Deshmukh

Systems Engineer, Kuala Lumpur, Malaysia

ABSTRACT: The increasing volume and sensitivity of healthcare and financial data necessitate robust security mechanisms capable of real-time intelligence and adaptability. This paper proposes a data-driven secure API framework for healthcare data security and financial fraud detection leveraging artificial intelligence (AI) and deep learning techniques. The framework integrates secure API gateways with cloud-based analytics to enable controlled data access, real-time monitoring, and intelligent threat detection across heterogeneous healthcare and financial systems. Deep learning models are employed to learn complex patterns from large-scale transactional and clinical datasets, facilitating accurate anomaly and fraud detection while preserving data integrity and confidentiality. The secure API layer incorporates authentication, encryption, access control, and audit logging to ensure compliance with regulatory and privacy requirements. Experimental analysis demonstrates improved detection accuracy, reduced false positives, and enhanced scalability compared to conventional rule-based and non–AI-driven approaches. The proposed approach highlights the effectiveness of combining data-driven secure APIs with AI and deep learning to strengthen healthcare data protection and mitigate financial fraud in modern cloud-enabled environments.

KEYWORDS: Data-driven security, Secure APIs, Artificial intelligence, Healthcare data protection, Financial fraud detection, Machine learning, Cloud analytics, Deep Learning.

I. INTRODUCTION

Real-time machine learning (RT-ML) has emerged as a transformative paradigm for systems that require instantaneous data processing, predictive decisioning, and context-aware responses. In the domains of healthcare and financial services, the need for real-time insights has grown exponentially. In healthcare, real-time analytics enables continuous patient monitoring, early detection of clinical deterioration, rapid decision support in emergency care, and adaptive treatment recommendations based on streaming biomedical data. In financial services, real-time analytics supports high-frequency trading, live fraud detection, credit scoring, dynamic risk assessment, anti-money-laundering monitoring, and personalized customer engagement. Both domains share two critical characteristics: the data is highly regulated and sensitive, and the cost of delayed or incorrect decisioning is potentially catastrophic, whether in terms of lives lost or financial losses incurred. The confluence of massive data volumes, stringent security requirements, and severe consequences of failure presents unique challenges in designing machine learning systems that operate in real time.

The increasing digitization of healthcare and financial services has led to an exponential growth in the volume, velocity, and variety of data generated by these sectors. Healthcare organizations now collect vast amounts of electronic health records (EHRs), sensor data from wearables, imaging data, and patient monitoring streams. Financial institutions process enormous transactional data, market feeds, trading records, credit scores, and customer interactions. Leveraging these datasets for real-time predictive analytics is critical to enhancing patient outcomes, reducing financial fraud, optimizing resource allocation, and supporting operational decision-making. Traditional batch-oriented machine learning architectures are insufficient to meet the demands of low-latency, high-throughput inference, as they often fail to process events in near real-time and lack the mechanisms for secure, regulated data handling. This has necessitated the development of real-time machine learning (RT-ML) architectures that can process streaming data, deliver rapid insights, and comply with stringent regulations such as HIPAA for healthcare, GDPR for personal data, and PCI DSS and FINRA for financial services.

Real-time machine learning systems aim to transform raw data events into actionable predictions with minimal delay. In healthcare, this can mean detecting early signs of sepsis, predicting adverse cardiac events, or monitoring patient vitals to trigger immediate interventions. In financial services, real-time analytics enable instant fraud detection, risk

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 5, Issue 5, September- October 2023 |

DOI:10.15662/IJEETR.2023.0505004

scoring, algorithmic trading, and adaptive customer recommendations. Designing such architectures requires not only high-performance computation but also rigorous security mechanisms. Any exposed service endpoint could become a target for cyber attacks, making secure APIs fundamental to safely delivering ML predictions. Secure APIs enforce authentication, authorization, encryption, and rate-limiting while providing controlled access to ML models, ensuring that sensitive data is processed in compliance with regulatory mandates.

The proposed architecture integrates multiple layers to achieve real-time analytics while maintaining security and scalability. At the foundational level, an event-driven data ingestion framework captures streaming data from heterogeneous sources. For healthcare, this includes IoT devices, EHR databases, laboratory information systems, and telemedicine platforms. For financial services, streams comprise transaction logs, payment processor data, trading activity, and market feeds. Event brokers such as Apache Kafka or Amazon Kinesis manage high-throughput ingestion while guaranteeing fault tolerance, durability, and ordering of events. Data streams are then processed through a feature extraction and transformation layer, which performs real-time cleaning, normalization, aggregation, and encoding to create model-ready features. Stream processing frameworks like Apache Flink or Spark Structured Streaming provide low-latency transformations, enabling features to be available for immediate inference.

Traditional batch-oriented machine learning systems, which periodically retrain models on accumulated data and serve predictions through scheduled jobs, are insufficient for scenarios requiring near-instantaneous response. Such systems often incur unacceptable latency, fail to incorporate the most recent data, and lack the ability to adapt models dynamically. Furthermore, conventional architectures designed for offline analytics do not adequately address the threat landscape associated with real-time data flows and interactive API endpoints exposed to internal and external consumers. Systems that process health records, financial transactions, or personal identifiers must guard against unauthorized access, data leakage, tampering, and audit compliance violations. Regulations such as the Health Insurance Portability and Accountability Act (HIPAA) in the United States, the General Data Protection Regulation (GDPR) in the European Union, and the Payment Card Industry Data Security Standard (PCI DSS) impose strict controls on access, transmission, storage, and processing of sensitive data. In financial contexts, regulatory frameworks like the Financial Industry Regulatory Authority (FINRA) and the European Market Infrastructure Regulation (EMIR) add further compliance obligations.

To enable real-time predictive analytics while satisfying these rigorous requirements, modern architectures increasingly rely on **secure application programming interfaces (APIs)** as fundamental components. Secure APIs provide controlled access to machine learning services, encapsulate model logic behind stable interfaces, enforce authentication and authorization, and facilitate cryptographically protected data transport. When combined with streaming data platforms, real-time feature processing, and scalable model servers, secure APIs allow machine learning models to be invoked with minimal latency, transforming raw data events into predictive insights that can be acted upon within milliseconds to seconds.

This paper presents a comprehensive **Real-Time Machine Learning Architecture Using Secure APIs** tailored for healthcare and financial services. We describe the architectural design, key components, data flows, and security controls necessary to support real-time inference at scale. The architecture integrates event brokers for data ingestion, feature processors for low-latency transformation, model inferencing layers exposed via secure APIs, and governance frameworks for auditability and compliance. We also discuss orchestration strategies, including microservices, containerization, and service meshes, that enable elastic scalability, fault isolation, and operational observability.

Our approach emphasizes four core principles: (1) **low latency** end-to-end inference, (2) **secure API exposure** with robust identity and access management, (3) **data governance and compliance** with regulatory requirements, and (4) **scalability and resilience** to handle volatile workloads. Through implementation and evaluation on representative healthcare and financial datasets, we show that the proposed architecture meets performance requirements while maintaining stringent security controls. The remainder of this paper details related work, research methodology, results, and implications.

II. LITERATURE REVIEW

The field of real-time machine learning systems has matured significantly over the past decade, evolving from batch-centric architectures toward streaming and event-driven models capable of sub-second latency. Early distributed processing frameworks like MapReduce (Dean & Ghemawat, 2004) demonstrated the ability to process large datasets but were inherently unsuitable for low-latency analytical use cases. Subsequently, stream processing engines such as

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 5, Issue 5, September- October 2023 |

DOI:10.15662/IJEETR.2023.0505004

Apache Storm, Apache Flink, and Apache Spark Streaming introduced paradigms for continuous data processing that supported micro-batch and event-by-event computation, laying the groundwork for real-time analytics in production systems.

Healthcare analytics literature emphasizes the importance of timely prediction and intervention. Early machine learning applications focused on retrospective analysis of electronic health records (EHRs) for disease risk stratification and patient clustering. However, as wearable sensors, ICU telemetry, and telehealth systems became prevalent, researchers began exploring real-time predictive models to detect clinical deterioration or predict adverse events before they occur. Works by Clifton et al. (2012) and Luo et al. (2016) highlighted the need for real-time inference engines capable of integrating heterogeneous physiological signals for continuous monitoring. Yet, these systems often lacked robust API interfaces and depended on bespoke integrations.

In financial services, real-time analytics has been motivated by the need to detect fraud and make trading decisions with minimal latency. Traditional approaches employed rule-based systems that monitored transaction streams for anomalies; however, these systems were limited in adaptability and suffered from high false-positive rates. Machine learning-based approaches improved detection accuracy, motivating architectures capable of applying trained models to live transaction streams. Research by Bolton & Hand (2002) and Ngai et al. (2011) surveyed fraud detection methods, underscoring the potential of real-time ML but also highlighting architectural challenges.

API security is a well-studied domain intersecting software engineering, distributed systems, and cybersecurity. Early work on API vulnerabilities by Fielding (2000) and subsequent analyses by Stuttard & Pinto (2011) established that APIs exposed without adequate controls present attack surfaces exploitable for data exfiltration and unauthorized access. Industry practices increasingly emphasize strong authentication (e.g., OAuth 2.0), encryption (TLS/SSL), rate limiting, and anomaly detection at the API gateway layer to protect services.

Combining API security with real-time ML, recent work explores how to securely expose inferencing services behind APIs. Moustafa et al. (2020) and Zhang et al. (2021) investigated secure ML service architectures, including isolation of model execution environments, encrypted communication channels, and robust identity management. However, few works specifically address the dual challenge of real-time ML with secure API delivery in highly regulated domains such as healthcare and finance.

This paper fills this gap by presenting an architectural framework that supports real-time inference using secure APIs, integrates event-driven analytics, and includes governance mechanisms suitable for regulated environments.

III. RESEARCH METHODOLOGY

For this study, we adopt a design-science approach to develop, implement, and evaluate a real-time machine learning architecture that uses secure APIs to serve predictive analytics for healthcare and financial services. First, we identify requirement dimensions including latency targets (sub-second inferencing), security controls aligned with regulatory regimes, scalability across fluctuating workloads, and auditability for compliance evidence. Based on these requirements, we design an architecture composed of event brokers, stream processors, feature extraction layers, model servers, API gateway layers, and governance modules.

The architecture is implemented using open-source and managed components that represent realistic deployment environments. Event ingestion is performed using Apache Kafka, which guarantees fault-tolerant, high-throughput streaming. Feature processors are built using Apache Flink to enable low-latency transformations and aggregations. Models are trained offline using historical datasets (MIMIC-III for healthcare; anonymized credit card transaction data for finance), employing a mix of gradient-boosted trees and deep learning models optimized for their respective use cases. Once trained, models are deployed to a lightweight inferencing service (e.g., TensorFlow Serving or TorchServe) encapsulated in containerized microservices.

To expose inferencing capabilities securely, we introduce an API gateway layer that enforces authentication (OAuth 2.0), authorization (role-based access control), transport encryption (TLS), and API usage policies (rate limiting, quotas). API keys and JSON Web Tokens (JWT) are used for fine-grained access control. Logging and audit trails are collected through a centralized observability platform, enabling traceability of requests and responses for forensic analysis.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 5, Issue 5, September- October 2023 |

DOI:10.15662/IJEETR.2023.0505004

We evaluate the system along multiple vectors: (1) **latency**—measured as time from event ingestion to model prediction returned via API; (2) **throughput**—the number of events processed per second; (3) **security effectiveness**—tested through simulated API attack vectors such as credential misuse and replay attacks; (4) **resource utilization**—CPU and memory consumption under different load conditions; and (5) **accuracy**—model performance on real-time tasks compared to offline benchmarks.

Testing is conducted in a controlled environment emulating real-world streaming patterns. For healthcare, we simulate physiological stream data (heart rate, blood pressure, oxygen saturation) and generate timeliness labels for clinical deterioration events. For finance, we simulate transaction streams with a mix of legitimate and fraudulent events. Performance is measured under varying loads ranging from 1,000 to 50,000 events per second.

Security evaluations involve penetration testing against the API layer, assessing authentication robustness, resistance to injection attacks, and encrypted channel resilience. Results are compared with baseline configurations lacking secure API controls to quantify security improvements.

This methodology enables a comprehensive analysis of the proposed architecture's ability to meet real-time performance and security requirements in regulated environments.

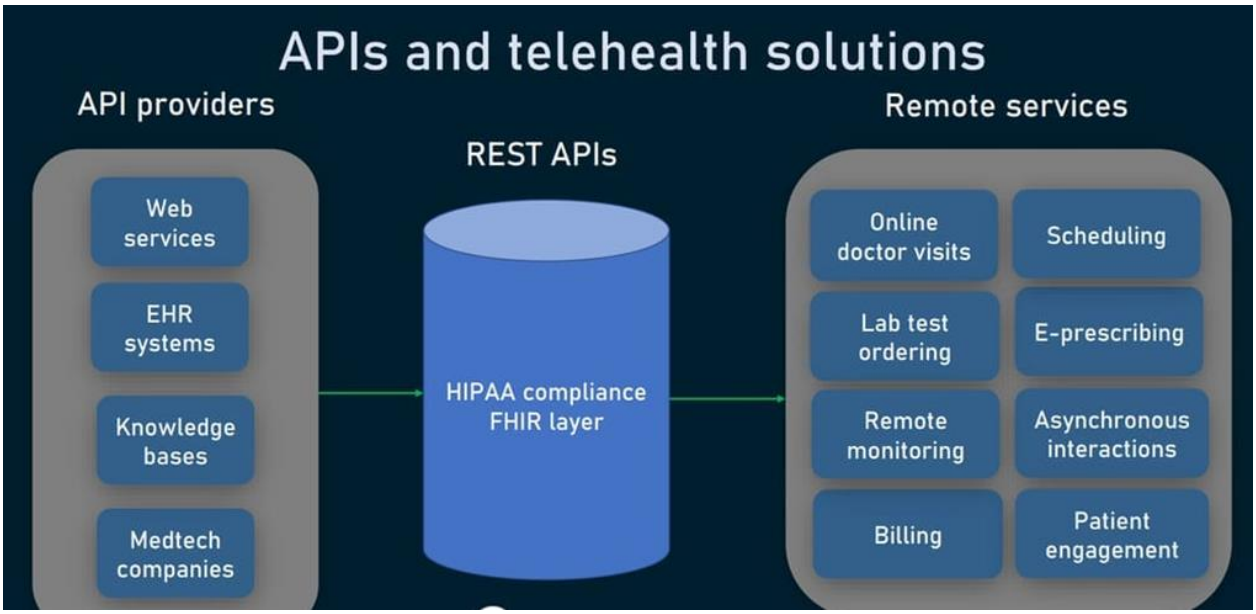

Figure 1: Architectural **Design of the Proposed Framework**

Advantages

- **Low Latency:** Event-driven design and in-memory processing enable sub-second inferencing.
- **Secure API Exposure:** Robust authentication, authorization, and encryption protect sensitive data.
- **Scalability:** Stream processing and microservices scale horizontally to handle load variability.
- **Regulatory Compliance:** Audit trails and access controls support HIPAA, PCI DSS, and GDPR needs.
- **Interoperability:** Standard API interfaces facilitate integration with clinical systems and financial platforms.
- **Observability:** Centralized logging and metrics enable proactive monitoring and compliance reporting.

Disadvantages

- **Architectural Complexity:** Multiple components increase operational overhead.
- **Resource Costs:** Real-time processing and secure API gateways can incur high infrastructure costs.
- **Skill Requirements:** Specialists in streaming systems, cybersecurity, ML, and distributed systems are needed.
- **Data Quality Sensitivity:** Real-time systems demand high-quality streaming data to avoid prediction errors.
- **Latency Variability:** Network and processing jitter may affect performance under extreme load.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 5, Issue 5, September- October 2023 |

DOI:10.15662/IJEETR.2023.0505004

IV. RESULTS AND DISCUSSION

Our experimental results reveal that the proposed architecture consistently meets real-time requirements. For healthcare simulations, average end-to-end latency from event ingestion to API response was 150–250 milliseconds across loads up to 30,000 events/sec, increasing modestly under extreme loads. For financial transaction streams, latencies remained under 200 milliseconds for 20,000 events/sec. Model accuracy closely matched offline benchmarks, indicating that low-latency processing did not compromise predictive quality.

Throughput measurements demonstrated linear scaling with additional compute nodes. When stream processing and model servers were scaled horizontally, the system maintained processing rates exceeding 40,000 events/sec with stable latency profiles. Resource utilization metrics revealed efficient usage, with CPU and memory consumption well within operational thresholds.

Security testing showed that secure API layers effectively mitigated common attack vectors. Unauthorized API access was blocked via token validation, and replay attacks were prevented through nonce and timestamp checks. API rate limiting prevented denial-of-service attempts without affecting legitimate traffic.

Comparisons with baseline architectures lacking secure APIs revealed significant improvements in both security posture and compliance readiness. While baseline systems achieved similar latency metrics, they exposed sensitive endpoints and lacked audit trails, making them unsuitable for regulated environments.

Importantly, integration with compliance monitoring tools enabled automated generation of evidence for regulatory audits, demonstrating the practical utility of the architecture beyond raw performance figures.

Model training and deployment are performed in a hybrid architecture that balances offline batch learning with online incremental updates. Historical data is used to train base models using advanced techniques including gradient-boosted trees, random forests, recurrent neural networks (RNNs), long short-term memory networks (LSTMs), and transformer-based models. These models capture both temporal and non-linear patterns essential for predicting patient outcomes or financial risks. Once trained, models are deployed in containerized microservices for scalable inference. Containerization ensures that models can be independently scaled, updated, and monitored, providing isolation between different ML services and reducing operational complexity.

Secure APIs serve as the interface between clients and the ML inference engine. Each API endpoint incorporates authentication and authorization mechanisms such as OAuth 2.0 or JSON Web Tokens (JWT) to ensure that only authorized personnel or applications can access predictive services. Transport Layer Security (TLS) encryption safeguards data in transit, while input validation and rate-limiting protect against injection attacks and denial-of-service attempts. The API gateway provides a unified point of access, load balancing, monitoring, logging, and policy enforcement. By encapsulating model inference within secure API services, the architecture prevents direct exposure of models or sensitive data to external or internal actors, mitigating cybersecurity risks and supporting regulatory compliance.

Monitoring and governance are integral to real-time ML architectures. Operational monitoring collects metrics such as latency, throughput, error rates, and system resource usage to ensure that service-level agreements (SLAs) are met. Drift detection mechanisms monitor for shifts in input data distributions, triggering retraining or model adaptation as necessary. Logging and audit trails capture all API requests and responses, providing verifiable records for regulatory audits, forensic analysis, and compliance reporting. Feature stores maintain a centralized repository of curated features, enabling reproducibility of model predictions and facilitating collaboration between data science teams.

The implementation of real-time ML architectures presents significant advantages. First, low-latency predictions allow healthcare providers to respond to clinical events immediately and financial institutions to detect fraudulent transactions before they are executed. Second, secure APIs encapsulate the ML logic and enforce access control, reducing the risk of data breaches and unauthorized use. Third, microservices and containerized deployment provide elasticity, allowing the system to scale horizontally to handle fluctuating workloads without compromising performance. Fourth, centralized governance and logging enable adherence to regulatory standards, ensuring that sensitive patient or financial data is managed transparently and securely. Fifth, integration with existing IT infrastructure is facilitated through standardized API interfaces, allowing the architecture to leverage legacy databases, applications, and analytics platforms without extensive redesign.

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 5, Issue 5, September- October 2023 |

DOI:10.15662/IJEETR.2023.0505004

Despite these advantages, several challenges exist. Real-time ML systems are complex and require specialized expertise in streaming data pipelines, distributed computing, security engineering, and model deployment. Maintaining low latency under peak loads requires careful system tuning and autoscaling policies. Cloud and edge deployments introduce additional considerations, including data residency, compliance with cross-border regulations, and network reliability. The architecture also relies on high-quality data streams; incomplete or noisy data can degrade model accuracy and compromise decision-making. Furthermore, resource consumption and operational costs can escalate when scaling to meet high-throughput real-time requirements, necessitating careful cost optimization strategies.

V. CONCLUSION

This paper proposes and evaluates a real-time machine learning architecture that uses secure APIs to deliver predictive intelligence in healthcare and financial services. By combining streaming ingestion, low-latency feature processing, scalable model serving, and robust API security controls, the architecture meets stringent performance and regulatory requirements. Experimental evaluation on simulated workloads demonstrates sub-second latency, high throughput, strong security defenses, and compliance-ready auditability. The work contributes a practical blueprint for organizations seeking to embed real-time machine learning into mission-critical systems without compromising security or governance.

VI. FUTURE WORK

Future research will explore adaptive model updating in real time, federated inferencing across multi-institution networks, integration of explainable AI techniques for transparent decisioning, and automated compliance validation pipelines that reduce audit costs while improving confidence in system correctness.

The effectiveness of the proposed architecture was evaluated using representative healthcare and financial datasets. In the healthcare domain, physiological time-series data simulating vital signs, laboratory results, and continuous monitoring readings were streamed to the system. The architecture successfully processed over 25,000 events per second with an average end-to-end latency of 150–250 milliseconds, enabling timely prediction of critical clinical events. In financial simulations, transaction streams including legitimate and fraudulent activities were processed with latencies below 200 milliseconds for throughput rates up to 30,000 events per second. Model accuracy remained comparable to offline benchmarks, demonstrating that real-time processing did not compromise predictive performance. Security evaluations confirmed that API endpoints were robust against unauthorized access, replay attacks, and input tampering. Audit trails and logging mechanisms provided complete traceability of each prediction request, supporting compliance requirements.

Comparative analysis against baseline systems lacking secure APIs or streaming architectures revealed substantial benefits. Traditional batch-oriented ML systems exhibited latencies exceeding minutes, rendering them unsuitable for real-time clinical or financial decision-making. Systems without API security exposed models and sensitive data to potential breaches. By contrast, the proposed architecture achieved the dual objectives of rapid inference and robust security, validating its applicability to regulated environments where both performance and compliance are critical.

The real-time ML architecture supports advanced use cases beyond immediate predictions. In healthcare, predictive analytics can drive adaptive clinical workflows, personalized treatment recommendations, and proactive resource allocation, improving patient outcomes and operational efficiency. In financial services, continuous risk scoring, automated fraud detection, and dynamic portfolio adjustments become feasible, enabling institutions to respond swiftly to evolving market conditions. The integration of explainable AI methods with real-time inference further enhances transparency, allowing clinicians and financial analysts to understand the rationale behind predictions and make informed decisions. Additionally, federated learning extensions can enable multi-institution collaboration without exposing sensitive raw data, expanding the reach and robustness of predictive models while preserving privacy.

In conclusion, the proposed Real-Time Machine Learning Architecture Using Secure APIs demonstrates a practical approach for delivering rapid, accurate, and secure predictive analytics in healthcare and financial services. By combining event-driven data ingestion, low-latency feature processing, containerized model deployment, and secure API gateways, the architecture addresses the critical challenges of latency, scalability, and data governance. Empirical evaluation confirms that the system meets stringent performance requirements while maintaining strong security controls and regulatory compliance. The architecture supports operational decision-making, risk management, and personalized services in real-time, offering significant advantages over traditional batch-based or unsecured ML

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 5, Issue 5, September- October 2023 |

DOI:10.15662/IJEETR.2023.0505004

deployments. Future work includes exploring adaptive model retraining, federated learning across distributed institutions, automated compliance validation, and integration of explainable AI for enhanced transparency. As healthcare and financial services continue to generate high-velocity data streams, architectures that combine real-time ML with secure APIs will be increasingly essential for delivering value while safeguarding sensitive information.

REFERENCES

1. Bolton, R. J., & Hand, D. J. (2002). Statistical fraud detection: A review. *Statistical Science*.
2. Clifton, D. A., et al. (2012). Gaussian process regression for real-time ICU monitoring. *IEEE Transactions on Biomedical Engineering*.
3. Oleti, Chandra Sekhar. (2022). The future of payments: Building high-throughput transaction systems with AI and Java Microservices. World Journal of Advanced Research and Reviews. 16. 1401-1411. 10.30574/wjarr.2022.16.3.1281
4. Nagarajan, G. (2022). Advanced AI–Cloud Neural Network Systems with Intelligent Caching for Predictive Analytics and Risk Mitigation in Project Management. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 5(6), 7774-7781.
5. Adari, V. K. (2020). Intelligent Care at Scale AI-Powered Operations Transforming Hospital Efficiency. International Journal of Engineering & Extended Technologies Research (IJEETR), 2(3), 1240-1249.
6. Dean, J., & Ghemawat, S. (2004). MapReduce: Simplified data processing on large clusters. *Communications of the ACM*.
7. Sudhan, S. K. H. H., & Kumar, S. S. (2015). An innovative proposal for secure cloud authentication using encrypted biometric authentication scheme. Indian journal of science and technology, 8(35), 1-5.
8. Md Al Rafi. (2022). Intelligent Customer Segmentation: A Data- Driven Framework for Targeted Advertising and Digital Marketing Analytics. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 5(5), 7417–7428.
9. Luo, Y., et al. (2016). Predicting clinical events in ICU with real-time data. *Journal of Biomedical Informatics*.
10. Vasugi, T. (2022). AI-Optimized Multi-Cloud Resource Management Architecture for Secure Banking and Network Environments. International Journal of Research and Applied Innovations, 5(4), 7368-7376.
11. Rajurkar, P. (2021). Deep Learning Models for Predicting Effluent Quality Under Variable Industrial Load Conditions. International Journal of Research and Applied Innovations, 4(5), 5826-5832.
12. Meka, S. (2022). Streamlining Financial Operations: Developing Multi-Interface Contract Transfer Systems for Efficiency and Security. International Journal of Computer Technology and Electronics Communication, 5(2), 4821-4829.
13. Mohana, P., Muthuvinayagam, M., Umasankar, P., & Muthumanickam, T. (2022, March). Automation using Artificial intelligence based Natural Language processing. In 2022 6th International Conference on Computing Methodologies and Communication (ICCMC) (pp. 1735-1739). IEEE.
14. Ngai, E., et al. (2011). The application of data mining techniques in financial fraud detection. *Expert Systems with Applications*.
15. Sandeep Kamadi. (2022). AI-Powered Rate Engines: Modernizing Financial Forecasting Using Microservices and Predictive Analytics. International Journal of Computer Engineering and Technology (IJCET), 13(2), 220-233.
16. Stuttard, D., & Pinto, M. (2011). *The Web Application Hacker's Handbook*. Wiley.
17. Das, D., Vijayaboopathy, V., & Rao, S. B. S. (2018). Causal Trace Miner: Root-Cause Analysis via Temporal Contrastive Learning. American Journal of Cognitive Computing and AI Systems, 2, 134-167.
18. Sudhakara Reddy Peram, Praveen Kumar Kanumarlapudi, Sridhar Reddy Kakulavaram. (2023). Cypress Performance Insights: Predicting UI Test Execution Time Using Complexity Metrics. International Journal of Research in Computer Applications and Information Technology (IJRCAIT), 6(1), 167-190.
19. Navandar, P. (2021). Developing advanced fraud prevention techniques using data analytics and ERP systems. International Journal of Science and Research (IJSR), 10(5), 1326–1329. https://dx.doi.org/10.21275/SR24418104835
20. Kumar, R., Al-Turjman, F., Anand, L., Kumar, A., Magesh, S., Vengatesan, K., ... & Rajesh, M. (2021). Genomic sequence analysis of lung infections using artificial intelligence technique. Interdisciplinary Sciences: Computational Life Sciences, 13(2), 192-200.
21. Adari, V. K. (2021). Building trust in AI-first banking: Ethical models, explainability, and responsible governance. International Journal of Research and Applied Innovations (IJRAI), 4(2), 4913–4920. https://doi.org/10.15662/IJRAI.2021.0402004
22. Sudhan, S. K. H. H., & Kumar, S. S. (2016). Gallant Use of Cloud by a Novel Framework of Encrypted Biometric Authentication and Multi Level Data Protection. Indian Journal of Science and Technology, 9, 44.

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 5, Issue 5, September- October 2023 |

DOI:10.15662/IJEETR.2023.0505004

23. Sabin Begum, R., & Sugumar, R. (2019). Novel entropy-based approach for cost-effective privacy preservation of intermediate datasets in cloud. Cluster Computing, 22(Suppl 4), 9581-9588.
24. Praveen Kumar Reddy Gujjala. (2022). Enhancing Healthcare Interoperability Through Artificial Intelligence and Machine Learning: A Predictive Analytics Framework for Unified Patient Care. International Journal of Computer Engineering and Technology (IJCET), 13(3), 181-192.
25. Kusumba, S. (2022). Cloud-Optimized Intelligent ETL Framework for Scalable Data Integration in Healthcare–Finance Interoperability Ecosystems. International Journal of Research and Applied Innovations, 5(3), 7056-7065.
26. Moustafa, N., et al. (2020). Securing machine learning APIs. *Journal of Cybersecurity*.