AI-Powered Fraud Detection and Risk Intelligence in Banking and Insurance

Dileep Valiki
Independent Researcher
valiki.dileep.researcher@gmail.com

Abstract

AI/ML technologies are exploited to detect, mitigate, control, and minimize risks in banking activities as they create, transmit, lend, invest, insure, and secure value. AI/ML paradigms and technologies based on microservice architecture support dynamic delivery of probabilistic native support of banking actions and transactions. Notable categories of AI/ML technologies are based on risk intelligence, fraud detection and prevention, and fraud control. Risk intelligence examines the risk area of the banking ecosystem. Risk intelligence warns participants about the risks involved in their actions or transactions, predicts a credit card operation's logical behaviour in order to set dedicated limits, and determines the best approach to set a limit for a credit score. Fraud detection identifies fraud in real time using analysis of users' attributes and account history, monitor behaviour changes or sudden changes of an attribute, and classify whether an operation will be done by a legit user or not. Fraud prevention detects fraud before it occurs using knowledge of historical cases to find the relationship between data in order to alert when a similar case may happen and detect unusual patterns to create rules for risky operations. Fraud control deals with the management of fraudulent approaches, actions or transactions for the banking sector. Practices for fraud control address identifying tendencies in the banking sector for fraud detection—vent consumers' intolerance for crime, minimize social disparity, increase competence to minimize price of fraud.

Keywords : Risk; fraud; banking; insurance; risk intelligence; fraud detection; fraud prevention.

1. Introduction

Risk intelligence and fraud control are critical issues for many organizations in banking and insurance services. A cloud-native big-data AI/ML framework is proposed to organize and visualize large volumes of data from banking and financial organizations and help monitor anomalous or fraudulent behaviors, detect warning signs, or develop predictive risk models. The use of cloud-native architecture principles and technologies for data processing and analysis allows servicing fraud detection and prevention needs with fast response times, automated model retraining cycles, and low maintenance costs. Banking sector applications are presented as examples of the framework's design and implementation.

Cloud computing has become the new paradigm for IT and software development. Cloud-native capabilities are integrated into platform, infrastructure, and software-as-a-service solutions. The cloud-native paradigm encompasses the entire life cycle of applications: development, testing, deployment, scaling, and maintenance. Cloud-native applications are managed as a set of microservices residing in containers. Every microservice implements a specific business feature, delivering a required module or function for business specifications. Interactions between microservices are automated and orchestrated by the cloud platform.The concept of fraud health is introduced and correlated to the fraud ecosystem. Once defined, fraud health can also be used for comparing banks or regions for their level of exposure to fraud. Fraud detection can complement fraud failure prevention to continuously assess the fraud ecosystem health of banks and other frauded organizations in micro, meso, and macro areas.

ISSN: 3067-283X
AMERICAN JOURNAL OF ANALYTICS AND ARTIFICIAL INTELLIGENCE
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 19 REVISED: MAY 09 ACCEPTED: MAY 24 PUBLISHED: JUNE 11

Fig 1: Artificial Intelligence (AI) and Machine Learning (ML) within the Financial Services (FS)

1.1. Background and Significance

Achieving resilient growth within alternative banking and insurance ecosystems formed by non-bank providers requires effective risk intelligence and fraud control instead. Fraud is increasingly becoming an organized cybercrime industry that sustains hackers and criminal organizations. Detecting, preventing, and controlling fraud and storm risk through intelligence comes under the umbrella of risk intelligence systems.

Cyber fraud detection, prevention, and control are very elaborate tasks and sensitive due to the volume and variety of incidents happening around the world. While many traditional systems and controls address the problem, it requires to be seen as a fraud ecosystem. An ecosystem approach is a framework that guides the systematic and comprehensive identification, analysis, and resolution of complex problems. An ecosystem of failure constitutes an environment that allows fraud to take place. Hence, fraud failure detection, prevention, and control functions should be facilitated as part of an ecosystem for effective closure.

Fraud ecosystems encompass the factors that influence, lower, and react to real incidents of fraud and, hence, assist in its management.

2. Context and Problem Space

Risk intelligence and fraud control in banking and insurance ecosystems rely on cloud-native architectures and Big Data AI/ML models. Research design involves building an AI-enabled framework using industry standards and best practices, which overlays cloud-native architecture across an ecosystem that includes a cloud provider along with ISV and market infrastructure that create, distribute, and use commercial insurance products. Both risk intelligence and fraud control use data management processes that safeguard data quality by leveraging best-of-breed data management technologies to acquire, cleanse, enrich, and store data in a Data Lake. Fraud detection protects against malicious use. Early-warning monitoring of insurance fraud relies on three types of models that record and synthesize detection-relevant information, examining new data using the information without ongoing reference to historic banking or insurance transactions.

The initial work focuses on applications in the banking sector, where COVID-19 has created new areas for risk and fraud as cybercriminals ramp up their activities. During 2020, the Risk Intelligence & Fraud Control AI ML project leveraged a cloud-native architecture to quickly produce such models, subsequently integrating them into a SaaS framework for banking and insurance ISVs and other financial institutions. The aim is to deploy the analysis results within secure FI ecosystems. Model experimentation and construction are supported by commercial-grade banking-domain data, extensively labelled for fraud detection, and by multiple past incidents. Proven AI ML methods have been applied and tested at a higher level.

2.1. Research design

A multi-disciplinary research design addresses the complex banking and insurance problems holistically by drawing on the fields of AI and cybersecurity in the context of Big Data. Although the design uses the highly visible and impactful domain of fraud detection and prevention to motivate a cloud-native framework for risk intelligence Cyber-AI/ML solutions address multi-purpose requirements of numerous risk-threat models and secure-business operations, the

achieved results are also relevant to internal and external fraud detection and prevention in many sectors.

The requirements for Cyber-AI/ML solutions fall into the three major categories of volumetric detection and mitigation, multi-attribute detection and mitigation, and abnormal event sequence detection and prediction. The Cyber-AI/ML solutions need to operate in response time-frames ranging from real-time to hours-days-weeks, and with both supervised detection-response and unsupervised detection-recovery capabilities. The use of the banking and insurance sectors for demonstration of the framework has the specific objective of showing that SEC, GRC, and ECS data may also be used to expose and stem risk-vulnerabilities of the business ecosystem ecosystem's fraud business ecosystem's fraud control processes.

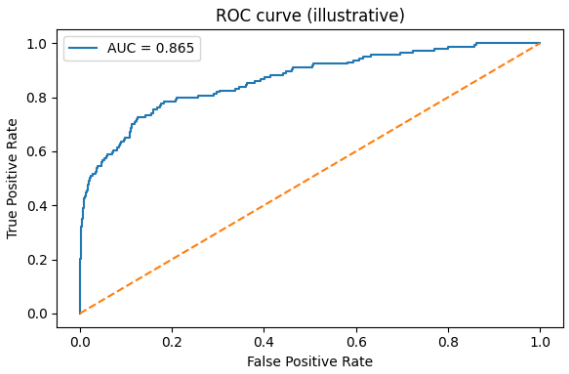

3. Technical Foundations

The framework is designed in accordance with Cloud-Native Architecture principles. The Cloud-Native concept is an evolutionary approach in the definition of applications able to take full advantage of Cloud Computing features and services. The main Cloud-Native characteristics are: Microservices architecture, i.e., applications are built as a set of loosely coupled microservices that can be developed, tested, deployed, and maintained independently of each other; Containerization, i.e., container technology is behind the packaging and isolation of microservices; Dynamic orchestration, i.e., microservices are dynamically created and destroyed according to application demands; Automation, i.e., everything from application deployment to scaling, monitoring, and management tasks must be automated; and Continuous delivery, i.e., DevOps concepts, practices, and tools are used to deliver code changes faster and more reliably by unifying software development (Dev) and operations (Ops). In these applications, the Cloud Provider is assumed to be friendly with the User. For Enterprise-B, the Architecture is dedicated to the Fingerprint Detection service, employing three ElasticSearch Modules for information storage, one S3 module for images storage, and a Pyspark Streaming service for real time fraud detection.

The Cloud-Native Architecture relies on a Cloud Application Provider positioned as a Service Provider that offers to other Enterprises the capability of developing and deploying Cloud-Native Applications. The Provider has a set of Infrastructure as a Service (IaaS) resources that act as a Cloud Infrastructure (IaaS layer) capable of hosting applications developed by enterprises. In this situation, the application can be assumed as a multi-user Cloud-Native application, serving a specific Cloud application Provider that is not the same as the Cloud Infrastructure owner. The Cloud-Native Application has been designed for the Banking sector, aiming the key Business to deliver a Fraud Detection and Prevention Model for Banking Services that operates as an independent service in the Banking Sector.

Equation 1: Risk score as a probability (logistic model)

Step 1: Linear score

Let transaction features be $x \in \mathbb{R}^d$(in the paper's example: $d = 51$)

Cloud-Native Big Data AI_ML Fra…

.
Define a linear model:

$$z = w^\top x + b$$

Step 2: Convert linear score to probability with sigmoid

We need a value in $[0, 1]$. Use the sigmoid:

$$\sigma(z) = \frac{1}{1 + e^{-z}}$$

So predicted fraud probability ("risk score") is:

$$p(y = 1 \mid x) = \sigma(w^\top x + b)$$

Step 3: Why sigmoid "works" (odds / log-odds derivation)

Define **odds**:

$$\text{odds} = \frac{p}{1 - p}$$

Take log:

$$\log\left(\frac{p}{1 - p}\right) = w^\top x + b$$

Now solve for p:

$$\frac{p}{1 - p} = e^{w^\top x + b}$$
$$p = (1 - p)e^{w^\top x + b}$$
$$p = e^{w^\top x + b} - pe^{w^\top x + b}$$
$$p(1 + e^{w^\top x + b}) = e^{w^\top x + b}$$
$$p = \frac{e^{w^\top x + b}}{1 + e^{w^\top x + b}} = \frac{1}{1 + e^{-(w^\top x + b)}}$$

3.1. Cloud-Native Architecture Principles Cloud-Native Big Data AI/ML Framework for Risk Intelligence and Fraud Control in Banking and Insurance Ecosystems

[Technical Foundations: Cloud-Native Architecture Principles] Cloud-native systems are developed and driven by a set of principles. The design takes advantage of the characteristics of cloud computing concepts; implementation exploits the capabilities of cloud technologies; and operation leverages cloud services. Following these guidelines ensures the ease and efficiency of adopting a cloud-native application. In addition, such principles enable cloud regeneration and full-stack contribution.

Cloud-native principles dictate that each system component is actually a small and decoupled application that performs a specific service or function. Units are small enough to be efficiently developed and managed by a single unit. Each component's lifecycle is decoupled from the overall system, allowing the unit to be deployed independently and enabling rapid iteration cycles. Such decomposition naturally leads to the construction of microservices-based heterogeneous architectures that rely on a container ecosystem for implementation, deployment, and execution.

In cloud-native systems, components are constructed by composing cloud services, and code is written only where cloud services do not provide the required functionality. Components interact exclusively through well-defined interfaces, and components are unseen outside of strictly defined execution environments. Tracing, debugging, and system diagnostics tools and services are built in from the inception of each unit without creating intrusive overhead. Components are treated as production systems and as part of the application stack from conception to implementation. No single unit can fail while remaining unnoticed, either through extensive monitoring or testing.

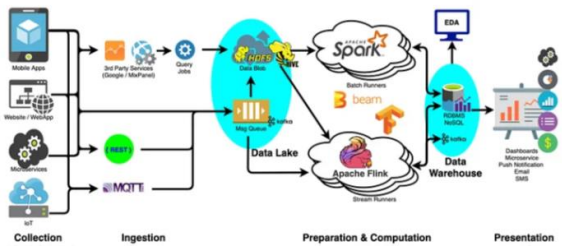

Fig 2: Big Data Architecture for Fraud Detection and Prevention in Banking Industry

4. Risk Intelligence and Fraud Control Paradigms

Technological advancements lead to a plethora of frameworks, methodologies, and techniques focused on intelligent control of risk in modern banking transactions. Such intelligent decisions require a comprehensive feature/indicator set that assists in imposing proactive rules for fraud detection and prevention. To manage and minimize information asymmetry in banking and insurance, risk intelligence and fraud control paradigms have grown in importance. Risk Intelligence refers to the implementation of a combination of Fraud Modeling and Risk Management techniques.

Intelligent Decision Control helps in generating alerts while executing financial transactions and minimizes losses compared to the current traditional semi-automated Business as Usual (BAU) strategy of reviewing suspicious transactions by Operations teams. Several concepts within the Framework assist in building the FDPA features, employing innovative techniques to intelligently control fraud risk. Organizations need to implement specialized data science units that work closely with these application teams for generating analytics-based decision-making models and tools.

4.1. Fraud Detection and Prevention Models

Traditionally, fraud detection relied on past records and indigenous business logic encapsulated in users' rules. As a consequence, these systems can manage recognition of previously identified fraud patterns, but can neither detect new, adapted schemes nor prevent fraud attempts. Machine learning models can advance this framework by learning from fraud samples and distinguishing them from normal behavior. Deep learning algorithms enable detection of the weakest points in both of those logical flows, opening space for hybrid fraud prediction-detection-prevention designs.

Based on the most suitable techniques, prediction and detection models can be built to protect a banking institution against internal and external fraud. The outcome of prediction models should be a validation of an operation before it reaches the backend, thus preventing fraud with intelligence built over many attempts. The detection model is a second level of protection, which investigates potential issues concentrated in certain areas and in operations with no previously implemented clickstream. Several solutions are available so the choice depends both on the requirement of applicable results in close to real-time (i.e., when a transaction is performed) and on the data volume. The prediction models could be based on lightweight models built on variables (numerical, categorical, and textual) with specific business rules from the organization. Risk scoring data preparation is capital for a successful model implementation; the risk score added to the operation is the guidance to finance investigation teams.

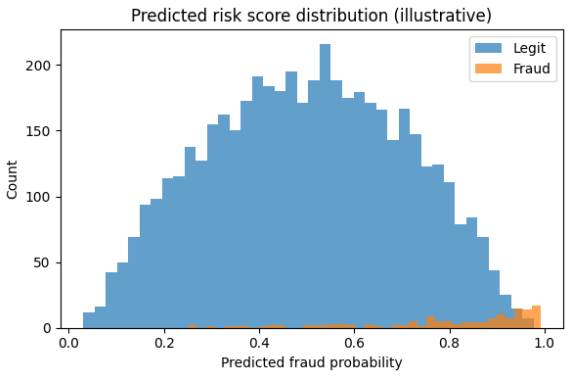

5. Ecosystem and Stakeholders

The proposed Risk Intelligence and Fraud Control Cloud-Native Big Data AI/ML Framework aligns and integrates the banking, digital commerce, insurance and telecommunications sectors, incorporating current industry

ISSN: 3067-283X
AMERICAN JOURNAL OF ANALYTICS AND ARTIFICIAL INTELLIGENCE
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 19 REVISED: MAY 09 ACCEPTED: MAY 24 PUBLISHED: JUNE 11

focuses such as Banking 4.0, digital banking transformation and customer centricity, thereby creating an Industry 4.0 risk ecosystem for Banking-as-a-Service, Payment-as-a-Service and Fraud-as-a-Service functions. Risk Intelligence and Fraud Control become relevant in an environment of highly accelerated transactional growth.

The Banking-as-a-Service paradigm is centred on customers and ecosystems, allowing banks to consider third-party products alongside their own, deepening customer relations. Banking-as-a-Service also refers to banks that facilitate transactions for non-bank corporations. The fast-moving universe of telecommunications revisits the fraud landscape for major risks. Fraud detection and prevention techniques are paramount for all provisioning players. Government, central banks and regulators also function within a risk and fraud control ecosystem. With the installation of 5G, the banking ecosystem extends into digital commerce, into payment, financial and insurance transactions, along 5G-based value-added services extending to public administration.

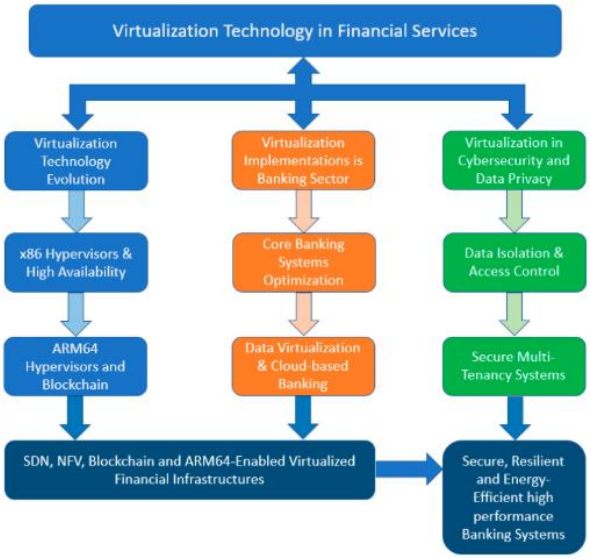

Fig 3: Ecosystem and Stakeholders of Cloud-Native Big Data

5.1. Banking Sector Applications

Risk intelligence and fraud control paradigms developed in this research are designed and adapted for a broad range of application domains. The banking sector is specifically chosen for detailed exploration of potential applications, reflecting a growing societal recognition of the banks' systemic importance that continues to intensify since the 2007–2008 financial crisis. Cyberattacks against payment service providers, high-profile data leaks, and continued regulatory enforcement actions under the Bank Secrecy Act worked in concert to sharpen the focus on risk management. New physical, credibility, market, operational, liquidity, and interest rate risks have subsequently emerged with the growth of crypto-assets, while risks also continue to be driven by the interconnectivity of the banking system and the increased interconnectedness of global financial markets. Within the banking sector, significant application opportunities also exist in credit card fraud detection, insurance claim fraud detection, insider fraud detection, remittance fraud detection, and know-your-customer violations, among others.

Credibility fraud, one prominent application area, is defined as the abuse of a legitimate relationship for illicit gain. It can be perpetrated by anyone from customers and suppliers to intermediaries and employees; perpetrators include customers working in tandem with bank employees or with accomplices posing as merchants. In banks and insurance companies, anomalies can arise for reasons outside the fraud domain; hence, the focus should be on correlating suspicious events and refining the screening list using heuristics to conceal less-probable transactions. To illustrate, A receives a call from a credit card customer about a potential fraudulent transaction on account B. Although this customer has no risk profile, the geographical location of account B and the fact that the account has a history of infractions from other customers with similar profiles indicate credibility fraud.

Equation 2: Training the classifier (maximum likelihood → cross-entropy loss)

Step 1: Bernoulli likelihood

For one example:

$$P(y_i \mid x_i) = p_i^{y_i}(1-p_i)^{(1-y_i)}$$

For dataset (independent samples):

$$\mathcal{L}(w,b) = \prod_{i=1}^{n} p_i^{y_i}\,(1-p_i)^{(1-y_i)}$$

Step 2: Log-likelihood

$$\log \mathcal{L} = \sum_{i=1}^{n} [y_i \log(p_i) + (1-y_i)\log(1-p_i)]$$

Step 3: Negative log-likelihood (loss to minimize)

$$J(w,b) = -\log \mathcal{L}$$

So:

$$J(w,b) = -\sum_{i=1}^{n} [y_i \log(p_i) + (1-y_i)\log(1-p_i)]$$

This is the **binary cross-entropy** used heavily in fraud detection.

Step 4: (Optional) Class-imbalance weighting

Fraud is rare; weight positives higher:

$$J(w,b) = -\sum_{i=1}^{n} [\alpha\, y_i \log(p_i) + \beta\,(1-y_i)\log(1-p_i)]$$

Where $\alpha > \beta$(common in fraud).

6. Cloud-Native Deployment and Operations

Cloud-native deployment and operations provide a set of modern approaches geared toward achieving efficiency in the design, development, testing, and management of applications deployed in heterogeneous multi-cloud environments. They enable industrialization of big data and AI/ML-based application creation. Consequently, the data processing, analytics, AI/ML processes powering the risk intelligence and fraud control paradigms are delivered as a set of container-based microservices deployed on Kubernetes platforms in a cloud-native fashion. These principles are applied to the implementation of the cloud-native banking and insurance solutions, delivering these services as microservices packaged for deployments on Kubernetes environments. Adoption of containerization of the applications also enables continuous integration, continuous delivery, monitoring, and management of the application microservices over the entire lifecycle.

In addition to application containers, Kubernetes platform and microservices, the overall architecture employs other supporting cloud-native components including message queues, distributed file storage, as well as source code repository, build pipeline, testing environments, monitoring services, and database services such as Redis and PostgreSQL. These components also leverage cloud-native principles such as automation of the operational operations, provisioning of resources in an on-demand manner, and ability to scale up and down elastically based on the workload to deliver increased efficiency and fault-tolerance besides reducing overall operational overhead.

6.1. Microservices and Containerization

Microservices and Containerization—The design and implementation of cloud-native systems for production should follow the principles outlined in the architecture. The technology components and supporting infrastructure should be separately provisioned and operationally managed. Each individual microservice within the framework can be

deployed using its most suitable run-time environment tool kit and underlying stack, independent of other microservices. This independent microservice implementation and run-time approach permits teams to adopt basic principles of agile software development and accelerate the product development cycle while enhancing the quality of deliverables.

Rapid scaling of cloud infrastructure should support burst provisioning, which enables immediate provisioning of additional instances of the service when the demand exceeds the threshold. In a highly parallel environment, such burst provisioning can happen without any business impact using cloud-native burst infrastructure provisioning and auto-scaling principles. Orchestrated data and service provisioning using containers and its orchestration tools further enhance the overall efficiency of operations.

Due to the complex and diverse nature of big data, any failure in the data processing pipeline will have a significant impact on the overall data quality and availability of data for analytics. Establishing a monitoring mechanism for each data processing stage will enable visibility and improve the quality of incoming data for analytics. Service orchestration to provide a data pipeline as a service to data scientists should be considered for the subsequent phases of product development once the initial data pipeline services are established and matured.

Equation 3: Confusion matrix + key metrics (what operations teams actually track)

Define counts:

- TP: predicted fraud and truly fraud
- FP: predicted fraud but truly legit
- TN: predicted legit and truly legit
- FN: predicted legit but truly fraud

Matrix:

$$\begin{pmatrix} TN & FP \\ FN & TP \end{pmatrix}$$

Metrics:

$$\text{Precision} = \frac{TP}{TP + FP} \quad \text{Recall (TPR)} = \frac{TP}{TP + FN}$$

$$\text{FPR} = \frac{FP}{FP + TN} \quad \text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

7. Data Management and Quality Assurance

Data management and quality assurance are crucial for developing robust risk intelligence and fraud detection systems, as they rely on large volumes of data from diverse sources, including historical data repositories and real-time transactional data feeds. Data management encompasses the processes of data acquisition, cleansing, transformation, and enrichment, while quality assurance ensures that incoming data meet required standards.

Data acquisition focuses on establishing connections with popular databases, data warehouses, and cloud data lakes using autonomous agents. Pre-existing databases provide historical records for model training, while current data flows ingest real-time transactional information. A supervised learning approach is used, training data consist of fifty-one relevant features extracted from historical fraud records for a two-year period. Data cleansing ensures the integrity of incoming data streams; integration, conversion, and transformation of external data sources into internal formats and standardization into the knowledge graph provide enrichment. A third-party cloud platform offers both tooling and capabilities to guarantee the success of these requirements.

ISSN: 3067-283X
AMERICAN JOURNAL OF ANALYTICS AND ARTIFICIAL INTELLIGENCE
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 19 REVISED: MAY 09 ACCEPTED: MAY 24 PUBLISHED: JUNE 11

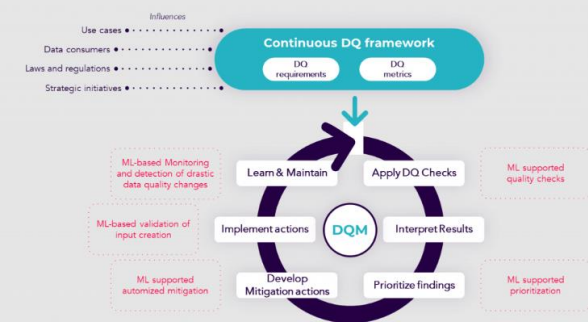

Fig 4: Data Quality Management

7.1. Data Acquisition, Cleansing, and Enrichment

Continuous collection of banking transaction data from a streaming source completes the ecosystem's supporting infrastructure. The mining of novel patterns within this transaction dataset supports algorithms for fraud detection and prevention. New transaction records pass through a set of rules to support data cleansing and quality improvement. The continuous nature of data arrival allows the implementation of a Lambda architecture, encompassing the simultaneous creation of batch and incremental datasets. These datasets feed the operation of two separate models intended for fraud detection and fraud prevention. Both models generate results that return to the source backend for further operationalization. The careful design of the data stream allows excluded transactions to re-enter the process on a different data pathway along the banking transaction lifecycle.

The multi-data-source configuration of the Pipeline Ecosystem encourages the resume of a continuous process to feed the supporting risk-intelligence and fraud-control pipelines for each detection and prevention domain. Batch processing integrates new data records hosted on the enterprise data warehouse before returning to a Hadoop-based operational staging area. A secondary data-cleansing routine removes stale data from the Hadoop wallet in preparation for replay episodes that explore episodic fraud detection and prevention.

8. Conclusion

By having analysed the approached ecosystems, it can be concluded that the hybrid design and architecture principles used for the framework should target three risk control core modules. These modules are deployed and operated in the core component of the framework whose microservices together with other core microservices provided by the framework's ecosystem and utilized in the banking sector and its insurance modules form the ecosystem. The ecosystem aims to provide innovative solutions for risk control and management in banking and insurance systems and utilize related solutions offered by the China Banking industry, ICAI group and other actors.

The defence and prevention models support the detection and prevention of attacks in banking systems through utilizing shared data sources that enhance and complement the fraud detection solutions. With a focus on Fraud Intelligence and Prevention the banking system can insulate itself from damage caused by fraud and attacks by other systems or actors. In the case of insurance investigations the focus is on providing evidence for forensic investigation for subsequent procedures of law enforcement bodies or watchdog organizations.

9. References

[1] Aven, T. (2016). Risk assessment and risk management: Review of recent advances on their foundation. European Journal of Operational Research, 253(1), 1–13.

[2] Baranwal, G., & Vidyarthi, D. P. (2019). A survey on cloud resource management. ACM Computing Surveys, 51(5), 1–36.

[3] Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5–32.

[3] Carcillo, F., Dal Pozzolo, A., Le Borgne, Y. A., Caelen, O., Mazzer, Y., & Bontempi, G. (2019). Combining unsupervised and supervised learning in credit card fraud detection. Information Sciences, 557, 317–331.

[4] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 785–794).

[5] Claessens, S., & Kose, M. A. (2018). Frontiers of macrofinancial linkages. BIS Papers, 95, 1–27.

[6] Dang, V., & Nguyen, M. (2020). A comprehensive survey of deep learning in fraud detection. IEEE Access, 8, 1–22.

[7] Dautov, R., & Distefano, S. (2018). Quantifying the performance of data stream processing systems: A systematic survey. IEEE Transactions on Big Data, 4(4), 1–18.

[8] Davis, M., & Dutta, K. (2020). Big data and analytics in banking: A review and opportunities. Journal of Management Analytics, 7(2), 1–24.

[9] Demertzis, K., Iliadis, L., Anezakis, V.-D., & Malli, E. (2020). Cybersecurity in banking with AI-driven fraud detection. Future Internet, 12(11), 1–23.

[10] Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of NAACL-HLT (pp. 4171–4186).

[11] European Banking Authority. (2021). EBA guidelines on ICT and security risk management. European Banking Authority.

[12] European Central Bank. (2020). Guide on outsourcing arrangements. European Central Bank.

[13] Fawcett, T. (2006). An introduction to ROC analysis. Pattern Recognition Letters, 27(8), 861–874.

[14] Finance Stability Board. (2020). Effective practices for cyber incident response and recovery. Financial Stability Board.

[15] Gai, K., Qiu, M., & Zhao, H. (2018). Cost-aware multimedia data hosting in cloud with big data processing. IEEE Transactions on Multimedia, 20(6), 1–14.

[16] Gartner. (2023). Hype cycle for fraud detection and prevention. Gartner Research.

[17] Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press.

[18] Hastie, T., Tibshirani, R., & Friedman, J. (2009). The elements of statistical learning: Data mining, inference, and prediction (2nd ed.). Springer.

[19] He, H., & Garcia, E. A. (2009). Learning from imbalanced data. IEEE Transactions on

Knowledge and Data Engineering, 21(9), 1263–1284.

[20] Hofmann, P., & Woods, D. (2010). Cloud computing: The limits of public clouds for business applications. IEEE Internet Computing, 14(6), 90–93.

[21] International Organization for Standardization. (2022). ISO/IEC 27001:2022 Information security, cybersecurity and privacy protection. ISO.

[22] Jayasekara, S., et al. (2020). Stream processing in the cloud: A survey. ACM Computing Surveys, 53(3), 1–36.

[23] Jurgovsky, J., Granitzer, M., Ziegler, K., Calabretto, S., Portier, P.-E., He-Guelton, L., & Caelen, O. (2018). Sequence classification for credit-card fraud detection. Expert Systems with Applications, 100, 234–245.

[24] Kelleher, J. D., Mac Namee, B., & D'Arcy, A. (2020). Fundamentals of machine learning for predictive data analytics (2nd ed.). MIT Press.

[25] Kiran, B. R., Thomas, D. M., & Parakkal, R. (2018). An overview of deep learning based methods for unsupervised and semi-supervised anomaly detection. Journal of Imaging, 4(2), 1–27.

[26] Kose, J., & Reddy, P. K. (2021). Big data analytics in financial risk management: A systematic review. Information Systems Frontiers, 23(4), 1–19.

[27] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. Advances in Neural Information Processing Systems, 25, 1097–1105.

[28] Kubernetes Authors. (2024). Kubernetes: Production-grade container orchestration documentation. Cloud Native Computing Foundation.

[29] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 4765–4774.

[30] Markowitz, H. (1952). Portfolio selection. The Journal of Finance, 7(1), 77–91.

[31] Microsoft. (2024). Azure Well-Architected Framework: Reliability, security, and cost optimization guidance. Microsoft Learn.

[32] NIST. (2018). Framework for improving critical infrastructure cybersecurity (Version 1.1). National Institute of Standards and Technology.

[33] NIST. (2020). Security and privacy controls for information systems and organizations (SP 800-53 Rev. 5). National Institute of Standards and Technology.

[34] Obermeyer, Z., & Emanuel, E. J. (2016). Predicting the future—Big data, machine learning, and clinical medicine. The New England Journal of Medicine, 375(13), 1216–1219.

[35] Pozzolo, A. D., Caelen, O., Le Borgne, Y.-A., Waterschoot, S., & Bontempi, G. (2015). Learned lessons in credit card fraud detection from

a practitioner perspective. Expert Systems with Applications, 41(10), 4915–4928.

[36] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1135–1144).

[37] Salinas, D., Flunkert, V., Gasthaus, J., & Januschowski, T. (2020). DeepAR: Probabilistic forecasting with autoregressive recurrent networks. International Journal of Forecasting, 36(3), 1181–1191.

[38] Van Vlasselaer, V., Eliassi-Rad, T., Akoglu, L., Snoeck, M., & Baesens, B. (2017). Gotcha! Network-based fraud detection for social security fraud. Management Science, 63(9), 3090–3110.

[39] Zohuri, B., & Moghaddam, M. (2020). Artificial intelligence driven by big data analytics: A deep learning approach. Springer.