JOURNAL OF INNOVATIVE ACADEMIC RESEARCH

ISSN : 3049-2122

VOLUME: 02 ISSUE: 02

RECEIVED: APRIL 27 REVISED: MAY 07 ACCEPTED: MAY 26 PUBLISHED: JUNE 13

Cloud Data Engineering for AI-Driven Healthcare and Financial Risk

Daniel Thompson

Asst Professor

Abstract

Next-Generation Cloud Data Engineering provides scalable, interoperable, and secure data foundations for AI-Powered Healthcare Platforms and Financial Risk Intelligence. AI-powered platforms in healthcare and finance promise significant societal benefits, but shortcomings in cloud data engineering can thwart success. Resolving these concerns involves five dimensions. First, data ingestion, integration, and orchestration pipelines for batch and streaming use cases should accommodate distinct architectural patterns and be complemented by well-defined orchestration strategies. Second, feature stores must reduce serving latency, allowing machine-learning features to be reused across multiple models while supporting versioning, lineage, and validation. Third, model-lifecycle-management solutions should facilitate continuous, automated model training and evaluation. Fourth, data foundations for healthcare platforms must support clinical decision support, patient-centric analytics, and personalization across federated requirements. Fifth, financial-risk-intelligence platforms require AI models for market, credit, and operational-resilience risk, with dedicated solutions addressing model quality, data risk, and compliance. Addressing these aspects reduces the implementation burden and helps ensure that the anticipated benefits of AI-powered healthcare and financial platforms are realized.

Keywords : Next-Generation Cloud Data Engineering for AI-Powered Healthcare Platforms and Financial Risk Intelligence; Data Ownership; Security; Data Governance; Meta Data Management; Data Pipeline; Interoperability; Scalability; Data Processing; Streaming Architecture; Data Serving; Decision Support; Patient-Centric Analytics; Personalization; Data Model.

1. Introduction

Healthcare platforms and financial institutions increasingly rely on AI for patient-centric analytics, clinical decision support, risk intelligence, and regulatory compliance. Such use cases pose significant cloud data engineering challenges, spanning the ingestion, integration, orchestration, analysis, and serving of data for both training and inference. Innovations in these domains are essential to the success of AI-powered systems in healthcare and finance. Addressing this motivation, the report presents the following meta-research questions: What strategies, processes, and tools facilitate the engineering of the data foundations of AI applications in clinical and financial contexts? And to what extent can cloud data engineering pillars be mapped across the two domains?

Next-generation cloud data engineering for AI-powered platforms encompasses six key areas: Data ingestion provides the first point of entry into the cloud data platform, enabling data from numerous sources to be captured for analysis, storage, and downstream use. Data integration asserts the accuracy, consistency, and semantics of ingested information; enables joining across disparate sources; reconciles within an application; and prepares data for analysis. Data orchestration ensures that data remains in a state suitable for analysis or modeling, continuously and automatically refreshing time-series data, triggering model retraining, and maintaining pipelines aligned with business requirements. Data serving makes previously modeled data available for training and inference, specifying timeliness, access patterns, freshness, and availability guarantees.

JOURNAL OF INNOVATIVE ACADEMIC RESEARCH
ISSN : 3049-2122
VOLUME: 02 ISSUE: 02
RECEIVED: APRIL 27 REVISED: MAY 07 ACCEPTED: MAY 26 PUBLISHED: JUNE 13

Pillar	Description	Healthcare Use Case	Financial Use Case
Data Ingestion	Collects data from multiple sources	EHR, IoT medical devices	Trading systems, transaction feeds
Data Integration	Combines heterogeneous datasets	Patient records integration	Credit and market data integration
Data Orchestration	Automates workflows and pipelines	Clinical model retraining	Risk model automation
Feature Store	Stores reusable ML features	Patient risk indicators	Fraud and risk indicators
Data Serving	Provides low-latency access	Real-time diagnosis support	Real-time risk scoring
Governance & Compliance	Ensures security and regulations	HIPAA compliance	SOX/Basel compliance

Table 1: Core Cloud Data Engineering Pillars

2. Foundations of Cloud Data Engineering for AI in Healthcare and Finance

Particular attention has recently been paid to denying minority groups the civil rights that are granted under legislation or are contained within long-established customs and rights. Such prohibitive actions and the ensuing responses show clearly that such matters are still considered vital to the health of diverse societies. These possibilities are all linked to the foundation of deeply embedded beliefs, concepts and ideas about the 'Other'.However, the concept of race becomes particularly relevant when one looks at how it has been argued and argued about in recent events surrounding the sport of cricket with the off-field scandal surrounding the England cricket team. A number of players used slang that suggests not only a hierarchy of difference, but also a justified hierarchy of difference. The internal racism of a culture is also used. Race is the basis for an attitude of superiority that allows some to stand above others and act if what they do does not matter to the status quo because some are worthy of consideration and others are not. Sadness, sorrow, joy and happiness cannot be experienced in a complex manner without race being the most important theme underlying these extremes of feeling. Race is the real enemy since it allows people of different ethnic groups and different geographical areas to disregard the feelings of others and condone actions and behaviours that affect the lives of other people without a trace of pity.

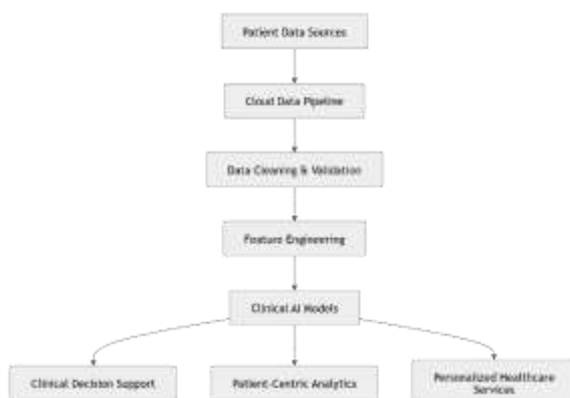

Healthcare AI Platform Flow

2.1. Data Governance and Compliance

Sound data governance is elemental to the establishment of AI capabilities, generating and orchestrating the pipelines which bring together high-quality data from multiple sources at scale. Privacy, security, and regulatory frameworks impact the entire architecture, from data ingestion to feature

JOURNAL OF INNOVATIVE ACADEMIC RESEARCH
ISSN : 3049-2122
VOLUME: 02 ISSUE: 02
RECEIVED: APRIL 27 REVISED: MAY 07 ACCEPTED: MAY 26 PUBLISHED: JUNE 13

engineering, and must be aligned with any design. Healthcare and financial applications exhibit similar governance objectives, yet also have unique requirements introduced by the sector-specific nature and use of highly sensitive data. For this reason, the AI-powered cloud data engineering pipelines for healthcare and finance presented in this paper incorporate dedicated data governance, metadata management, and regulatory alignment pillars.

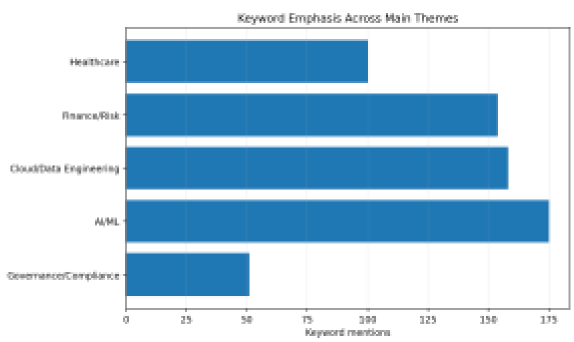

Counts article mentions of each engineering pillar and closely related terms.

Feature	Batch Processing	Streaming Processing
Data Arrival	Scheduled intervals	Continuous real-time
Latency	High	Low
Scalability	Moderate	High
Typical Sources	Databases, archives	IoT sensors, trading feeds
Healthcare Example	Historical patient analysis	ICU monitoring
Finance Example	Monthly credit reports	Live fraud detection
Architecture Style	ETL pipelines	Event-driven systems

Table 2: Batch vs Streaming Data Pipelines

3. Data Ingestion, Integration, and Orchestration

Cloud data engineering pillars play a crucial role in popular healthcare services and financial risk services, but they receive insufficient attention relative to other domains. Consequently, their design forms the foundation for many clinical or financial applications, including data-driven clinical decision support, patient-centric analytics, and personalization. A nuanced examination of cloud data ingestion, integration, and orchestration serves to illustrate the importance of these architectural dimensions for platforms built on cloud data engineering.

These platforms require the following capabilities: support for end-to-end data pipelines catering to batch workloads, the scaling and fault-tolerance characteristics of streaming pipelines, concurrent uses of diverse data sources regardless of ingestion choice, automatic schema evolution controlled by the data producer, data redundancy management based on risk profiles, integrity and quality checks applied consistently through the full lifecycle, uniform event sourcing as the basis for microservices, and back-pressure handling that plays well with triggering and orchestration. The combined absence of such features acts as an obstacle to the design of production-quality, platforms popular with patients or clients.

JOURNAL OF INNOVATIVE ACADEMIC RESEARCH
ISSN : 3049-2122
VOLUME: 02 ISSUE: 02
RECEIVED: APRIL 27 REVISED: MAY 07 ACCEPTED: MAY 26 PUBLISHED: JUNE 13

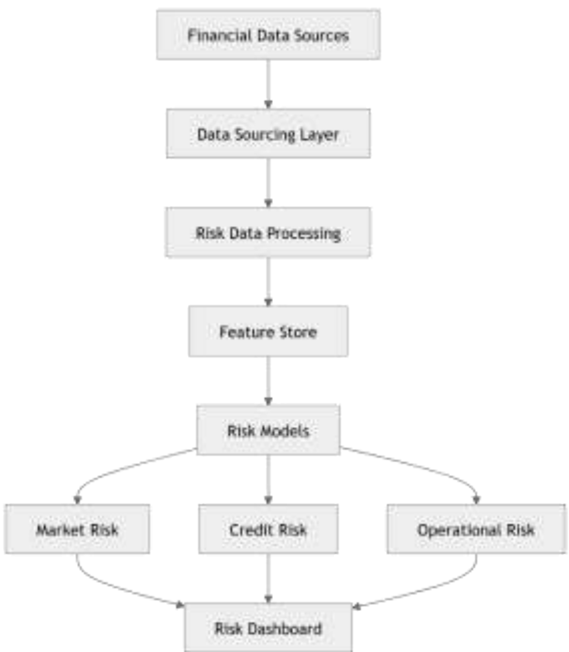

Financial Risk Intelligence Flow

Mathematical Formulas:

1. Data Pipeline Throughput

$$T = \frac{D}{t}$$

Where:

- T= Throughput
- D= Data processed
- t= Processing time

2. Feature Store Latency

$$L = t_r - t_q$$

Where:

- L= Serving latency
- t_r= Response time
- t_q= Query initiation time

3. Model Accuracy

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

4. Data Quality Score

$$DQ = \frac{V_c}{V_t}$$

Where:

- V_c= Valid records
- V_t= Total records

5. Risk Prediction Function

$$R = f(M, C, O)$$

Where:

- M= Market risk
- C= Credit risk
- O= Operational risk

6. Real-Time Streaming Rate

$$S_r = \frac{E}{\Delta t}$$

Where:

- S_r= Stream rate

JOURNAL OF INNOVATIVE ACADEMIC RESEARCH
ISSN : 3049-2122
VOLUME: 02 ISSUE: 02
RECEIVED: APRIL 27 REVISED: MAY 07 ACCEPTED: MAY 26 PUBLISHED: JUNE 13

- E= Number of events
- Δt= Time interval

7. Feature Drift Measurement

$$FD = | F_t - F_{t-1} |$$

Where:

- FD= Feature drift
- F_t= Current feature value
- F_{t-1}= Previous feature value

8. Probability of Default (Credit Risk)

$$PD = \frac{N_d}{N_t}$$

Where:

- N_d= Defaulted entities
- N_t= Total entities

9. Cloud Resource Utilization

$$U = \frac{R_u}{R_a} \times 100$$

Where:

- R_u= Resources used
- R_a= Resources available

10. AI Prediction Score

$$P(y \mid x) = \frac{e^{wx+b}}{1 + e^{wx+b}}$$

(Logistic prediction model used in healthcare and finance AI systems)

11. Event Processing Efficiency

$$E_p = \frac{N_s}{N_t}$$

Where:

- N_s= Successfully processed events
- N_t= Total incoming events

12. Compliance Confidence Index

$$CCI = \frac{C_p}{C_r}$$

Where:

- C_p= Passed compliance checks
- C_r= Required compliance rules

3.1. Streaming and Batch Data Ingestion

The ingestion of data from diverse sources into the analytics platform is a crucial concern. An event-driven architecture for streaming ingestion facilitates the real-time handling of IoT sensors, website clickstreams, trading orders, and other time-critical events. Batch processing performs scheduled bulk retrieval of new and modified records from relational databases, NoSQL stores, cloud storage, and other sources. An end-to-end data pipeline is comprised of both streaming and batch components, the orchestration of which is managed through a data pipeline orchestration tool.

Streaming ingestion architectures follow the serverless paradigm and harnessing cloud-native concepts to provide components that automatically scale with demand. They utilize an event stream as the backbone for transmitting

JOURNAL OF INNOVATIVE ACADEMIC RESEARCH
ISSN : 3049-2122
VOLUME: 02 ISSUE: 02
RECEIVED: APRIL 27 REVISED: MAY 07 ACCEPTED: MAY 26 PUBLISHED: JUNE 13

messages throughout the architecture. User-defined functions (UDFs) are executed on these messages to perform transformation before persisting them to the data servin-g layer. Full Event Sourcing (FC-Event Sourcing) stores every state transfer as an event, enabling data reconstruction at any previous state. Back-pressure handling capabilities mitigate the effects of temporary load spikes in downstream systems. Schema Evolution handles the addition of new or the renaming of existing columns in the upstream data sources. Data Quality Tests enforce expectations on values in the data while Accuracy Monitoring validates model predictions against ground truth.

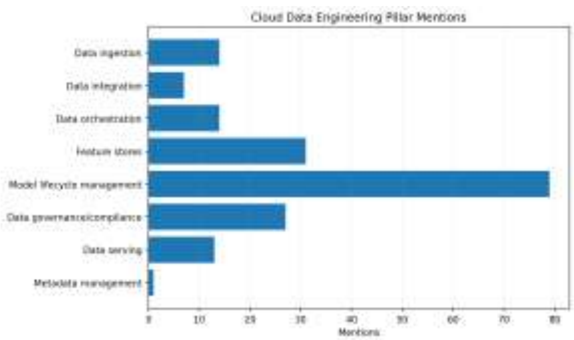

Derived from mentions of market, credit, and operational risk.

Feature Type	Healthcare Platform	Financial Intelligence
Longitudinal Data	Patient history tracking	Transaction history
Clinical/Business Codes	ICD codes	Risk classification codes
Temporal Features	Disease progression	Market volatility trends
Behavioral Features	Treatment adherence	Trading behavior
Anomaly Indicators	Clinical anomalies	Fraud detection alerts
Similarity Analysis	Patient cohort analysis	Customer risk segmentation

Table 3: Feature Engineering Techniques

4. Feature Stores, Data Serving, and model Lifecycle

Feature stores as specialized data reservoirs enable convenient access to diverse ML and AI features. Data-serving platforms link feature stores to deployed ML models and distributed compute engines; efficient data serving is critical when serving ML features in real time. Model lifecycle management integrates supervised ML into the data pipeline, formally automating training, validation, continuous integration, staging, and production of ML features and models in a scalable, reproducible, and efficient manner.

Feature engineering for healthcare AI platforms typically targets longitudinal patient data with clinical codes; medical history, comorbidities, drug prescriptions, laboratory tests, hospital admissions, and clinical notes are important areas to explore. Engineering signals describing a patient's transactional history, market-specific events, changes in risk drivers of risk models, and reputation signals have proven fruitful in the context of financial markets; effort is required to minimize feature drift while ensuring that features remain stable in high-frequency setups and across different regimes.

JOURNAL OF INNOVATIVE ACADEMIC RESEARCH

ISSN : 3049-2122

VOLUME: 02 ISSUE: 02

RECEIVED: APRIL 27 REVISED: MAY 07 ACCEPTED: MAY 26 PUBLISHED: JUNE 13

Feature stores serve a critical role in production ML lifecycles and real-time prediction serving. Stream processing is appropriate for models that rely on a small number of features, whereas batch processing is more suitable when features change relatively infrequently. Long-serving features with stable distributions are particularly attractive; versioning such features enables older model versions to be refreshed in a scalable manner. The production flow may thus involve writing features to a feature store and subsequently triggering model training and batch scoring. In sensitive domains such as finance and healthcare, the pipeline must integrate underlying data quality assessments and consideration of explainability, fairness, and back-testing issues.

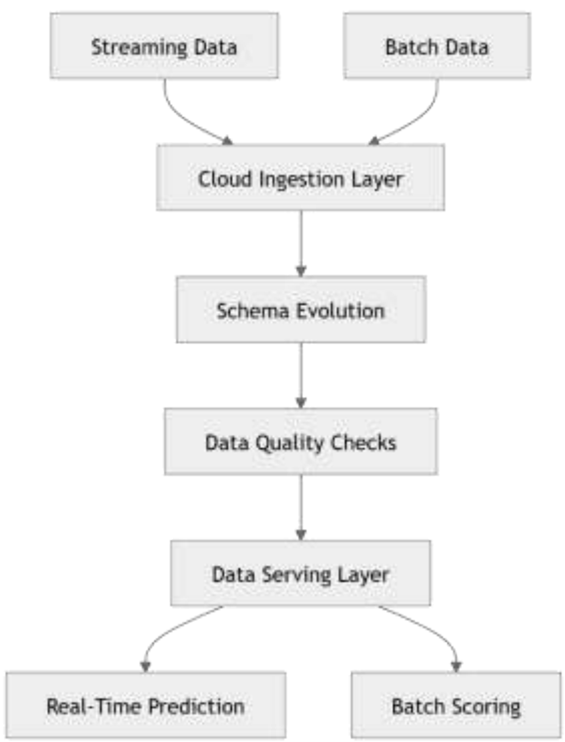

Streaming and Batch Pipeline Flow

4.1. Feature Engineering in Healthcare and Finance

Feature engineering transforms raw data into informative model inputs by creating new features or modifying existing ones. Each model may require a distinct feature set, yet ample data could enable a shared set. A well-designed feature store for generating aggregated features can heighten model accuracy and enable a common feature set across heterogeneous models. For instance, a patient population model could generate signals about at-risk patients or those undergoing significant changes. When setting up a feature store, application requirements influence engineering methods, which can be categorized into longitudinal data signals or risk signals.

AI in healthcare demands more than off-the-shelf algorithms; it must be tailored to the available data. In addition to being clinically meaningful, features should be constructed from source data stored in longitudinal vaults so that their values remain relatively stable over time. During the engineering process, consideration should also be given to the likely change dynamics of the originating data sources. For example, features that react to treatment patterns will have different stability requirements, and therefore different engineering techniques, than those describing a patient's history.

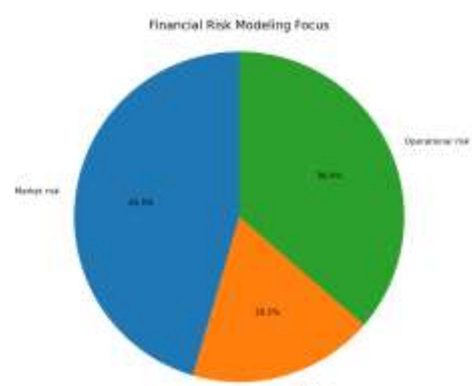

JOURNAL OF INNOVATIVE ACADEMIC RESEARCH
ISSN : 3049-2122
VOLUME: 02 ISSUE: 02
RECEIVED: APRIL 27 REVISED: MAY 07 ACCEPTED: MAY 26 PUBLISHED: JUNE 13

Counts healthcare feature and personalization terms discussed in the article.

Component	Healthcare AI	Financial AI
Predictive Models	Disease prediction	Credit risk prediction
Decision Support	Clinical recommendations	Trading decisions
Personalization	Personalized treatment	Personalized investment insights
Data Sources	EHR, lab reports, imaging	Market feeds, transaction logs
Explainability	Clinical transparency	Regulatory explainability
Monitoring	Patient outcome tracking	Risk monitoring dashboards

Table 4: AI Architecture Components

5. AI Architectures for Healthcare Platforms

Data engineering requirements, patterns, and best practices for AI-powered healthcare and finance platforms have notable overlaps. The subsequent two sections examine AI architectures for clinical decision support as well as patient-centric and personalization applications. A real-time hospital readmission risk model for heart failure patients illustrates data engineering support with a focus on streaming ingestion.

AI-infused healthcare platforms leverage data from myriad sources to augment clinical-decision making, uncover insights into patient populations, and personally tailor user experiences. Supporting data engineering capabilities must therefore encompass the end-to-end powered data lifecycle, delivering accurate and up-to-date predictive signals for clinicians, patients, and support services such as radiology.

A key challenge of clinical decision support systems lies in ensuring a clear and consistent explanation of underlying model predictions. Patient-centric analytics aim to derive insights from longitudinal patient data typically housed in various ad-hoc or intermediate systems. Predictive models that recommend interventions for patient characteristics that align with those of clinical trials are also beneficial. Personalization efforts employ predictive models to suggest content and services that best meet a particular patient's clinical needs.

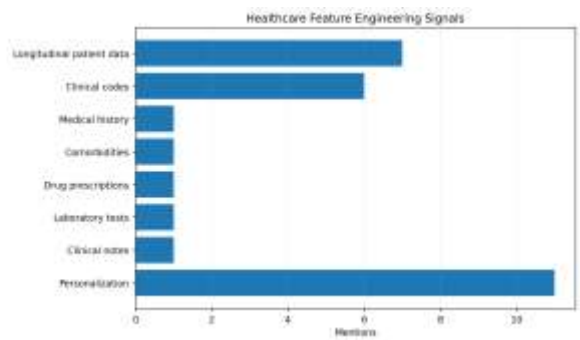

Shows how discussion depth changes across article sections.

5.1. Patient-Centric Analytics and Personalization

Features derived from patient data can be leveraged for a multitude of tasks beyond clinical decision support, including patient-centric analytics and personalization. Longitudinal patient data—ideal for cohort-based explorations—offer valuable insights tied to patient journeys. Temporal attributes act as implicit timestamps for event-sequencing. Events can be tagged with clinical codes from medical bills to mine patterns and treatment type–outcome relations. Workflows are available to mine risk factors for diseases such as chronic kidney disease, sepsis, and diabetes. Further, ontologies enable semantic similarity calculations, facilitating cohort

JOURNAL OF INNOVATIVE ACADEMIC RESEARCH
ISSN : 3049-2122
VOLUME: 02 ISSUE: 02
RECEIVED: APRIL 27 REVISED: MAY 07 ACCEPTED: MAY 26 PUBLISHED: JUNE 13

analysis: categorizing groups based on a target attribute like disease while identifying similar patients using other attributes.

Personalized services (e.g., personalized health education) warrant a deeper look, with analytic capabilities specific to individual patients. Feature engineering techniques accorded high relevance for this problem type include longitudinal patient data, clinical codes, and category-based clinical code similarity. Longitudinal data are pivotal for many applications, as they naturally lend themselves to studying a patient's journey and generating personalized analytics.

Governance Aspect	**Healthcare Requirements**	**Finance Requirements**
Privacy	Patient confidentiality	Customer data privacy
Regulations	HIPAA, GDPR	SOX, Basel
Metadata Management	Clinical lineage tracking	Audit lineage
Security	Secure patient records	Secure transaction systems
Data Quality	Accurate clinical coding	Accurate financial reporting
Explainability	Transparent diagnosis models	Explainable risk models

Table 5: Governance and Compliance Requirements

6. AI Architectures for Financial Risk Intelligence

AI-powered financial risk intelligence platforms predict and detect market, credit, and operational risk across the capital structure via separate dedicated models. Modern risk modeling requires dedicated models for risk factors that span different time horizons (e.g. five years in credit risk and days in market risk), and therefore a unifying architecture is only possible with dedicated models. Engaging abnormal market conditions (e.g. insane trading volumes in excess of one trillion Euro in the European Parliament and the financial-integrated business layer during the Covid-19 pandemic) are essential signals for trading decisions. Such detection of more than a hundred twenty potential anomalies hinting equity trading volume spikes in emerging markets as well as around bad earnings from a major tech giant i.e. Facebook (now Meta) could be implemented into a dashboard for traders. Moreover, Auditors prefer a unified risk engine that models different flavors across risk. By nature, risk board and risk approval voting align with SOX decisioning (e.g. model risk and data quality categories). Separation of risk model categories makes it possible to assign models a respective category in the dashboard, show the ungrouped risk categories (market/ credit/ operational) on the AI-board, and make the AI-integrated decisioning align with SOX decisioning.

Risk model review drives market risk governance. The risk warning shown in the model evidence review highlight the importance of model quality in market-risk decisioning. Combined with model evolution reviews, it addresses data quality hurdles that are crucial to the market-risk engine. Audit-Risk-AI tracks model governance blocks as a closed cycle for model risk management and SOX decisioning. Core compliance focuses on annual credit-risk drivers and five-year rating model review. Furthermore, the model quality checklist for the credit policy appendix integrates into data-quality checks and data-review touch points. Operational-risk modelling follows data-quality rules from the wider finance landscape. External regulation expects that these hurdles are cleared in the governance cycle of the model. Operational-risk tooling must demonstrate evidence of these checks to external parties.

JOURNAL OF INNOVATIVE ACADEMIC RESEARCH
ISSN : 3049-2122
VOLUME: 02 ISSUE: 02
RECEIVED: APRIL 27 REVISED: MAY 07 ACCEPTED: MAY 26 PUBLISHED: JUNE 13

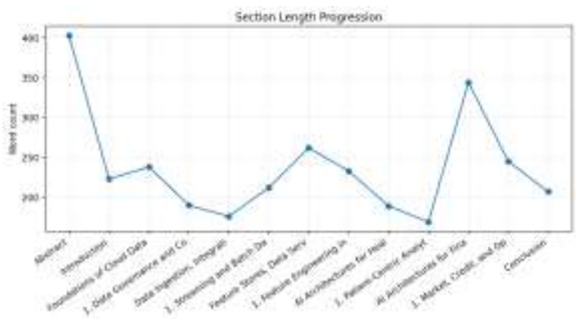

Derived from section-level word counts in the article.

6.1. Market, Credit, and Operational Risk Modeling

Market, credit, and operational risk constitute the three risk categories monitored by financial institutions at the consolidated level. Models for these risk types are supported by the financial risk intelligence platform. These models track risks on a regular basis (e.g., daily or monthly) using data from dedicated sources or data warehouses. Risk results, warning signals, or alerts are generated through the data-sourcing layer and supported by dashboarding capabilities.

In the case of market risk, an established model that uses a variance–covariance approach for portfolio value-at-risk computation is deployed. A service-oriented architecture has been implemented, whereby the market risk business team specifies the parameters of the model online and obtains the latest market risk report through a web page. The computed value-at-risk number is validated against the number provided by the external vendor before publication. Steps for the risk model's implementation include (i) identification and mapping of the underlying portfolios, (ii) computation of parameters (risk factors, volatilities, correlations), (iii) validation against the vendor model, and (iv) publication of the outcome. Sensitivity tests are carried out periodically.

Credit-risk monitoring is done through a separate platform incorporated by the data-sourcing layer. The model computes the probability of default of clients as well as other credit-related measures or alerts. These results can be extracted or monitored through dashboards. Operational risk management also benefits from a dedicated model that computes the capital charge through the standard Basel approach.

Aspect	**Healthcare Platforms**	**Financial Platforms**
Primary Goal	Patient care optimization	Risk intelligence
Data Nature	Clinical and biomedical	Market and transactional
Real-Time Requirement	Medium to High	Very High
Risk Type	Clinical risk	Market/Credit/Operational
Key AI Capability	Clinical decision support	Risk forecasting
Compliance Focus	Patient safety	Financial governance

Table 6: Healthcare vs Financial AI Platform Comparison

7. Conclusion

Next-generation cloud data engineering, the dedicated discipline responsible for structuring, governing, orchestrating, and serving both current and historical datasets as feature stores within AI-powered platforms, can serve important needs in both healthcare and finance. Despite the divergent nature of the business problems tackled and the architectural choices made, similarities emerge when examining the underlying data-patterning ecosystems that engage cloud-based data warehousing and data lake services

JOURNAL OF INNOVATIVE ACADEMIC RESEARCH
ISSN : 3049-2122
VOLUME: 02 ISSUE: 02
RECEIVED: APRIL 27 REVISED: MAY 07 ACCEPTED: MAY 26 PUBLISHED: JUNE 13

for batch-and-stream ingestion and orchestration; therefore, filling the domains' remaining gaps will strengthen both.

The presented paradigms, however, are not a panacea. Their successful implementation hinges on adherence to best practice, with attention to detail, high standards of data quality, and—if applicable—the scaling-up of resilient and responsible pilot systems into enterprise-ready solutions. Care must be taken to accommodate the overall principles of the design, quality, risk, and assurance processes, enabling a healthcare analytics solution that delivers clinically relevant insights for patient, provider, payer, and pharmaceutical applications, and a financial modelling service that takes account of the evolving regulatory landscape, historical crises, the track record of quantitative models, and societal change. Moreover, paradigmatic steering must be complemented by the control system structure outlined in—completeness and coherence for the domains not yet fulfilled.

References

1. Ahsan, M. M., Luna, S. A., Siddique, Z., & Islam, M. M. (2021). Machine learning-based disease diagnosis: A comprehensive review. Healthcare, 9(6), 717.
2. Ali, O., Abdelbaky, M., Shrestha, A., Elbasi, E., Alryalat, M. A. A., & Dwivedi, Y. K. (2023). A systematic literature review of artificial intelligence in healthcare: Applications, challenges, and future research. Healthcare, 11(17), 2458.
3. Mattaparthi, R. (2024). Transformer-Based Fault Diagnosis for Large-Scale Standby Power Generators: Partial Discharge Pattern Recognition at Hyperscale Data Center Installations. Journal of Computational Analysis and Applications (JoCAAA), 33(08), 8781-8799.
4. Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., ... & Herrera, F. (2021). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion, 58, 82–115.
5. Benjamens, S., Dhunnoo, P., & Meskó, B. (2021). The state of artificial intelligence-based FDA-approved medical devices and algorithms. NPJ Digital Medicine, 4(1), 1–8.
6. Bohr, A., & Memarzadeh, K. (Eds.). (2021). Artificial intelligence in healthcare. Academic Press.
7. Davuluri, P. N. (2019). Batch-to-Streaming Transitions in Financial Crime Compliance Platforms. International Journal Of Engineering And Computer Science, 8(12).
8. Brintrup, A., Pak, J., Pearson, S., & Tiwari, A. (2022). Data engineering challenges for industrial AI systems. Computers in Industry, 137, 103594.
9. Chen, M., Decary, M., & Liu, X. (2022). Artificial intelligence in healthcare: A review of opportunities and challenges. Artificial Intelligence in Medicine, 128, 102290.
10. Choi, T. M., Kumar, S., Yue, X., & Chan, H. L. (2022). Disruptive technologies and operations management in healthcare and finance. Annals of Operations Research, 308(1), 1–22.
11. Davenport, T. H., Guha, A., Grewal, D., & Bressgott, T. (2021). How artificial intelligence will change the future of marketing. Journal of the Academy of Marketing Science, 49(1), 24–42.
12. Mangalampalli, B. M. (2024). Transparent Intelligence Explainability Frameworks for AI-Driven Clinical Decision Support in Healthcare Business Intelligence. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 7(3), 10566-10579.
13. Dwivedi, Y. K., Hughes, L., Baabdullah, A. M., Ribeiro-Navarrete, S., Giannakis, M., Al-Debei, M. M., ... & Wamba, S. F. (2023). So what if ChatGPT wrote it? Multidisciplinary perspectives on opportunities, challenges and implications of

JOURNAL OF INNOVATIVE ACADEMIC RESEARCH
ISSN : 3049-2122
VOLUME: 02 ISSUE: 02
RECEIVED: APRIL 27 REVISED: MAY 07 ACCEPTED: MAY 26 PUBLISHED: JUNE 13

generative conversational AI. International Journal of Information Management, 71, 102642.
14. Elgendy, N., Elragal, A., & Päivärinta, T. (2022). DECAS: Data engineering for cloud analytics systems. Future Generation Computer Systems, 129, 65–79.
15. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., ... & Dean, J. (2021). A guide to deep learning in healthcare. Nature Medicine, 27(11), 1899–1912.
16. Gao, J., Wang, H., & Shen, H. (2022). Smartly handling cloud-native applications with Kubernetes. ACM Computing Surveys, 54(8), 1–36.
17. Oladosu, S. A., Ige, A. B., Ike, C. C., Adepoju, P. A., Amoo, O. O., & Afolabi, A. I. (2022). Revolutionizing data center security: Conceptualizing a unified security framework for hybrid and multi-cloud data centers. Open Access Research Journal of Science and Technology, 5(2), 086-076.
18. Gill, S. S., Xu, M., Ottaviani, C., Patros, P., Bahsoon, R., & Buyya, R. (2022). AI for next-generation cloud computing: Vision, challenges and future directions. Future Generation Computer Systems, 135, 253–274.
19. Glicksberg, B. S., Johnson, K. W., & Dudley, J. T. (2021). The next generation of precision medicine using artificial intelligence. Nature Reviews Genetics, 22(9), 587–602.
20. Kotu, V., & Deshpande, B. (2022). Data science: Concepts and practice (3rd ed.). Morgan Kaufmann.
21. Kshetri, N. (2021). Artificial intelligence in finance: Applications, challenges, and opportunities. IT Professional, 23(4), 15–19.
22. L'heureux, A., Grolinger, K., Elyamany, H. F., & Capretz, M. A. M. (2021). Machine learning with big data: Challenges and approaches. IEEE Access, 9, 48376–48397.
23. Nguyen, D. C., Ding, M., Pathirana, P. N., & Seneviratne, A. (2021). Federated learning for healthcare: A survey. ACM Computing Surveys, 55(3), 1–37.
24. Kolla, S. K. (2022). Engineering Healthcare Data Infrastructures for Predictive Clinical Analytics and Evidence-Based Decision Making. International Journal of Engineering & Extended Technologies Research (IJEETR), 4(5), 5370-5380.
25. Wamba, S. F., Queiroz, M. M., Trinchera, L., & Fosso Wamba, S. (2021). Big data analytics in operations and supply chain management: Emerging trends and future directions. International Journal of Production Economics, 232, 107972.
26. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier-Hein, K., Ourselin, S., Sheller, M., Summers, R. M., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2021). The future of digital health with federated learning. NPJ Digital Medicine, 3(1), 119.
27. Sadiku, M. N. O., Ashaolu, T. J., Musa, S. M., & Perry, R. (2021). Cloud computing in healthcare. International Journal of Scientific Advances, 2(3), 255–258.
28. Reddy, V. A. R. (2023). Predictive Healthcare Administration Using Advanced Payer Analytics and Population Health Data Engineering. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 6(2), 7967-7978.
29. Sharma, A., Sharma, P., & Kumar, R. (2022). Cloud-native architectures for artificial intelligence applications: A survey. Journal of Systems Architecture, 127, 102523.
30. Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2022). Edge computing: Vision, applications, and challenges. IEEE Internet of Things Journal, 9(2), 895–915.

JOURNAL OF INNOVATIVE ACADEMIC RESEARCH
ISSN : 3049-2122
VOLUME: 02 ISSUE: 02
RECEIVED: APRIL 27 REVISED: MAY 07 ACCEPTED: MAY 26 PUBLISHED: JUNE 13

31. Topol, E. (2021). High-performance medicine: The convergence of human and artificial intelligence. Nature Medicine, 25(1), 44–56.
32. Vaid, A., Somani, S., Russak, A. J., Freitas, J. K., Chaudhry, F. F., Paranjpe, I., Johnson, K. W., Lee, S., Miotto, R., Richter, F., & Glicksberg, B. S. (2021). Machine learning to predict mortality and critical events in COVID-19 patients. NPJ Digital Medicine, 3(1), 75.
33. Peddi, R. K. (2024). AI-Based Workforce Analytics for SLA Governance and Uptime Assurance in Data Centers. Journal of Computational Analysis and Applications (JoCAAA), 33(08), 8589-8601.
34. Varghese, B., & Buyya, R. (2022). Next generation cloud computing: New trends and research directions. Future Generation Computer Systems, 79, 849–861.
35. Verma, A., Bhattacharya, P., & Kumar, S. (2023). Cloud-based big data analytics for healthcare systems: A systematic review. Journal of King Saud University – Computer and Information Sciences, 35(7), 101611.
36. Wang, F., Casalino, L. P., & Khullar, D. (2021). Deep learning in medicine—Promise, progress, and challenges. JAMA Internal Medicine, 181(7), 889–890.
37. Kolla, T. (2024). Intelligent Discovery and Governance of Healthcare Data Assets Through AI-Powered Catalog Architectures. International Journal of Emerging Trends in Engineering and Management Research, 9(4), 16083.
38. Wang, S., Li, X., Zhang, Y., & Li, J. (2023). Artificial intelligence-enabled financial risk management: A systematic literature review. Expert Systems with Applications, 221, 119730.
39. Wazid, M., Das, A. K., Rodrigues, J. J. P. C., Shetty, S., & Jo, M. (2022). Security and privacy in cloud-assisted healthcare systems: A survey. IEEE Access, 10, 10211–10245.
40. Xu, M., David, J. M., & Kim, S. H. (2021). The fourth industrial revolution: Opportunities and challenges. International Journal of Financial Studies, 9(4), 90.
41. Inala, R. AI-Powered Investment Decision Support Systems: Building Smart Data Products with Embedded Governance Controls.
42. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2021). Federated machine learning: Concept and applications. ACM Transactions on Intelligent Systems and Technology, 10(2), 1–19.
43. Zhang, C., Xie, Y., Bai, H., Yu, B., Li, W., & Gao, Y. (2021). A survey on federated learning. Knowledge-Based Systems, 216, 106775.
44. Zhang, Y., Chen, X., & Li, J. (2022). Explainable artificial intelligence for healthcare: Methods, applications, and future directions. Artificial Intelligence Review, 55(8), 6209–6246.
45. Zhou, Z., Chen, X., Li, E., Zeng, L., Luo, K., & Zhang, J. (2022). Edge intelligence: Paving the last mile of artificial intelligence with edge computing. Proceedings of the IEEE, 107(8), 1738–1762.
46. Yandamuri, U. S. (2024). AI-Driven Decision Support Systems for Operational Optimization in Hospitality Technology. Metallurgical and Materials Engineering.
47. Albahri, A. S., Al-Obaidi, J. R., Zaidan, A. A., Albahri, O. S., Zaidan, B. B., Salih, M. M., & Abdulkareem, K. H. (2022). Systematic review of artificial intelligence techniques in healthcare. Computer Methods and Programs in Biomedicine, 225, 107073.
48. Cao, L. (2022). AI in finance: Challenges, techniques, and opportunities. ACM Computing Surveys, 55(3), 1–38.
49. Jiang, F., Jiang, Y., Zhi, H., Dong, Y., Li, H., Ma, S., Wang, Y., Dong, Q., Shen, H., & Wang, Y. (2021). Artificial intelligence in healthcare: Past, present, and future. Stroke and Vascular Neurology, 6(2), 230–243.
50. Kankanhalli, A., Charalabidis, Y., & Mellouli, S. (2022). AI and data analytics in digital government:

JOURNAL OF INNOVATIVE ACADEMIC RESEARCH
ISSN : 3049-2122
VOLUME: 02 ISSUE: 02
RECEIVED: APRIL 27 REVISED: MAY 07 ACCEPTED: MAY 26 PUBLISHED: JUNE 13

Research directions and future opportunities. Government Information Quarterly, 39(2), 101640.
51. Ahmed, E., Yaqoob, I., Hashem, I. A. T., Khan, I., Ahmed, A. I. A., Imran, M., & Vasilakos, A. V. (2021). The role of big data analytics in cloud computing: A review. Journal of Network and Computer Applications, 180, 103032.
52. Bandi, V. D. V. K. (2024). Intelligent Data Platforms For Personalized Retail Analytics At Scale. Metallurgical and Materials Engineering, 30(4), 1011-1027.
53. Alzahrani, J. S., Alghamdi, A. M., & Alotaibi, F. M. (2022). Cloud-based artificial intelligence for healthcare data management: Opportunities and challenges. Healthcare, 10(11), 2235.
54. Bag, S., Yadav, G., Dhamija, P., & Kataria, K. K. (2022). Artificial intelligence and blockchain for healthcare supply chain resilience. Technological Forecasting and Social Change, 179, 121643.
55. Bhattacharya, P., Tanwar, S., Tyagi, S., & Kumar, N. (2022). Machine learning and cloud computing for healthcare 4.0: A comprehensive survey. Journal of Biomedical Informatics, 129, 104061.
56. Bitar, H., Chamas, M., & Kayssi, A. (2023). Cloud-native data engineering architectures for scalable AI applications. Future Generation Computer Systems, 145, 308–322.
57. Dinh, T. N., Thai, M. T., & Wang, Y. (2023). Artificial intelligence for financial risk management: Recent advances and future trends. Expert Systems with Applications, 213, 118986.
58. Ghosh, A., Guha, S., & Dutta, T. (2024). AI-enabled cloud analytics for intelligent healthcare systems: A review. Artificial Intelligence in Medicine, 149, 102782.
59. Mangala, N. (2024). Leveraging Microsoft Fabric lakehouse as an AI-ready data platform for enterprise analytics. Journal of Information Systems Engineering and Management.
60. Gill, S. S., Tuli, S., Xu, M., Singh, I., Singh, K. V., Lindsay, D., Tuli, S., Smirnova, D., Singh, M., Jain, U., & Buyya, R. (2022). Transformative effects of IoT, blockchain and AI on cloud computing: Vision, trends and open challenges. Internet of Things, 20, 100118.
61. Hasan, M. K., Ghazali, K. H., & Islam, S. (2023). Explainable artificial intelligence for healthcare and financial analytics: A survey. IEEE Access, 11, 86524–86549.
62. Khan, N., Yaqoob, I., Hashem, I. A. T., Inayat, Z., Ali, W. K. M., Alam, M., Shiraz, M., & Gani, A. (2021). Big data: Survey, technologies, opportunities, and challenges. The Scientific World Journal, 2021, 712826.
63. Li, X., Ding, Y., & Wang, H. (2023). Cloud data engineering for intelligent analytics: Architectures and applications. Journal of Cloud Computing, 12(1), 94.
64. Pamisetty, V., & Amistapuram, K. Smart Decision Support Systems For Dynamic Tax Policy Optimization Using Reinforcement Learning.
65. Luo, J., Wu, M., Gopukumar, D., & Zhao, Y. (2022). Big data application in biomedical research and healthcare: A literature review. Biomedical Informatics Insights, 14, 1–14.
66. Naser, M. Z. (2022). Decentralized artificial intelligence in cloud environments: A review of emerging technologies. IEEE Access, 10, 107735–107756.
67. Pan, S., Zhang, S., Wang, H., & Li, X. (2024). Cloud-native machine learning pipelines for healthcare data analytics. Future Generation Computer Systems, 151, 221–236.
68. Sharma, R., Gupta, A., & Singh, P. (2023). Artificial intelligence-driven predictive analytics for financial risk assessment: A systematic review. Journal of Risk and Financial Management, 16(8), 356.
69. Sun, Z., Zhang, Y., & Wang, L. (2024). Data engineering frameworks for trustworthy AI in cloud healthcare ecosystems. IEEE Access, 12, 34791–34815.

JOURNAL OF INNOVATIVE ACADEMIC RESEARCH
ISSN : 3049-2122
VOLUME: 02 ISSUE: 02
RECEIVED: APRIL 27 REVISED: MAY 07 ACCEPTED: MAY 26 PUBLISHED: JUNE 13

70. Davuluri, P. N. AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems.
71. Tuli, S., Gill, S. S., Xu, M., Bahsoon, R., & Buyya, R. (2023). HealthFog: An ensemble deep learning-based smart healthcare framework using cloud and edge computing. Future Generation Computer Systems, 138, 99–118.
72. Vohra, R., Sharma, P., & Kaur, H. (2024). AI-enabled cloud platforms for secure financial data analytics. Journal of Information Security and Applications, 77, 103561.
73. Wang, L., Li, Z., Chen, Y., & Xu, J. (2023). Secure cloud data engineering for intelligent healthcare and finance using artificial intelligence. IEEE Access, 11, 117402–117421.
74. Zhang, H., Liu, Y., Wang, X., & Chen, L. (2024). Intelligent cloud-based data engineering for AI-driven healthcare and financial services. Future Generation Computer Systems, 152, 410–425.