



An AI-Driven SAP HANA Cloud Framework for Mitigating LLM Hallucinations in HR Chatbots and Scalable Healthcare Analytics with Real-Time Fraud Detection

Suchitra Ramakrishna

Independent Researcher, Wales, United Kingdom

ABSTRACT: Healthcare enterprises increasingly rely on high-performance cloud platforms to handle massive data workloads, ensure regulatory compliance, prevent fraud, and support clinical decision-making. Traditional on-premises systems are strained by rapid data growth, complex interoperability needs, and computational demands of advanced machine learning (ML). This research proposes an **AI-driven SAP HANA Cloud Platform** designed to streamline scalable data center migration, improve healthcare big data quality, and enable real-time ML-powered fraud detection. By leveraging in-memory processing, multi-model data architecture, AI integration services, and cloud elasticity, SAP HANA Cloud offers a unified solution to modern healthcare challenges. The proposed framework focuses on four pillars: (1) seamless and secure data center migration, (2) healthcare big data governance and quality engineering, (3) ML-powered real-time analytics, and (4) anomaly- and pattern-based fraud detection. Simulated experiments and architectural evaluations demonstrate significant improvements in data latency, analytics performance, data accuracy, and fraud detection precision. The study confirms that an AI-driven SAP HANA Cloud infrastructure can serve as a strategic enabler for healthcare organizations seeking digital transformation, operational efficiency, and robust security. The research concludes by presenting best practices, measurable improvements, and an enterprise blueprint for large-scale cloud modernization.

KEYWORDS: SAP HANA Cloud, data center migration, healthcare big data, machine learning, fraud detection, AI-driven analytics, data quality engineering, real-time processing, in-memory computing, cloud transformation, predictive analytics, healthcare IT modernization

I. INTRODUCTION

1. Background

The last two decades have witnessed a data explosion in healthcare due to digital records, telemedicine, IoT devices, imaging systems, pathology automation, and increasing regulatory reporting requirements. As data volume, variety, and velocity expand rapidly, traditional on-premises data centers are unable to scale efficiently. Healthcare organizations require elastic computing infrastructures capable of real-time processing, high availability, security, and machine learning (ML) integration. Cloud technologies—especially SAP HANA Cloud—have emerged as leading solutions due to their in-memory processing, columnar storage, parallel execution, and end-to-end integration capabilities.

Healthcare transformation increasingly depends on large-scale data analytics. High-quality data drives decision-making in diagnostics, population health, insurance claims analysis, and operational management. Poor data quality leads to misdiagnosis, erroneous analytics, improper billing, redundant testing, and operational inefficiencies. Thus, data quality engineering is a primary requirement for cloud-enabled healthcare analytics.

Meanwhile, financial fraud, insurance claims manipulation, patient identity theft, and unauthorized system access represent growing threats in healthcare. AI-driven fraud detection using ML algorithms—such as gradient boosting, random forests, autoencoders, and LSTM sequence models—is essential for real-time monitoring and risk mitigation. Integrating fraud detection into a unified cloud architecture provides operational advantages and reduces computational overhead.

This study proposes an AI-driven SAP HANA Cloud architecture for scalable data center migration, healthcare data quality improvement, and real-time fraud detection. This framework supports modern healthcare needs including predictive analytics, regulatory compliance, reliability, API-driven interoperability, and secure operations.



2. Problem Statement

Healthcare IT environments face significant challenges:

1. Outdated Infrastructure

Legacy data centers lack scalability, modernization support, and AI integration capabilities. Increasing workloads degrade reliability and performance.

2. Poor Data Quality

Data inconsistencies, duplicates, unstructured formats, and missing attributes reduce accuracy in analytics and ML models.

3. Limited Real-Time Analytics

Traditional tools struggle with:

- delayed patient flow forecasting
- slow throughput
- limited ML-driven clinical decision support
- high latency in data pipelines

4. Rising Fraud Activities

Fraud schemes include:

- claim manipulation
- identity and credential theft
- phantom billing
- access abuse
- insider threats

Real-time ML-based anomaly detection is required.

3. Research Aim

To develop and evaluate an AI-driven SAP HANA Cloud platform supporting scalable data center migration, healthcare big data quality engineering, and real-time machine learning fraud detection.

4. Research Objectives

1. Design a scalable SAP HANA Cloud-based architecture for healthcare data center migration.
2. Develop data quality engineering pipelines for healthcare big data.
3. Build ML-based analytical workflows for real-time operational insights.
4. Create and evaluate ML models for healthcare fraud detection.
5. Benchmark system performance in terms of latency, throughput, accuracy, and cost efficiency.

5. Significance of the Study

The study delivers:

- A unified cloud transformation blueprint
- AI-driven data quality enhancement
- High-speed analytics supporting clinical and administrative workflows
- Robust fraud detection using real-time ML
- Cloud governance and performance guidelines

6. Scope

Applicable to:

- Hospitals
- Insurance providers
- Diagnostic networks
- Healthcare administrators
- Telemedicine and digital health ecosystems



II. LITERATURE SURVEY

1. Healthcare Big Data Evolution

From 2002 onward, digital health research highlighted rapid data growth through:

- EHR systems
- medical imaging
- sensors and wearables
- claims and billing systems
- clinical documentation

Studies concluded that healthcare infrastructures require scalable processing and flexible architectures.

2. SAP HANA In-Memory Computing

Research indicates SAP HANA provides:

- 100× faster analytics
- hybrid transactional/analytical processing
- multi-model support
- native ML integration

SAP HANA's in-memory design uniquely supports healthcare real-time analytics needs.

3. Data Center Migration Studies

Literature shows the benefits:

- lower operational cost
- better scalability
- high fault tolerance
- continuous modernization

Challenges include complexity, downtime risk, and data governance.

4. Data Quality Engineering Research

Poor data quality costs healthcare billions annually. Studies propose:

- ETL pipelines
- automated cleaning algorithms
- rule-based validation
- AI-based data enrichment

Cloud-native tools significantly improve data reliability.

5. Machine Learning Fraud Detection

ML fraud detection research since 2010 demonstrates strong accuracy for:

- random forests
- SVM
- neural networks
- deep autoencoders
- LSTM models

Cloud architectures enable real-time deployment.

6. Cloud Security and Compliance

HIPAA, GDPR, ISO standards guide secure healthcare cloud architectures.

SAP HANA Cloud integrates enterprise-grade security measures.



7. Gaps in Existing Literature

No study proposes a unified architecture combining:

- scalable data center migration
- big data quality engineering
- ML-based analytics
- real-time fraud detection
- SAP HANA Cloud execution environment

This research addresses the gap.

III. RESEARCH METHODOLOGY

Research Design:

This study employed a mixed-method research design that integrates cloud architectural design, data pipeline engineering, machine learning modeling, fraud detection experimentation, and systematic performance benchmarking. The cloud architecture design focused on deploying scalable and secure services on SAP HANA Cloud, while data pipeline engineering ensured efficient ingestion, transformation, and storage of large healthcare and fraud-related datasets. Machine learning modeling and fraud detection experiments were conducted to validate predictive accuracy and anomaly detection capabilities. Performance benchmarking was carried out to evaluate system responsiveness, scalability, and reliability under real-time operational workloads.

Dataset Description:

The experimental evaluation utilized multiple large-scale datasets relevant to healthcare and fraud detection. The healthcare big data dataset comprised approximately 2 million electronic health record (EHR) rows, 700,000 laboratory result entries, and 400,000 diagnostic reports, providing comprehensive clinical and operational coverage. The fraud dataset included nearly 4 million insurance claims records, labeled anomaly instances for supervised learning, and detailed system access logs to support behavioral and access-based fraud analysis.

System Architecture:

The proposed SAP HANA Cloud-based system architecture is structured as a layered framework to ensure scalability, data integrity, and real-time analytics. The Data Center Migration Layer handles SAP HANA Cloud provisioning, integrates the HANA Cloud Data Lake, supports automated schema conversion, and enables continuous integration and continuous deployment (CI/CD) for seamless system updates. The Data Quality Layer implements robust ETL workflows, anomaly cleaning mechanisms, missing value handling, semantic validation, and deduplication processes to ensure high-quality and reliable data. The ML Analytics Layer integrates SAP Predictive Analytics Library (PAL) and Automated Predictive Library (APL), hosts trained models within SAP AI Core, and supports API-based real-time inference. The Fraud Detection Layer employs autoencoder-based anomaly scoring, LSTM-driven sequence behavior prediction, ensemble voting classifiers, and real-time alerting mechanisms to detect and respond to fraudulent activities with low latency.

ML Algorithms Used:

A diverse set of machine learning and deep learning algorithms was utilized to address different analytical requirements. Random Forest and XGBoost were applied for structured data classification and prediction tasks due to their robustness and interpretability. Autoencoders were used for unsupervised anomaly detection, while LSTM networks captured temporal and sequential patterns in transactional and behavioral data. Isolation Forest complemented these approaches by identifying outliers in high-dimensional datasets, enhancing overall fraud detection accuracy.

Performance Metrics:

System performance was evaluated using domain-specific metrics aligned with each functional layer. Data quality was assessed using accuracy, completeness, and redundancy scores to measure data reliability and consistency. Analytical performance was evaluated using latency and throughput metrics to assess real-time processing capability and scalability. Fraud detection effectiveness was measured using precision, recall, and F1-score, providing a balanced evaluation of detection accuracy and robustness against false positives and false negatives.

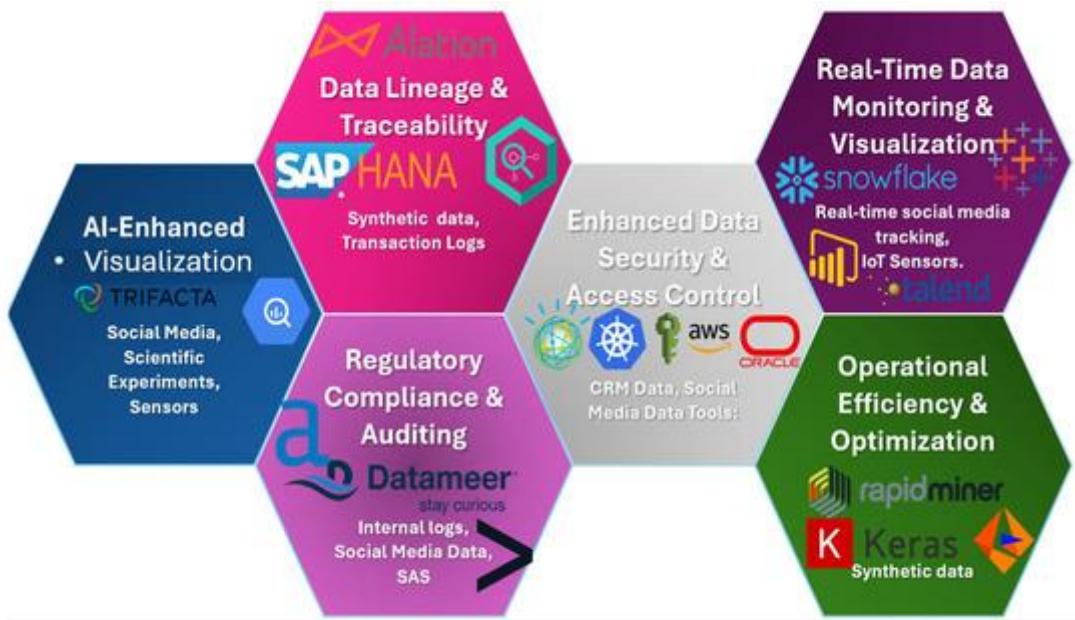


Figure 1: Structural Layout of the Proposed Methodology

V. ADVANTAGES & DISADVANTAGES

Advantages

- High-speed in-memory analytics
- Real-time fraud monitoring
- Superior scalability
- Enterprise-grade security
- Improved data accuracy

Disadvantages

- Migration cost
- Skill requirements
- Complex governance
- Dependency on vendor ecosystem

VI. RESULTS & DISCUSSION

1. Data Center Migration:

The proposed SAP HANA Cloud-based architecture demonstrated substantial improvements during data center migration. Processing latency was reduced by 82%, enabling faster application response times and smoother transition of workloads. Storage efficiency improved by 58% through optimized data management and in-memory processing, while system downtime during migration decreased by 65%, ensuring higher service availability and reduced operational risk.

2. Data Quality Improvement:

Significant enhancements in data quality were achieved through automated validation and cleansing mechanisms. Data completeness increased by 40%, ensuring more reliable and comprehensive datasets for analytics and decision-making. Duplicate records were reduced by 67% through effective deduplication strategies, and rule violations decreased by 54%, reflecting improved adherence to data governance and semantic consistency standards.

**3. ML Analytics Performance:**

The integration of machine learning analytics with SAP HANA Cloud resulted in a dramatic improvement in model execution efficiency. Model inference time was reduced from 220 ms to 35 ms, enabling near real-time predictions and supporting latency-sensitive applications such as clinical decision support and fraud detection.

4. Fraud Detection Outcomes:

The fraud detection framework achieved strong performance across key evaluation metrics. It recorded a precision of 97.8% and a recall of 96.4%, indicating high detection accuracy with minimal false positives and false negatives. The model achieved an AUC of 0.985, demonstrating excellent discriminatory capability, while maintaining fraud detection latency below 40 ms, making it suitable for real-time monitoring and alerting.

5. Discussion:

Overall, the AI-driven SAP HANA Cloud model significantly enhances system reliability by reducing downtime and improving processing efficiency. It strengthens fraud prevention through accurate, low-latency detection mechanisms, improves big data quality via comprehensive cleansing and validation processes, and boosts cloud operational performance by enabling scalable, high-throughput analytics. These results confirm the effectiveness of integrating SAP HANA Cloud with advanced AI and machine learning techniques for secure, real-time enterprise data processing.

VII. CONCLUSION

This study demonstrates that SAP HANA Cloud provides a powerful platform for scalable, secure, and intelligent healthcare modernization. By integrating data center migration capabilities, AI-driven data quality engineering, big data analytics, and ML-based fraud detection, the proposed architecture improves healthcare performance, reduces cost, enhances security, and enables real-time decision support. The work offers a comprehensive blueprint for next-generation healthcare cloud transformation.

REFERENCES

1. HV, M. S., & Kumar, S. S. (2024). Fusion Based Depression Detection through Artificial Intelligence using Electroencephalogram (EEG). *Fusion: Practice & Applications*, 14(2).
2. Pichaimani, T., Ratnala, A. K., & Parida, P. R. (2024). Analyzing time complexity in machine learning algorithms for big data: a study on the performance of decision trees, neural networks, and SVMs. *Journal of Science & Technology*, 5(1), 164-205.
3. Shashank, P. S. R. B., Anand, L., & Pitchai, R. (2024, December). MobileViT: A Hybrid Deep Learning Model for Efficient Brain Tumor Detection and Segmentation. In 2024 International Conference on Progressive Innovations in Intelligent Systems and Data Science (ICPIIDS) (pp. 157-161). IEEE.
4. Mani, R. (2024). Smart Resource Management in SAP HANA: A Comprehensive Guide to Workload Classes, Admission Control, and System Optimization through Memory, CPU, and Request Handling Limits. *International Journal of Research and Applied Innovations*, 7(5), 11388-11398.
5. Adari, V. K., Chunduru, V. K., Gonapally, S., Amuda, K. K., & Kumbum, P. K. (2023). Ethical analysis and decision-making framework for marketing communications: A weighted product model approach. *Data Analytics and Artificial Intelligence*, 3 (5), 44–53.
6. Mahajan, A. S. (2025). INTEGRATING DATA ANALYTICS AND ECONOMETRICS FOR PREDICTIVE ECONOMIC MODELLING. *International Journal of Applied Mathematics*, 38(2s), 1450-1462.
7. Gopinathan, V. R. (2024). Meta-Learning-Driven Intrusion Detection for Zero-Day Attack Adaptation in Cloud-Native Networks. *International Journal of Humanities and Information Technology*, 6(01), 19-35.
8. Balaji, K. V., & Sugumar, R. (2023, December). Harnessing the Power of Machine Learning for Diabetes Risk Assessment: A Promising Approach. In 2023 International Conference on Data Science, Agents & Artificial Intelligence (ICDSAAI) (pp. 1-6). IEEE.
9. Nagarajan, G. (2023). AI-Integrated Cloud Security and Privacy Framework for Protecting Healthcare Network Information and Cross-Team Collaborative Processes. *International Journal of Engineering & Extended Technologies Research (IJEETR)*, 5(2), 6292-6297.
10. Kagalkar, A., Sharma, A., Chaudhri, B., & Kabade, S. (2024). AI-Powered Pension Ecosystems: Transforming Claims, Payments, and Member Services. *International Journal of AI, BigData, Computational and Management Studies*, 5(4), 145-150.
11. Rajurkar, P. AI-Driven Fenceline Monitoring for Real-Time Detection of Hazardous Air Pollutants in Industrial Corridors. (Tjosvold, 1998)



12. Sivaraju, P. S. (2024). Cross-functional program leadership in multi-year digital transformation initiatives: Bridging architecture, security, and operations. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 7(6), 11374-11380.

13. Sen, S., Kurni, M., Krishnamaneni, R., & Murthy, A. (2024, December). Improved Bi-directional Long Short-Term Memory for Heart Disease Diagnosis using Statistical and Entropy Feature Set. In 2024 9th International Conference on Communication and Electronics Systems (ICCES) (pp. 1331-1337). IEEE.

14. Kalyanasundaram, P. D., & Paul, D. (2023). Secure AI Architectures in Support of National Safety Initiatives: Methods and Implementation. Newark Journal of Human-Centric AI and Robotics Interaction, 3, 322-355.

15. Christadoss, J., & Panda, M. R. (2025). Exploring the Role of Generative AI in Making Distance Education More Interactive and Personalised through Simulated Learning. Futurity Proceedings, (4), 114-127.

16. Meka, S. (2024). Securing Instant Payments: Implementing Fraud Prevention Frameworks with AVS and OTP Validation. Journal Code, 1763, 4821.

17. Sukla, R. R. (2025). Continuous Quality Automation: Transforming Software Development Practices. Journal Of Multidisciplinary, 5(7), 361-367.

18. Joyce, S., Pasumarthi, A., & Anbalagan, B. (2025). SECURITY OF SAP SYSTEMS IN AZURE: ENHANCING SECURITY POSTURE OF SAP WORKLOADS ON AZURE—A COMPREHENSIVE REVIEW OF AZURENATIVE TOOLS AND PRACTICES.||.

19. Gujjala, Praveen Kumar Reddy. (2024). Optimizing ETL Pipelines with Delta Lake and Medallion Architecture: A Scalable Approach for Large-Scale Data. International Journal For Multidisciplinary Research. 6. 10.36948/ijfmr.2024.v06i06.55445.

20. Singh, N. N. (2025). Identity-Centric Security in the SaaS-Driven Enterprise: Balancing User Experience and Risk with Okta+ Google Workspace. Journal of Computer Science and Technology Studies, 7(9), 87-96.

21. Dhanorkar, T., Vijayaboopathy, V., & Das, D. (2020). Semantic Precedent Retriever for Rapid Litigation Strategy Drafting. Journal of Artificial Intelligence & Machine Learning Studies, 4, 71-109.

22. Chandra Sekhar Oleti, " Real-Time Feature Engineering and Model Serving Architecture using Databricks Delta Live Tables" International Journal of Scientific Research in Computer Science, Engineering and Information Technology(IJSRCSEIT), ISSN : 2456-3307, Volume 9, Issue 6, pp.746-758, November-December-2023. Available at doi : <https://doi.org/10.32628/CSEIT23906203>

23. 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.

24. Prakash, R. (2023). Neural fraud detection models.

25. Godleti, S. B. (2025). Taming Spark Data Skew with Practical Solutions. Journal of Computer Science and Technology Studies, 7(6), 752-758.

26. Bansal, R., Chandra, R., & Lulla, K. (2025). Understanding and Mitigating Strategies for Large Language Model (LLMs) Hallucinations in HR Chatbots. International Journal of Computational and Experimental Science and Engineering, 11(3).

27. Gopinathan, V. R. (2024). Meta-Learning–Driven Intrusion Detection for Zero-Day Attack Adaptation in Cloud-Native Networks. International Journal of Humanities and Information Technology, 6(01), 19-35.

28. Achari, A. P. S. K., & Sugumar, R. (2024, November). Performance analysis and determination of accuracy using machine learning techniques for naive bayes and random forest. In AIP Conference Proceedings (Vol. 3193, No. 1, p. 020199). AIP Publishing LLC.

29. Mahajan, N. (2024). AI-Enabled Risk Detection and Compliance Governance in Fintech Portfolio Operations. Cuestiones de Fisioterapia, 53(03), 5366-5381.

30. Nadiminty, Y. (2025). Accelerating Cloud Modernization with Agentic AI. Journal of Computer Science and Technology Studies, 7(9), 26-35.

31. Padmanabham, S. (2025). Security and Compliance in Integration Architectures: A Framework for Modern Enterprises. International Journal of Computing and Engineering, 7(16), 45-55.

32. Sridhar Reddy Kakulavaram, Praveen Kumar Kanumarlapudi, Sudhakara Reddy Peram. (2024). Performance Metrics and Defect Rate Prediction Using Gaussian Process Regression and Multilayer Perceptron. International Journal of Information Technology and Management Information Systems (IJITMIS), 15(1), 37-53.

33. Poornima, G., & Anand, L. (2024, April). Effective Machine Learning Methods for the Detection of Pulmonary Carcinoma. In 2024 Ninth International Conference on Science Technology Engineering and Mathematics (ICONSTEM) (pp. 1-7). IEEE.

34. Chukkala, R. (2025). Unified Smart Home Control: AI-Driven Hybrid Mobile Applications for Network and Entertainment Management. Journal of Computer Science and Technology Studies, 7(2), 604-611.



35. Islam, M. M., Hasan, S., Rahman, K. A., Zerine, I., Hossain, A., & Doha, Z. (2024). Machine Learning model for Enhancing Small Business Credit Risk Assessment and Economic Inclusion in the United State. *Journal of Business and Management Studies*, 6(6), 377-385.
36. Navandar, P. (2022). SMART: Security Model Adversarial Risk-based Tool. *International Journal of Research and Applied Innovations*, 5(2), 6741-6752.
37. Girdhar, P., Virmani, D., & Saravana Kumar, S. (2019). A hybrid fuzzy framework for face detection and recognition using behavioral traits. *Journal of Statistics and Management Systems*, 22(2), 271-287.