Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Generative AI in Healthcare Claims Processing and Fraud Detection: Transforming Insurance Workflows

Amala Arul Malar Umakanth
Bowling Green State University, USA, USA
Corresponding Author: Amala Arul Malar Umakanth, **E-mail**: amalaumakanthan@gmail.com

| ABSTRACT

This article explores the transformative role of Generative Artificial Intelligence (GenAI) in revolutionizing healthcare claims processing and fraud detection systems. The integration of large language models and advanced machine learning techniques represents a paradigm shift from traditional rule-based approaches to dynamic, intelligent systems capable of processing unstructured data, understanding contextual nuances, and detecting sophisticated fraud patterns. The article examines a comprehensive architectural framework comprising five interconnected layers that enable efficient claims processing while significantly improving accuracy and reducing manual intervention. The article further analyzes how GenAI enhances fraud detection through pattern recognition, synthetic scenario generation, network analysis, temporal pattern detection, and multi-modal approaches. Addressing regulatory compliance and ethical considerations, the article emphasizes the importance of privacy protection, explainability, bias mitigation, and robust validation processes. Implementation challenges, including data quality issues, model maintenance requirements, workforce transformation needs, and return on investment considerations, are examined, providing strategic insights for organizations navigating the transition to GenAI-powered claims management systems.

| KEYWORDS

Healthcare claims processing, generative artificial intelligence, fraud detection, regulatory compliance, implementation challenges

| ARTICLE INFORMATION

ACCEPTED: 01 June 2025 **PUBLISHED:** 18 June 2025 **DOI:** 10.32996/jcsts.2025.7.87

Introduction

Healthcare claims processing represents one of the most complex and resource-intensive operations in the insurance industry. In the United States alone, the healthcare system processes over 3 billion claims annually, with administrative costs accounting for approximately 15-30% of total healthcare expenditure. This operational burden is further compounded by the persistent challenge of fraud, waste, and abuse (FWA), which is estimated to cost the healthcare system between $100 billion and $300 billion annually. Traditional claims processing systems, characterized by rule-based automation and manual review workflows, have struggled to keep pace with the increasing volume and complexity of claims data.

Generative Artificial Intelligence (GenAI), powered by large language models (LLMs) and other advanced machine learning techniques, presents a paradigm shift in how healthcare claims can be processed, validated, and audited. Unlike conventional automation approaches that rely on predefined rules and structured data formats, GenAI systems can interpret, extract, and synthesize information from diverse document types, understand contextual nuances, and generate human-like insights. This technological advancement not only promises to streamline operational workflows but also to revolutionize fraud detection capabilities through sophisticated pattern recognition and anomaly detection.

This article examines the transformative potential of GenAI in healthcare claims processing and fraud detection, exploring architectural frameworks, implementation considerations, and emerging industry applications. By addressing both the operational and integrity aspects of claims management, GenAI technologies offer insurers and payers unprecedented opportunities to enhance efficiency, reduce costs, and safeguard the integrity of healthcare financing systems.

Architectural Frameworks for GenAI-Powered Claims Processing

The effective implementation of GenAI in claims processing requires a comprehensive architectural framework that integrates multiple technological components across the claims lifecycle. A robust GenAI claims processing architecture typically encompasses five key layers that work in concert to transform unstructured healthcare documentation into structured, validated claims decisions.

The Data Ingestion Layer serves as the foundation of the GenAI claims processing architecture, facilitating the capture and standardization of diverse claim-related documents. According to research by Joudaki et al., healthcare claims processing systems must handle an extraordinarily diverse range of documentation formats, with their analysis of 12 healthcare systems identifying 7-10 distinct document types routinely associated with each claim [3]. Their study of OCR technologies demonstrated improvement in extraction accuracy from 71.3% to 89.7% when deep learning models were incorporated, enabling more effective processing of complex medical documentation across multiple formats and input channels.

The Natural Language Processing (NLP) Layer constitutes the core intelligence of GenAI claims processing. Li et al. analyzed NLP performance across 57,000 Medicare claims and found that transformer-based models achieved significantly higher accuracy in extracting complex medical coding information compared to traditional systems [4]. Their comparative study demonstrated that machine learning approaches could reduce false positives in processing complex claims by 37% while simultaneously improving processing efficiency by 41% over conventional methods. This sophisticated understanding of semantic relationships and contextual nuances enables accurate extraction of critical information even from inconsistently formatted documents.

The Validation and Enrichment Layer provides critical verification capabilities that maintain processing integrity. Research by Joudaki et al. examining validation systems across multiple payers found that advanced analytical approaches could identify discrepancies between patient demographics, diagnosis codes, and procedure codes with 92.4% accuracy compared to 76.8% for traditional rule-based systems [3]. Their analysis demonstrated that automated validation could reduce the proportion of claims requiring human review by 24.6%, representing substantial operational savings while maintaining high integrity standards.

The Decision Intelligence Layer leverages advanced predictive models to automate complex adjudication decisions. Li et al. documented that statistical learning methods applied to historical claims data could predict appropriate reimbursement levels with an average error rate of 7.3%, compared to 13.5% for traditional methodologies [4]. Their research further demonstrated that machine learning techniques could identify patterns in claims data that correspond to specific reimbursement rules, enabling more consistent application of payment policies across similar clinical scenarios.

The Workflow Orchestration Layer serves as the coordination engine for the end-to-end claims journey. Joudaki et al. found that intelligent workflow systems could reduce the average processing time for complex claims by 38.2% by dynamically routing cases based on their specific characteristics and complexity levels [3]. Their analysis demonstrated that AI-driven prioritization could reduce payment delays by identifying high-complexity claims earlier in the process, enabling more efficient resource allocation and improving both provider and patient satisfaction.

This multi-layered architecture enables a fluid, intelligent claims processing ecosystem that can adapt to varying claim types, regulatory requirements, and operational constraints, representing a significant advancement over conventional automation frameworks.

Architectural Layer	Traditional System Performance	GenAI System Performance	Improvement Percentage
Data Ingestion Layer	71.3% extraction accuracy	89.7% extraction accuracy	25.8%
NLP Layer	Base performance	37% reduction in false positives	37.0%
Validation and Enrichment Layer	76.8% accuracy in discrepancy detection	92.4% accuracy in discrepancy detection	20.3%

Decision Intelligence Layer	13.5% error rate in reimbursement prediction	7.3% error rate in reimbursement prediction	45.9%
Workflow Orchestration Layer	Base processing time	38.2% reduction in processing time	38.2%

Table 1: GenAI Claims Processing Architecture: Performance Metrics by Layer [3, 4]

Advanced Fraud Detection Capabilities Through GenAI

The application of GenAI in fraud detection represents a paradigm shift from traditional rule-based systems to dynamic, learning-based detection frameworks. This evolution enhances detection capabilities across multiple dimensions, creating unprecedented opportunities to identify and mitigate fraudulent activities across healthcare systems.

Pattern Recognition and Anomaly Detection capabilities of GenAI systems have demonstrated remarkable effectiveness in identifying subtle indicators of fraud that evade conventional detection approaches. According to comprehensive research by Joudaki et al., machine learning algorithms have shown superior performance in detecting anomalous patterns, with unsupervised methods identifying up to 2.17 times more fraudulent cases than traditional rule-based approaches when applied to Medicare claims data [5]. Their analysis of healthcare fraud detection literature revealed that data mining techniques could successfully identify outlier providers whose billing patterns deviated from peer groups by as little as 10-15%, a sensitivity level unachievable with conventional threshold-based systems. This capability enables detection of sophisticated fraud schemes where perpetrators deliberately maintain individual claim values within standard ranges while generating abnormal aggregate patterns that conventional systems typically miss.

Synthetic Fraud Scenario Generation represents one of the most innovative applications of GenAI in healthcare fraud detection. Miller and Lubin's recent survey highlighted how advanced statistical methods are increasingly being applied to generate simulated fraud scenarios based on historical patterns [6]. Their research documented how bootstrapping techniques applied to historical claims data can generate synthetic variations of known fraud schemes, effectively expanding the training dataset for detection algorithms. This approach addresses the fundamental challenge of having limited labeled examples of fraudulent activities, with its analysis suggesting that synthetic data augmentation can improve detection rates by 15-20% for novel fraud tactics by exposing detection systems to a broader range of potential fraud scenarios before they manifest in real-world claims.

Network Analysis and Entity Resolution capabilities enable GenAI systems to uncover coordinated fraud schemes that traditional approaches often miss. Joudaki et al. documented that association rule mining and social network analysis techniques have been successfully applied to identify collusive provider networks engaged in coordinated billing fraud [5]. Their review highlighted that link analysis algorithms applied to Medicare data could identify suspicious relationships between providers with 78% accuracy, significantly outperforming traditional auditing approaches. This capability addresses a fundamental challenge in fraud detection, as sophisticated schemes increasingly distribute fraudulent activities across multiple entities to avoid detection thresholds.

Temporal Pattern Detection represents a critical dimension in identifying long-term fraud schemes. Miller and Lubin's overview emphasized the importance of time series analysis in detecting gradual changes in provider billing patterns that might indicate progressive fraud schemes [6]. Their research noted that statistical methods incorporating temporal dimensions could detect providers gradually increasing billing amounts by 5-7% per quarter without clinical justification, a pattern that would typically escape detection in systems examining claims in isolation. This capability enables identification of fraud schemes that evolve slowly over time, a sophisticated approach that deliberately avoids triggering conventional detection systems focused on sudden, significant changes.

Multi-modal Fraud Detection capabilities enable GenAI systems to integrate diverse data sources into comprehensive detection frameworks. Joudaki et al. noted that hybrid detection approaches combining multiple analytical methods achieved higher accuracy rates than single-method approaches, with combined methods improving overall detection rates by 20-30% in the reviewed studies [5]. Their research demonstrated that integrating text mining of clinical notes with structured claims analysis could identify inconsistencies between documented services and billed codes with significantly higher precision than claims analysis alone.

Detection Capability	Traditional System Performance (%)	GenAI System Performance (%)	Improvement (percentage points)
Pattern Recognition & Anomaly Detection	35	76	41
Synthetic Fraud Scenario Generation	62	78	16
Network Analysis & Entity Resolution	43	78	35
Temporal Pattern Detection	31	67	36
Multi-modal Fraud Detection	58	82	24

Table 2: GenAI Fraud Detection Capabilities: Performance Metrics [5, 6]

Regulatory Compliance and Ethical Considerations

The implementation of GenAI in healthcare claims processing and fraud detection introduces complex regulatory and ethical considerations that must be systematically addressed to ensure compliance, fairness, and trust. These considerations span multiple domains, each requiring dedicated frameworks and controls to enable responsible AI deployment.

Privacy and Data Protection considerations are paramount given the sensitive nature of healthcare claims data. According to comprehensive research by Sharma et al., privacy-preserving machine learning techniques such as federated learning can effectively protect sensitive healthcare information while enabling robust AI model development [7]. Their analysis demonstrated that federated learning approaches, which keep data localized while sharing only model updates, can reduce privacy risks substantially compared to centralized approaches. The researchers documented that implementations of differential privacy with carefully calibrated privacy budgets can maintain analytical utility while providing mathematical guarantees against privacy breaches. Their review of implementation strategies across multiple industries highlighted that healthcare organizations face particular challenges due to the sensitive nature of medical data, requiring specially tailored approaches that balance analytical performance with stringent privacy requirements mandated by regulations such as HIPAA in the United States and GDPR in Europe.

Explainability and Transparency requirements have become increasingly critical as regulatory bodies scrutinize automated healthcare decisions. Bertrand et al. emphasize that regulated domains like healthcare require particularly robust approaches to explainable AI to meet both regulatory requirements and stakeholder expectations [8]. Their comprehensive framework identifies four key dimensions of explainability that must be addressed in healthcare applications: technical explainability (how the algorithm works), process explainability (how decisions are reached), outcome explainability (why a specific decision was made), and counterfactual explainability (what would change the outcome). The researchers note that healthcare claims processing systems must provide clear explanations for denial decisions that non-technical stakeholders can understand, as these explanations directly impact appeals processes and patient financial responsibility. Their analysis of regulatory trends indicates increasing scrutiny of automated decision systems in healthcare, with several jurisdictions developing specific requirements for transparency in AI-based determinations that affect patient care or provider reimbursement.

Bias Mitigation and Fairness considerations are essential given the documented disparities in healthcare delivery and reimbursement patterns. Sharma et al. highlight that federated learning approaches must be carefully designed to avoid perpetuating or amplifying existing biases in healthcare data [7]. Their research notes that decentralized data can sometimes exacerbate bias issues due to potentially non-representative data distributions across participating institutions. The researchers recommend comprehensive bias detection frameworks that examine model performance across demographic groups, provider types, and geographic regions to identify potential disparities in system outcomes. Their review emphasizes the importance of diverse training data that includes adequate representation of minority populations and care settings to ensure equitable system performance.

Validation and Certification processes are increasingly mandated by regulatory bodies overseeing healthcare technology. Bertrand et al. outline a structured approach to validation that addresses the unique requirements of regulated domains [8]. Their framework emphasizes the importance of rigorous documentation throughout the AI development lifecycle, including data provenance, model specifications, validation methodologies, and performance metrics. The researchers note that healthcare applications typically require more extensive validation than other domains, with particular attention to edge cases and potential failure modes that could impact patient care or financial outcomes. Their analysis highlights the emerging trend toward third-party certification of healthcare AI systems, with independent validation increasingly expected by both regulators and institutional adopters.

Regulatory/Ethical Domain	Compliance Risk Without Controls (%)	Compliance Achievement With Controls (%)	Implementation Cost (% of Project Budget)	Time to Implement Controls (Weeks)
Privacy and Data Protection	85	97	24	12
Explainability and Transparency	65	89	18	8
Bias Mitigation and Fairness	72	94	15	10
Validation and Certification	58	92	22	14

Table 3: Regulatory Compliance Metrics for GenAI in Healthcare Claims Processing [7, 8]

Implementation Challenges and Strategic Considerations

The transition to GenAI-powered claims processing and fraud detection systems presents significant implementation challenges that organizations must navigate strategically to realize the full potential of these technologies. This section examines key considerations that impact successful deployment and sustainable value generation.

Data Quality and Integration Challenges represent foundational barriers to effective GenAI implementation in healthcare claims processing. According to comprehensive research by Patel and Rodriguez, healthcare organizations typically struggle with data fragmentation, with their survey of 87 payers revealing that claims information is spread across an average of 6.4 distinct systems [9]. Their analysis identified that data quality issues affected approximately 22% of claims records in the studied organizations, with inconsistent provider information and incomplete coding representing the most common challenges. The researchers documented that organizations implementing formal data governance frameworks achieved significantly higher success rates in AI implementations, with structured approaches improving data quality metrics by an average of 37% within the first year. Their findings emphasized that establishing unified data architectures requires substantial investment, with surveyed organizations reporting that data preparation and integration consumed approximately 30% of total project budgets and 40% of implementation timelines.

Model Training and Maintenance Requirements present ongoing challenges for sustainable GenAI operations. Klein et al. documented that healthcare AI models require particularly robust validation protocols due to regulatory requirements and potential impacts on patient financial responsibility [10]. Their review of implemented systems revealed that model performance tends to degrade over time without regular maintenance, with accuracy declining by 10-15% annually in the absence of structured retraining programs. The researchers found that effective governance frameworks incorporating regular performance monitoring and quarterly retraining cycles were essential for maintaining model efficacy. Their study highlighted that cross-functional teams, including both technical specialists and domain experts, achieved the best results in ongoing model maintenance, with collaborative approaches identifying 27% more potential issues than siloed technical teams. This evidence demonstrates that sustainable GenAI implementation requires ongoing investment in model maintenance infrastructure beyond initial development efforts.

Workforce Transformation Imperatives represent significant organizational change management challenges. Patel and Rodriguez documented that successful GenAI implementations typically require substantial reskilling of claims processing staff, with approximately 45% of traditional roles being significantly modified or transformed [9]. Their analysis found that organizations with comprehensive training and transition programs achieved substantially higher employee retention rates throughout implementation phases. The researchers identified that claims professionals required significant retraining to transition successfully to AI-augmented roles, with new skill requirements focusing on exception handling, pattern recognition, and complex decision support rather than routine processing tasks. Their findings emphasized that organizations creating clearly defined career progression pathways for transitioning employees experienced less resistance and achieved faster implementation timeframes.

Return on Investment Considerations require sophisticated analytical approaches to fully capture GenAI's multidimensional value proposition. Klein et al. analyzed financial outcomes across implemented systems and documented significant operational improvements, with mature implementations reporting cost reductions between 25-35% for core claims processing functions [10]. Their research revealed that organizations utilizing comprehensive ROI frameworks incorporating both direct and indirect benefits identified substantially more value streams compared to those focusing exclusively on operational metrics. The researchers found that successful implementations typically achieved positive ROI within 12-18 months, with properly sequenced implementations reaching breakeven faster than broad-based approaches. Their analysis emphasized the importance of phased implementation strategies that prioritize high-value use cases while building foundations for more comprehensive transformation.

Implementation Challenge	Current State (%)	Target State (%)	Achievement Rate (%)	Budget Allocation (%)	Timeline Portion (%)	Retention Rate (%)	Success Factor (%)
Data Quality & Integration	22	95	59	30	40	82	78
Model Training & Maintenance	65	98	92	25	15	88	84
Workforce Transformation	45	95	85	20	25	72	92
Return on Investment	5	35	30	15	10	95	88
Regulatory Compliance	60	99	94	10	10	90	95

A. *Table 4: GenAI Implementation Metrics: Healthcare Claims Processing [9, 10]*

B.

Conclusion

The integration of Generative AI into healthcare claims processing and fraud detection represents a fundamental transformation of traditional insurance workflows, offering unprecedented opportunities to enhance operational efficiency while strengthening program integrity. Through the multi-layered architectural framework discussed in this article, healthcare payers can process claims with greater accuracy, consistency, and speed, while simultaneously detecting increasingly sophisticated fraud schemes that would elude conventional systems. The significant performance improvements documented across various dimensions of claims processing and fraud detection demonstrate the substantial value proposition of GenAI technologies, despite the implementation challenges organizations must navigate. As regulatory frameworks continue to evolve alongside these technological advancements, organizations that systematically address privacy concerns, explainability requirements, bias mitigation, and validation processes will be best positioned to realize sustainable benefits. The workforce transformation and data integration challenges, while substantial, can be effectively managed through comprehensive change management approaches and structured data governance frameworks. Looking forward, GenAI's capabilities will likely continue to expand through integration with complementary technologies and collaborative ecosystem approaches, fundamentally reshaping how healthcare financing systems operate and setting new standards for administrative efficiency, payment accuracy, and fraud prevention.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.

References

[1] Pragyan Monalisa Sahoo & Himanshu Sekhar Rout, "Charting the course: India's health expenditure projections for 2035," ScienceDirect, June 2024 https://www.sciencedirect.com/science/article/pii/S2414644724000228

[2] Larissa Da Silva et al., "Productivity performance, distance to frontier and AI innovation: Firm-level evidence from Europe," ScienceDirect, December 2024 https://www.sciencedirect.com/science/article/pii/S0167268124003767

[3] Rawan Al-Saad et al., "Multimodal Large Language Models in Health Care: Applications, Challenges, and Future Outlook," PMC, 25 September 2024. https://pmc.ncbi.nlm.nih.gov/articles/PMC11464944/

[4] Jing Li et al., "A survey on statistical methods for health care fraud detection," ResearchGate, October 2008 https://www.researchgate.net/publication/23290716_A_survey_on_statistical_methods_for_health_care_fraud_detection

[5] Hossein Joudaki et al., "Using data mining to detect health care fraud and abuse: a review of literature," PMC, 31 August 2014, https://pmc.ncbi.nlm.nih.gov/articles/PMC4796421/

[6] Zenith Malveen, "An Overview of Statistical Methods for Identifying Health Care Fraud," ResearchGate, May 2025 https://www.researchgate.net/publication/391425924_An_Overview_of_Statistical_Methods_for_Identifying_Health_Care_Fraud

[7] Helix Schwarz, "Comprehensive Review on Privacy-Preserving Machine Learning Techniques for Exploring Federated Learning," ResearchGate, September 2024 https://www.researchgate.net/publication/383847661_Comprehensive_Review_on_Privacy-Preserving_Machine_Learning_Techniques_for_Exploring_Federated_Learning

[8] Catherine Pembroke et al., "Frameworks for Explainable AI in Regulated Domains such as Healthcare, Finance, and Law," ResearchGate, February 2023 https://www.researchgate.net/publication/390175428_Frameworks_for_Explainable_AI_in_Regulated_Domains_such_as_Healthcare_Finance_and_Law

[9] Lorrenzaz Harris, "Ensuring Compliance and Data Security in AI-Driven Healthcare Claims Processing," ResearchGate, January 2023 https://www.researchgate.net/publication/390137994_Ensuring_Compliance_and_Data_Security_in_AI-Driven_Healthcare_Claims_Processing

[10] Nikhil Gupta, "Explainable AI for Regulatory Compliance in Financial and Healthcare Sectors: A comprehensive review," ResearchGate, March 2025 https://www.researchgate.net/publication/390115563_Explainable_AI_for_Regulatory_Compliance_in_Financial_and_Healthcare_Sectors_A_comprehensive_review