

Enhancing Healthcare Claim Processing with Generative AI: Leveraging AWS Bedrock and SageMaker for Efficiency and Accuracy

Perumalsamy Ravindran

Anna University, India

Abstract

Healthcare claims processing faces significant challenges in managing operational efficiency and cost-effectiveness while maintaining regulatory compliance. Integrating Generative AI technologies, specifically through AWS Bedrock and SageMaker platforms, offers transformative solutions for automating and optimizing claims processing workflows. Implementing these advanced technologies substantially improves processing speed, accuracy, and cost reduction across various healthcare organizations. Through intelligent document processing, anomaly detection, fraud prevention, and predictive adjudication capabilities, healthcare providers have significantly reduced administrative overhead while improving claims accuracy and processing efficiency. AI-driven solutions have enabled healthcare organizations to streamline operations, reduce manual intervention requirements, and enhance overall claims management effectiveness. These implementations have demonstrated marked improvements in provider satisfaction, reduced processing times, and increased operational efficiency while maintaining high accuracy standards and regulatory compliance.

Keywords: Healthcare Claims Processing, Generative AI, AWS Bedrock, Machine Learning, Fraud Detection

Enhancing Healthcare
Claim Processing with
Generative AI:
Leveraging AWS
Bedrock and
SageMaker for
Efficiency and
Accuracy



Introduction

Healthcare claim processing represents a mounting challenge in the U.S. healthcare system, with recent data from CAQH revealing the processing of over 1.8 billion claims annually as of 2024. Traditional manual processing methods burden the healthcare industry with significant operational inefficiencies and costs. According to the 2024 CAQH Index Report, the healthcare industry could save \$25.8 billion annually by fully transitioning from manual to electronic transactions for claims-related processes. Manual claim submission costs an average of \$7.31 per transaction, compared to just \$1.12 for electronic submissions, representing a potential savings of \$6.19 per claim through automation [1].

The complexity of claims processing extends beyond basic submission costs. The Optum 2024 Revenue Cycle Denials Index highlights that healthcare organizations face an average claim denial rate of 11%, with a staggering 67% of these denials being avoidable through improved front-end processes and automation. The financial impact is substantial, with providers spending an average of \$118 per claim appeal, typically taking 45 to 60 days to resolve denied claims. Furthermore, larger healthcare organizations managing high claim volumes report spending up to \$8.6 million annually on denial management processes alone [2].

Evolving regulatory requirements and the increasing complexity of healthcare services compound these operational challenges. The 2024 CAQH Index indicates that the adoption of fully electronic processes varies significantly across different types of transactions, with claim status inquiries achieving 89% electronic adoption, while prior authorization remains at only 29% electronic adoption. This disparity in automation creates significant operational bottlenecks and increases the potential for errors in the claims-processing workflow [1].

Generative AI emerges as a transformative solution to these challenges, offering sophisticated capabilities in document processing, natural language understanding, and predictive analytics. By leveraging AWS Bedrock and SageMaker, healthcare organizations can implement GenAI solutions that address the pain points identified in current industry reports. These technologies show particular promise in addressing the high-cost, manual-intensive processes that currently dominate healthcare claims management, potentially significantly reducing the \$372 billion spent annually on administrative complexity in the U.S. healthcare system [1].

Metric	Traditional Processing
Claim Submission Cost	\$7.31 per transaction
Denial Rate	11%
Appeal Resolution Time	45-60 days
Electronic Adoption (Claims Status)	89%
Electronic Adoption (Prior Auth)	29%

Table 1: Healthcare Claims Processing Metrics and Cost Implications [1, 2]

Technical Architecture

AWS Bedrock Integration

The proposed architecture leverages AWS Bedrock as the foundational framework for deploying and managing large language models (LLMs) in the healthcare claims processing pipeline. According to implementation studies by Min et al., AWS Bedrock's foundation models can process unstructured insurance documents with a mean latency of 2.8 seconds per document while maintaining a document

parsing accuracy of 92.7%. Their research demonstrated the system could handle complex insurance documentation, including policy documents averaging 15-20 pages and containing over 50 unique data fields [3].

The Document Processing Layer implements Claude and other foundation models through Bedrock's unified API, which has shown particular effectiveness in processing medical terminology and insurance-specific language. The system demonstrated an 87% reduction in manual document review time, with the ability to accurately extract and classify key insurance parameters from various document formats. Integration testing revealed that the system could maintain consistent performance while processing up to 10,000 documents daily, with error rates below 0.5% even during peak processing periods [3].

The Inference Pipeline establishes comprehensive validation mechanisms that have proven particularly effective in insurance underwriting workflows. Min et al.'s implementation showed that the system could reduce the average underwriting decision time from 3-5 days to approximately 4.8 hours while maintaining accuracy above 91%. The automated validation system successfully processed 94.3% of straightforward cases without human intervention, allowing underwriters to focus on complex cases requiring specialized attention [3].

SageMaker Implementation

Amazon SageMaker provides the machine learning infrastructure for processing high volumes of insurance claims data. Research by Nerella et al. demonstrated that their SageMaker-based implementation achieved a 78% reduction in claim processing time compared to traditional methods. Their system processed an average of 850 claims per hour with an accuracy rate of 93.5% in automated classification tasks. The deployment of fraud detection models showed promise, with the system correctly identifying 96% of fraudulent claims while maintaining a false positive rate of just 2.3% [4].

The custom model training pipeline integrates historical claim data through SageMaker's distributed training capabilities. Nerella et al.'s implementation successfully processed training datasets comprising 1.2 million historical claims, with model training completed in under 48 hours. Their feature engineering pipeline automatically processed 127 unique claim attributes, leading to a 34% improvement in classification accuracy compared to previous rule-based systems. The research demonstrated that continuous model retraining using SageMaker's A/B testing capabilities resulted in a sustained accuracy improvement of 0.8% per month over the six-month study period [4].

Metric	Performance Value
Document Processing Latency	2.8 seconds
Document Parsing Accuracy	92.7%
Manual Review Reduction	87%
Daily Document Processing	10,000
Peak Error Rate	<0.5%

Table 2: AWS Implementation Performance Indicators [5, 6]

Key Features and Capabilities

Intelligent Document Processing

The system's intelligent document processing capabilities significantly advance healthcare claims management efficiency. According to Singhal et al.'s research, healthcare organizations implementing digital-first claims processing solutions have achieved a 30-40% reduction in administrative costs. Their

analysis demonstrates that automated document processing can reduce claims handling time by 50-70% compared to traditional manual methods. The study further shows that providers utilizing advanced digital solutions have improved their clean claims rate to over 90%, significantly reducing the need for manual intervention and rework [5].

Digital transformation of claims processing has shown a substantial impact on operational efficiency. McKinsey's analysis reveals that organizations implementing comprehensive digital solutions have reduced their claims processing workforce requirements by 20-30% while improving accuracy rates. The research indicates that automated validation against provider databases has reduced error rates by 40-50% compared to manual processing. Furthermore, healthcare providers leveraging these technologies have reported a 15-20% reduction in their claims denial rates, translating to significant improvements in revenue cycle efficiency [5].

Anomaly Detection and Fraud Prevention

Implementing machine learning models for fraud detection has demonstrated remarkable effectiveness in healthcare claims processing. Herland et al.'s comprehensive study of healthcare fraud detection systems shows that machine learning approaches achieve detection accuracy rates of up to 95% for known fraud patterns. Their analysis of supervised learning methods reveals that gradient-boosting algorithms successfully identified fraudulent claims with a sensitivity of 77.4% and a specificity of 95.7%. The study also demonstrates that hybrid detection models combining multiple machine learning techniques improved overall fraud detection rates by 23% compared to single-model approaches [6].

The research indicates that pattern recognition algorithms are particularly effective in identifying emerging fraud schemes. Healthcare organizations implementing these systems reported a 60% improvement in early fraud detection capabilities. The study shows that machine learning models trained on historical claims data can process and analyze hundreds of features per claim in real-time, with decision times averaging less than one second per claim. Furthermore, the continuous learning capabilities of these systems demonstrated a 15% year-over-year improvement in fraud detection accuracy by incorporating newly identified fraud patterns [6].

Predictive Adjudication

Predictive analytics in claims processing has shown significant promise in improving operational efficiency. Singhal et al.'s research indicates that organizations implementing predictive adjudication systems have reduced their average claims processing time by 40-50%. Their analysis shows that automated systems can accurately process up to 85% of routine claims without human intervention, allowing claims professionals to focus on complex cases requiring specialized attention. Healthcare providers utilizing these predictive capabilities have reported a 25-30% reduction in their claims backlog and a 20-25% improvement in their first-pass payment rates [5].

Operational Area	Improvement Range
Administrative Cost Reduction	30-40%
Claims Handling Time Reduction	50-70%
Clean Claims Rate	>90%
Claims Denial Rate Reduction	15-20%
Workforce Requirement Reduction	20-30%

Table 3: Digital Transformation Impact on Healthcare Operations [5, 6]

Case Studies and Results

Implementation at Major Healthcare Provider

UnitedHealth Group's implementation of digital claims processing solutions demonstrates the transformative impact of AI-powered systems in healthcare operations. According to PYMNTS' analysis, UnitedHealth's digital transformation initiatives have led to significant operational efficiencies in their claims processing workflows. The organization processed over 1.7 billion claims digitally in the past year, achieving a 35% reduction in administrative costs compared to traditional processing methods. The implementation of automated systems has enabled UnitedHealth to process 98% of claims within 24 hours, significantly improving their payment velocity and provider satisfaction rates [7].

The provider's digital transformation extends beyond basic claims processing. Through the implementation of AI-driven solutions, UnitedHealth has reduced its per-member administrative costs by 15% year over year. The organization reported that its digital platforms now handle more than 80% of provider transactions automatically, with automated systems managing an average of 4.7 million claims daily. These improvements have contributed to a 20% reduction in provider administrative burden while maintaining high accuracy rates in claims processing [7].

Technology Provider Implementation Results

Google Cloud's Claims Acceleration Suite implementation across various healthcare organizations provides compelling evidence of the effectiveness of AI-driven claims processing. According to their comprehensive analysis, organizations utilizing the AI-enabled claims processing solution have reduced prior authorization processing times by up to 50%. The system has demonstrated the ability to process claims up to six times faster than traditional methods while maintaining high accuracy rates in automated decision-making [8].

The implementation shows particular strength in streamlining complex claims workflows. Healthcare organizations using the Claims Acceleration Suite reported that their claims specialists could handle 25% more cases daily than traditional processing methods. The AI-enabled system successfully automated up to 80% of administrative tasks associated with claims processing, allowing healthcare professionals to focus on more complex cases requiring human expertise. Furthermore, the implementation improved the accuracy of claims documentation and reduced the need for manual review of routine cases, leading to significant operational efficiencies across the claims processing workflow [8].

Implementation Metric	UnitedHealth
Claims Processed Annually	1.7 billion
Administrative Cost Reduction	35%
Processing Time Improvement	98% within 24h
Automation Level	80%

Table 4: Implementation Results Across Healthcare Organizations [7, 8]

Conclusion

Implementing Generative AI technologies in healthcare claims processing represents a significant advancement in addressing operational inefficiencies and cost challenges. The integration of AWS Bedrock and SageMaker has demonstrated remarkable improvements in processing efficiency, accuracy, and cost reduction across multiple healthcare organizations. Adopting digital-first solutions has enabled healthcare providers to significantly reduce administrative burdens while improving claims processing

accuracy and speed. Through advanced capabilities in document processing, fraud detection, and predictive analytics, healthcare organizations have achieved substantial improvements in operational efficiency and provider satisfaction. The successful implementations at major healthcare providers and technology platforms demonstrate the transformative potential of AI-driven solutions in modernizing claims processing workflows. The continued evolution and adoption of these technologies suggest a promising future for healthcare claims management, with opportunities for further enhancements in processing efficiency, fraud prevention, and regulatory compliance.

References

1. CAQH, "2024 CAQH Index Report: Benchmarks for Electronic Administrative Transactions and Opportunities for Automation," CAQH, Washington, DC, 2024. [Online]. Available: https://www.caqh.org/hubfs/Index/2024%20Index%20Report/CAQH_IndexReport_2024_FINAL.pdf
2. Optum, "The Optum 2024 Revenue Cycle Denials Index," Optum Healthcare Operations, Eden Prairie, MN, 2024. [Online]. Available: <https://www.optum.com/content/dam/o4-dam/resources/pdfs/e-books/2024-denials-index-for-providers.pdf>
3. Paul Min et al., "Streamline insurance underwriting with generative AI using Amazon Bedrock – Part 1," AWS Machine Learning Blog, Amazon Web Services, 01 Aug 2024. [Online]. Available: <https://aws.amazon.com/blogs/machine-learning/streamline-insurance-underwriting-with-generative-ai-using-amazon-bedrock-part-1/>
4. Harshavardhan Nerella et al., "Integrating AWS AI for Automated Insurance Claims Processing," ResearchGate, September 2024. [Online]. Available: https://www.researchgate.net/publication/385049423_Integrating_AWS_AI_for_Automated_Insurance_Claims_Processing
5. Shubham Singhal et al., "For better healthcare claims management, think 'digital first'," McKinsey & Company, June 2019. [Online]. Available: <https://www.mckinsey.com/industries/healthcare/our-insights/for-better-healthcare-claims-management-think-digital-first>
6. M Herland et al., "Healthcare Fraud Detection Using Machine Learning," ResearchGate, December 2024. [Online]. Available: [https://www.researchgate.net/publication/387470643_Healthcare_Fraud_Detection_Using_Machine_Learning#:~:text=Machine%20learning%20\(ML\)%20has%20emerged,patient%20records%2C%20and%20provider%20data](https://www.researchgate.net/publication/387470643_Healthcare_Fraud_Detection_Using_Machine_Learning#:~:text=Machine%20learning%20(ML)%20has%20emerged,patient%20records%2C%20and%20provider%20data)
7. PYMNTS, "UnitedHealth Sees Digital Solutions Slash Onboarding Costs," PYMNTS.com, January 16, 2025. [Online]. Available: <https://www.pymnts.com/earnings/2025/unitedhealth-bets-on-digital-innovation-to-improve-care/>
8. PR Newswire, "Google Cloud Unveils New AI-enabled Claims Acceleration Suite to Streamline Health Insurance Prior Authorization and Claims Processing, Helping Experts Make Faster, More Informed Decisions," Apr 13, 2023. [Online]. Available: <https://www.prnewswire.com/news-releases/google-cloud-unveils-new-ai-enabled-claims-acceleration-suite-to-streamline-health-insurance-prior-authorization-and-claims-processing-helping-experts-make-faster-more-informed-decisions-301796467.html>