

ORIGINAL RESEARCH

Open access

Explainable Gradient Boosting Model for Predicting Health Insurance Claim Denials Using Documentation Completeness, Procedure Codes, Diagnosis-Code Consistency, Payer Rules, and Prior Authorization History

Oliver Grant¹, David Clark^{1*}, Sophia Nguyen²

Abstract

Claim denials represent a major source of lost or delayed healthcare revenue. They are commonly driven by documentation gaps, coding inconsistencies, payer rule violations, and missing or invalid prior authorizations. Current denial prevention often depends on manual review and deterministic claim-scrubbing rules. These approaches may not capture complex payer-specific interactions among documentation quality, procedure codes, diagnosis codes, and authorization history. This article proposes an explainable gradient boosting model for estimating the probability that a health insurance claim could be denied before submission. The model is intended to identify the specific claim-level factors contributing to denial risk. The proposed framework uses a gradient-boosted tree ensemble trained conceptually on historical claims and remittance data. SHAP-based explanations would provide both global and local interpretability for denial-risk predictions. Conceptually, the model would return a denial risk score alongside an explanation of contributing factors such as incomplete documentation, diagnosis-code mismatch, expired authorization, or payer rule conflict. These outputs would support targeted pre-billing review rather than broad manual auditing. An explainable denial prediction model could help shift revenue cycle management from reactive appeals toward proactive prevention. Transparent reasoning would be essential for revenue cycle staff, clinical documentation teams, coders, and compliance stakeholders.

Keywords Explainable artificial intelligence, SHAP, Gradient boosting, Claim denial prediction, Healthcare revenue cycle, Prior authorization

*Correspondence:

David Clark
david.clark@gmail.com

¹ Department of Health Informatics and Smart Healthcare Systems, Faculty of Engineering, University of Glasgow, Glasgow, United Kingdom

² Department of Digital Medicine and Health Analytics, Faculty of Medicine, National University of Singapore, Singapore, Singapore

Introduction

Health insurance claim denials create administrative burden, delayed reimbursement, and avoidable revenue leakage for hospitals, physician practices, and integrated delivery systems. Predictive revenue cycle analytics has therefore become a growing area of interest, especially where historical claims and remittance data can be used to anticipate reimbursement outcomes before claim submission [1, 2]. Although denial rates vary across commercial, Medicare, and Medicaid populations, the underlying operational challenge is consistent: providers must identify preventable denials early

enough to intervene. Machine learning approaches for claims and cost prediction suggest that tabular administrative data can support prospective risk stratification when features are aligned with real adjudication processes [3-6].

Denials are rarely caused by a single isolated defect; they often reflect interactions among missing documentation, coding inconsistency, medical necessity logic, prior authorization failure, and payer-specific policy requirements. Clinical documentation integrity research emphasizes that documentation gaps can weaken the evidentiary basis for billing, while automated coding research shows that diagnosis and procedure representations are central to downstream reimbursement logic [7-10]. Prior authorization adds another layer of denial risk because coverage approval may depend on timing, service category, payer rule interpretation, and whether authorization details are correctly linked to the claim [11, 12]. A denial prediction model for revenue cycle use should therefore represent the claim as an integrated administrative, clinical, and payer-policy object rather than as a simple billing transaction.

Existing claim scrubbers and rule engines are useful for deterministic edits, but they are typically constrained by predefined logic and may not learn complex interactions from historical adjudication behavior. A rigid rule can flag a missing field, but it may not estimate how strongly that defect matters for a particular payer, procedure, diagnosis combination, or authorization history [1-3]. Broader healthcare machine learning literature also warns that predictive systems must be designed for operational use, interpretability, and stakeholder trust rather than for prediction alone [13-16]. In revenue cycle settings, this means that a model should not merely label a claim as high risk; it should explain what a specialist could review or correct before submission.

This article proposes an explainable gradient boosting model that could predict pre-submission denial risk while producing claim-level explanations for revenue cycle teams. Gradient-boosted tree models are well suited to heterogeneous tabular features such as payer identifiers, procedure codes, diagnosis codes, authorization attributes, and documentation completeness indicators. SHAP-based explanation methods could identify whether denial risk is primarily driven by incomplete documentation, diagnosis-procedure inconsistency, payer coverage rules, or prior authorization history. By combining predictive modeling with explainability, the proposed framework aims to support proactive denial prevention without replacing coder, biller, compliance, or clinical documentation judgment.

Background

The healthcare claim lifecycle and denial epidemiology

The healthcare claim lifecycle begins with clinical documentation and charge capture, continues through coding and claim creation, and culminates in payer adjudication, remittance advice, payment, denial, or request for correction. Denial categories may include eligibility failures, authorization defects, coding edits, medical necessity disputes, missing documentation, duplicate billing, or payer-specific coverage exclusions [1, 2]. Historical remittance data can provide denial labels and reason codes, but those labels must be interpreted within the timing and policy context of each payer. Because denial prevention depends on acting before submission, predictive models should be aligned with the point in the workflow where claim fields, documentation status, and authorization information are available.

Documentation completeness and claim adjudication

Documentation completeness affects claim adjudication because payers often require clinical evidence that the billed service was ordered, performed, medically necessary, and properly authenticated. Missing signatures, incomplete notes, insufficient procedure detail, weak medical necessity narratives, or inconsistent clinical context can undermine the claim even when the billing codes are technically present [9, 10]. Natural language processing and computer-assisted coding methods show how clinical text can be transformed into structured signals, but revenue cycle use requires those signals to reflect billing-relevant documentation requirements [7, 8, 17]. A documentation completeness feature should therefore represent the presence, specificity, and consistency of required documentation elements rather than simply the length of the note.

Procedure-diagnosis code consistency and payer edits

Procedure-diagnosis code consistency is central to medical necessity validation because payers often evaluate whether a CPT or HCPCS procedure is supported by the submitted ICD-10 diagnosis context. Automated clinical coding research demonstrates that diagnosis assignment is difficult, context-sensitive, and dependent on accurate extraction of clinical meaning from notes and structured data [7, 8, 17]. Coding inconsistency may also arise when procedure codes conflict with coverage logic, bundling expectations, or payer edits, producing denial risk even when documentation appears complete. A denial prediction framework should therefore encode both the raw procedure and diagnosis codes and derived consistency measures that approximate how the payer might evaluate their relationship.

Payer-specific rules and prior authorization requirements

Payer-specific rules create substantial variability in adjudication because two claims with similar clinical content may receive different outcomes depending on coverage policy, local or national determinations, plan design, and authorization rules. Prior authorization is especially important because a claim may be denied if approval was required but missing, expired, mismatched to the billed service, or not linked correctly to the submitted claim [11, 12]. Policy-aware artificial intelligence in healthcare must account for insurance coverage decisions and the operational consequences of payer-specific requirements [12]. For this reason, a denial model should treat payer rules and authorization history as structured risk signals rather than as external context.

Explainable machine learning in healthcare finance

Explainable machine learning is particularly relevant to healthcare finance because billing decisions must be auditable, actionable, and acceptable to non-technical stakeholders. Gradient boosting can model nonlinear interactions in tabular healthcare data, while SHAP, LIME, partial dependence, and related explanation methods can help users understand why a prediction was made [16-22]. Healthcare AI guidance emphasizes that trustworthiness depends not only on predictive validity but also on transparency, documentation, workflow fit, and evaluation of human use [15, 16, 23, 24]. In a revenue cycle setting, explainability should translate model output into specific review actions, such as verifying authorization, improving documentation, or checking a diagnosis-code pairing.

Model Overview and Design Principles

High-level predictive framework

At the moment of claim creation or pre-billing claim review, the proposed model would ingest structured and derived claim features and output a conceptual probability that the claim could be denied. The framework would combine administrative fields from claims, denial labels from remittance advice, documentation completeness indicators, procedure-diagnosis consistency features, payer rule flags, and prior authorization history [1-3]. A gradient-boosted tree ensemble would be appropriate because it can represent nonlinear interactions without requiring the model designer to specify every payer-specific interaction in advance [18, 19]. The prediction would be used as a triage signal for human review rather than as an autonomous billing decision.

Core input features

The core input features would include documentation completeness score, CPT or HCPCS procedure codes, ICD-10 diagnosis codes, diagnosis-procedure consistency metrics, payer identifiers, payer rule flags, prior authorization status, authorization age, and historical authorization outcomes. These features reflect the major denial pathways emphasized in revenue cycle analytics, clinical documentation integrity, automated coding, and prior authorization research [7-11]. The model could also include derived indicators for medical necessity support, missing documentation elements, code mismatch risk, and payer-specific coverage conflicts. By representing both raw billing fields and adjudication-oriented

derived features, the model would be expected to align more closely with actual denial mechanisms than a generic claim scrubber.

Design principles

The design should be interpretable, auditable, payer-aware, and compatible with low-latency pre-billing workflows. Explainable healthcare AI literature emphasizes that models must be designed around user trust, regulatory expectations, and meaningful evaluation of explanations, not simply around technical model sophistication [15, 23-27]. In this context, payer awareness means that the model should learn and expose differences in denial risk across payer policies, coverage requirements, and authorization practices. The model should also be updateable as payer rules, coding guidelines, authorization requirements, and organizational billing workflows evolve.

Figure 1 presents the proposed explainable gradient boosting framework for transforming claim, documentation, coding, payer-rule, and authorization data into interpretable denial-risk predictions for targeted pre-submission review.




Figure 1. Explainable gradient boosting framework for pre-submission health insurance claim denial prediction and human-governed revenue cycle intervention.

Data Sources and Feature Engineering
Claims and remittance data extraction

The proposed data foundation would include structured claim fields from electronic claim transactions and adjudication outcomes from remittance advice. Historical denial labels could be derived from payer response codes, denial reason categories, adjustment codes, and internal denial management classifications, while payment or clean-claim outcomes would provide comparison cases [1-3]. Because claims data are temporal, feature extraction should preserve the chronology between service date, claim creation, authorization activity, submission, adjudication, denial, appeal, and payment. This design would reduce the risk of using information that would not have been available at the pre-submission decision point.

Constructing documentation completeness and code consistency metrics

Documentation completeness could be represented using rule-based checks, natural language processing, or hybrid methods that detect required signatures, medical necessity statements, procedure details, order documentation, and clinical support for billed services. NLP and computer-assisted coding research shows that clinical text can be transformed into structured representations, but those representations should be adapted to billing and adjudication requirements rather than treated as general clinical summaries [7-10, 17]. Code consistency metrics could compare procedure codes with diagnosis codes, payer coverage logic, and medical necessity rules to identify potential mismatch patterns before submission. These derived features would make the model more actionable because they point to correctable claim components rather than opaque statistical associations.

Encoding payer rules and prior authorization history

Payer policies, coverage determinations, and prior authorization requirements could be encoded as structured feature flags indicating whether a service is covered, restricted, authorization-dependent, diagnosis-dependent, or subject to additional documentation review. Prior authorization history would include whether authorization was required, requested, approved, denied, expired, matched to the billed service, and linked to the claim record [11, 12]. The model could also represent elapsed time since authorization, prior denials for similar payer-service combinations, and discrepancies between approved and billed procedure details. This feature engineering approach would allow the model to learn both standalone authorization risk and interactions between authorization status, payer rules, documentation, and coding.

Table 1 maps the major conceptual denial pathways to model features, explanation outputs, and corrective revenue cycle actions, clarifying how prediction is translated into operational intervention.

Table 1. Conceptual Mapping of Denial Drivers to Model Features, Explanations, and Corrective Revenue Cycle Actions

Denial-risk domain	Model feature representation	Explanation output	Corrective action supported	Primary stakeholder
Documentation incompleteness	Missing signature indicators, absent order documentation, weak procedure detail, insufficient medical necessity narrative, documentation completeness score	Claim risk increased by incomplete or insufficient documentation elements	Request clarification, add missing documentation, verify required attachments	Clinical documentation integrity team
Diagnosis-procedure inconsistency	CPT/HCPCS–ICD-10 pairing metrics, medical necessity alignment score, mismatch flags	Risk increased by diagnosis code not supporting billed service	Review diagnosis selection, verify procedure indication, correct coding relationship	Coder / coding auditor
Payer-specific rule conflict	Payer rule flags, coverage restriction indicators, plan-specific medical necessity requirements	Risk increased by payer coverage or rule conflict	Review payer policy, verify coverage logic, escalate payer-specific edit	Payer-policy specialist
Prior authorization defect	Authorization required flag, approval status, expiration status, service match, claim-linkage indicator	Risk increased by missing, expired, unmatched, or incorrectly linked authorization	Verify authorization, correct linkage, update authorization documentation	Authorization specialist

Historical denial pattern	Prior denial rates by payer, procedure, diagnosis, service line, authorization category	Risk increased by historical denial pattern for similar claim profile	Prioritize for manual review before submission	Revenue cycle analyst
Model uncertainty or low confidence	Calibration output, probability confidence band, sparse-pattern indicator	Prediction should be interpreted cautiously because confidence is limited	Route to human review rather than automated action	Revenue cycle supervisor / compliance team

Explainable Gradient Boosting Architecture

Model structure and training regimen

The proposed architecture would use an XGBoost or LightGBM-style classifier trained conceptually on historical claims and remittance-derived denial labels. Gradient-boosted tree ensembles are suitable for mixed structured features and can capture nonlinear interactions among payer, procedure, diagnosis, documentation, and authorization signals [18, 19]. A temporal training design would be preferred so that claims from earlier periods inform predictions for later claims, reducing the risk of leakage from future adjudication behavior into pre-submission predictions. Because denials are often a minority outcome, the model should be evaluated and calibrated in ways that reflect the operational cost of both missed denials and unnecessary manual review.

Handling categorical and sparse features

Revenue cycle data contain many categorical and sparse features, including payer identifiers, plan types, rendering providers, service locations, CPT or HCPCS codes, ICD-10 codes, denial reason codes, and authorization categories. Machine learning with electronic health and administrative data often requires careful representation of high-cardinality variables so that rare but important patterns are not lost [13, 14, 28, 29]. Target encoding, hierarchical grouping, payer-specific aggregation, or learned representations could be considered, provided that they are designed to avoid leakage and remain explainable to revenue cycle users. For billing operations, feature representation should preserve enough semantic meaning that explanations can be translated into practical review steps.

Output: denial probability and confidence

The model output would be a calibrated denial probability accompanied by an uncertainty or confidence indication suitable for workflow triage. In healthcare AI, calibrated and interpretable outputs are important because users need to understand not only the direction of risk but also whether a prediction should trigger operational intervention [15, 16, 24, 25]. The probability should be presented with explanatory context, such as whether risk is driven mainly by authorization, documentation, coding, or payer rule features. Rather than functioning as a final adjudication decision, the score would support pre-billing prioritization, escalation, and documentation of why a claim was selected for review.

Incorporating Payer Rules and Documentation Completeness

Payer rule library construction

A payer rule library would translate payer policies, coverage determinations, prior authorization requirements, medical necessity criteria, and documentation expectations into structured indicators that can be consumed by the denial prediction model. These indicators could identify whether a procedure is subject to authorization, whether diagnosis support is required, whether documentation attachments are expected, and whether local or national coverage rules

appear relevant to the claim [11, 12]. Because payer policies evolve over time, the library should be versioned so that predictions are tied to the rule environment that existed when the claim was created. This design would make payer-rule features auditable and would allow revenue cycle teams to update the model environment without treating policy changes as invisible statistical drift.

Documentation completeness scoring

Documentation completeness scoring would estimate whether the clinical record contains the evidence needed to support the billed service. The score could be generated through structured-data checks, natural language processing, or hybrid review logic that detects required signatures, medical necessity narratives, order details, procedure descriptions, and relevant clinical findings [7-10, 17]. Rather than replacing clinical documentation integrity review, the score would function as a pre-billing signal that identifies claims needing targeted documentation attention. Explainability is important here because revenue cycle staff need to know which documentation elements appear incomplete, not merely that the claim is statistically high risk [20-22].

Code consistency and medical necessity checks

Code consistency features would represent the relationship between CPT or HCPCS procedure codes, ICD-10 diagnosis codes, payer medical necessity requirements, and rule-based billing edits. Automated coding research shows that diagnosis and procedure coding depend on accurate extraction of clinical meaning from the record, while revenue cycle models must also account for payer adjudication logic [7, 8, 17]. A rule-based validation layer could therefore generate indicators for diagnosis-procedure mismatch, insufficient medical necessity support, bundling conflict, or payer-specific coverage inconsistency. These indicators would be expected to improve actionability because they connect model risk to correctable coding or documentation review tasks [1, 2].

Prior authorization as a modifier and standalone risk

Prior authorization should be treated both as a standalone denial-risk factor and as a modifier of other claim features. A missing, expired, unmatched, or incorrectly linked authorization could make a claim high risk even when coding and documentation are otherwise strong, while an approved authorization may not fully protect a claim if the billed service, diagnosis, or documentation conflicts with payer requirements [11, 12]. The model should therefore learn interactions among authorization status, payer policy, procedure category, diagnosis support, and documentation completeness. In practice, the explanation layer should distinguish between claims that need authorization correction and claims that remain risky despite authorization because of coding or medical necessity concerns [18, 19].

Explainability Methods for Denial Predictions

Global feature importance for denial risk management

Global explainability would help leaders understand which factors most commonly drive denial risk across the payer mix, service lines, and claim categories. SHAP summary analyses could show whether features such as expired prior authorization, diagnosis-procedure inconsistency, payer-specific medical necessity conflict, or missing documentation elements are major contributors to predicted denials [18-20]. These insights would support operational prioritization, such as strengthening documentation workflows, revising coding review rules, or improving authorization tracking. Because healthcare AI must be accountable and interpretable, global explanations should be reviewed as management intelligence rather than as definitive causal proof [15, 23-25].

Local explanation for the revenue cycle specialist

Local explanations would translate each claim's denial prediction into a case-specific rationale for the revenue cycle specialist. For example, a SHAP waterfall-style explanation could indicate that incomplete documentation, a diagnosis

code not aligned with payer coverage logic, and an authorization mismatch are the main reasons a claim is being flagged [18, 20, 21]. This local explanation would be more useful than a generic risk score because it directs the specialist toward the fields, documents, or policy requirements that should be reviewed before submission. In a healthcare finance workflow, the goal of local explainability is not simply transparency but timely correction of preventable denial drivers [1, 2, 16].

Counterfactual explanations for correction guidance

Counterfactual explanations would describe how denial risk could change conceptually if specific claim defects were corrected. For instance, the system could indicate that attaching a valid prior authorization, correcting a diagnosis-procedure mismatch, or adding missing medical necessity documentation would be expected to reduce denial risk, without presenting the statement as a guaranteed payer outcome [9-12]. Counterfactual reasoning is particularly valuable in revenue cycle operations because users need guidance on what action could make the claim more defensible before submission [20-22]. The model should present such explanations as decision support, since payer adjudication remains policy-dependent and may involve information not captured by the model.

Audit trail and payer dispute support

The explanation layer should create an audit trail that records the model version, feature state, payer rule version, denial-risk rationale, and any recommended pre-billing correction at the time of review. If the claim is later denied, this record could support internal root-cause analysis and help organize appeal documentation by showing which risks were detected and which corrective steps were taken [1, 2]. Explainable AI guidance emphasizes that auditability and documentation are central to trustworthy healthcare AI, particularly when model output influences operational decisions [23-26]. In payer disputes, the explanation should not be framed as evidence that the payer must pay the claim, but as structured support for why the provider believed the claim met documentation, coding, and authorization requirements.

Revenue Cycle Integration and Pre-Billing Intervention

Embedding in claim scrubbing and RCM workflow

The model would be embedded within the pre-billing claim edit process, where it could run after charge capture and coding but before claim submission. Claims with elevated predicted denial risk would be routed to appropriate work queues, such as coding review, authorization verification, clinical documentation integrity, or payer-policy review [1-3]. Because operational AI in healthcare must fit existing workflows, the interface should provide concise explanations, responsible work queues, and recommended review actions rather than technical model artifacts alone [15, 16]. This integration would allow the model to complement deterministic claim scrubbers by identifying interaction-based risk that static rules may not capture.

Closed-loop learning from denials and appeals

A closed-loop learning process would feed adjudication outcomes, denial reasons, appeal decisions, and post-review corrections back into the model governance cycle. When payer behavior changes, denial patterns shift, or authorization requirements are updated, the system should detect that previously reliable risk patterns may no longer apply [11, 12, 24]. This feedback loop would allow revenue cycle teams to refine payer rule libraries, documentation checks, coding consistency features, and model monitoring practices over time. The process should remain governed by human review because machine learning models can amplify data-quality problems or institutional bias if denial labels are accepted uncritically [15, 16, 23].

Evaluation Strategy

Predictive performance metrics

The proposed model should be evaluated with performance metrics appropriate for imbalanced denial prediction and pre-billing triage, including discrimination, precision-recall behavior, calibration, and decision-analytic usefulness. No single metric would be sufficient because revenue cycle teams must balance the cost of missed preventable denials against the burden of unnecessary manual review [1-3]. Calibration would be especially important because a denial probability should correspond to meaningful operational risk rather than merely ranking claims in relative order [15, 16, 24]. Evaluation should also compare the model conceptually against existing claim scrubbing and rule-based review workflows without treating predictive performance as the only criterion for adoption.

Explanation quality and user acceptance

Explanation quality should be evaluated by whether coders, billers, authorization specialists, and documentation integrity staff understand the explanation and can translate it into corrective action. Healthcare explainability literature emphasizes that explanation methods must be assessed in relation to user needs, decision context, trust, and potential misuse [20-22, 25-27]. In this setting, useful explanations would identify actionable denial drivers such as missing authorization evidence, documentation insufficiency, or diagnosis-code inconsistency rather than merely listing abstract feature names. User acceptance should therefore be evaluated through workflow-centered assessment, including whether explanations support appropriate review and avoid overreliance on model output [15, 16, 23].

Financial and operational impact

Financial and operational evaluation should examine whether the model would be expected to reduce preventable denials, accelerate correction before submission, improve work-queue prioritization, and reduce avoidable rework in accounts receivable. The assessment should consider denial prevention, appeal preparation, days in accounts receivable, write-offs, staff workload, and comparison with existing claim-scrubbing processes [1-3]. However, these outcomes should be interpreted cautiously because changes in payer policy, coding practices, documentation workflows, and patient mix can affect observed impact. A responsible evaluation strategy would therefore combine quantitative monitoring with qualitative review by revenue cycle, compliance, coding, and clinical documentation stakeholders [15, 16, 24].

Table 2 provides a multidimensional evaluation framework for assessing whether the proposed model is predictive, calibrated, explainable, operationally useful, and governable in revenue cycle practice.

Table 2. Evaluation Framework for an Explainable Claim Denial Prediction Model

Evaluation dimension	Key question	Suggested assessment approach	Interpretation for adoption
Predictive discrimination	Can the model distinguish claims likely to be denied from clean claims?	AUROC, precision-recall analysis, payer- and service-line stratified performance	Supports triage only if high-risk claims are reliably separated from low-risk claims
Calibration	Does the predicted denial probability reflect actual observed denial risk?	Calibration plots, Brier score, calibration by payer and denial category	Necessary for using the score as an operational risk estimate
Operational usefulness	Does the model reduce unnecessary review while capturing preventable denials?	Work-queue simulation, threshold analysis, decision-curve analysis	Determines whether model use improves pre-billing prioritization

Explanation actionability	Do explanations identify correctable denial drivers?	User review by coders, billers, authorization staff, and documentation teams	Explanations are valuable only if they guide concrete correction
Payer-rule robustness	Does performance remain stable across payer policies and rule changes?	Temporal validation, rule-version stratification, drift monitoring	Required because payer behavior and authorization rules change over time
Equity and consistency	Are some services, specialties, or patient groups disproportionately flagged without justification?	Subgroup analysis by payer, service line, specialty, claim type, and patient category	Reduces risk of biased workload allocation or unfair denial-prevention practices
Governance and auditability	Can decisions be reconstructed after review, denial, or appeal?	Model-version logs, payer-rule-version logs, explanation archive, reviewer action record	Supports compliance, internal audit, and appeal preparation

Limitations

Data quality and payer rule variability

The proposed framework would depend heavily on the quality, completeness, and timeliness of claims, remittance, authorization, documentation, and payer-policy data. Payer rules may change frequently, may be inconsistently documented, and may not always be reducible to clean structured features [11, 12]. Documentation completeness is also context-dependent because different procedures, specialties, and payers may require different kinds of evidence [9, 10]. These limitations mean that model explanations should be interpreted as structured decision support, not as definitive statements about payer adjudication [15, 24, 25].

Generalizability across providers and payer mixes

A model trained on one health system's payer mix, specialty distribution, documentation practices, coding workflows, and denial management processes may not transfer directly to another institution. Healthcare machine learning research repeatedly emphasizes that models can be affected by local data-generating processes, workflow differences, and population or institutional variation [13-15, 28, 29]. In revenue cycle settings, payer contracts, regional coverage practices, authorization processes, and internal coding policies can further limit generalizability [11, 12]. Multi-site evaluation and local calibration would therefore be necessary before the framework could be considered broadly reliable across providers.

Conclusion

An explainable gradient boosting model for claim denial prediction could provide a structured way to identify denial risk before claim submission. By combining documentation completeness, procedure codes, diagnosis-code consistency, payer rules, and prior authorization history, the model would represent the claim as an integrated adjudication object rather than a set of isolated billing fields.

The main strength of the proposed approach is that it links prediction with explanation. A denial risk score alone would have limited value, but a risk score paired with claim-level reasoning could guide coders, billers, authorization specialists, and documentation integrity teams toward targeted corrective action.

Important challenges would remain, including changing payer policies, inconsistent documentation quality, variation across specialties, and limited transferability across institutions. These issues reinforce the need for model governance, local validation, human review, and careful monitoring of unintended operational consequences.

Future work should pursue multi-site validation, shared denial-prediction benchmarks, and collaboration between providers, payers, informatics teams, and revenue cycle leaders. Such efforts could help move denial management from reactive appeals toward transparent, proactive, and evidence-supported prevention.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 03 Apr 2023

Revised: 21 Apr 2023

Accepted: 20 Jun 2023

Published online: 20 July 2023

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- Johnson M, Albizri A, Harfouche A. Responsible artificial intelligence in healthcare: predicting and preventing insurance claim denials for economic and social wellbeing. *Inf Syst Front.* 2023;25(6):2179-95.
<https://doi.org/10.1007/s10796-022-10311-4>.
- Pal S, Gaur M, Chaudhuri R, Kalaivanan R, Chetan KV, Praneeth BH, et al. Driving impact in claims denial management using artificial intelligence. In: Singh PK, Pawłowski W, Tanwar S, Sur A, Sharma R, editors. *Advances in Computing and Data Sciences*. Cham: Springer International Publishing; 2022. p. 107-20.
https://doi.org/10.1007/978-3-031-16364-7_9.
- Chaudhuri R, Parsa SP, Nagpal D, R K. Time to response prediction for following up on account receivables in healthcare revenue cycle management. In: Singh PK, Pawłowski W, Tanwar S, Sur A, Sharma R, editors. *Advances in Computing and Data Sciences*. Cham: Springer International Publishing; 2022. p. 124-37.
https://doi.org/10.1007/978-3-031-16364-7_11.

Permai SD, Herdianto K. Prediction of health insurance claims using logistic regression and XGBoost methods. *Procedia Comput Sci.* 2023;227:1012-9.
<https://doi.org/10.1016/j.procs.2023.10.598>.

Kaushik K, Bhardwaj A, Dwivedi AD, Singh R. Machine learning-based regression framework to predict health insurance premiums. *Int J Environ Res Public Health.* 2022;19(13):7898.
<https://doi.org/10.3390/ijerph19137898>.

Drewe-Boss P, Enders D, Walker J, Ohler U. Deep learning for prediction of population health costs. *BMC Med Inform Decis Mak.* 2022;22(1):32.
<https://doi.org/10.1186/s12911-022-01748-z>.

Dong H, Falis M, Whiteley W, Alex B, Matterson J, Ji S, et al. Automated clinical coding: what, why, and where we are? *NPJ Digit Med.* 2022;5(1):159.
<https://doi.org/10.1038/s41746-022-00691-2>.

Diao X, Huo Y, Zhao S, Yuan J, Cui M, Wang Y, et al. Automated ICD coding for primary diagnosis via clinically interpretable machine learning. *Int J Med Inform.* 2021;153:104543.
<https://doi.org/10.1016/j.ijmedinf.2021.104543>.

Campbell S, Giadresco K. Computer-assisted clinical coding: a narrative review of the literature on its benefits, limitations, implementation and impact on clinical coding professionals. *Health Inf Manag J.* 2020;49(1):5-18.
<https://doi.org/10.1177/1833358319851585>.

Pine KH, Landon LA, Bossen C, VanGelder ME. Innovations in clinical documentation integrity practice: continual adaptation in a data-intensive healthcare organisation. *Health Inf Manag J.* 2023;52(2):119-24.
<https://doi.org/10.1177/18333583221074707>.

Lenert LA, Lane S, Wehbe R. Could an artificial intelligence approach to prior authorization be more human? *J Am Med Inform Assoc.* 2023;30(5):989-94.

Park SH, Choi J, Byeon JS. Key principles of clinical validation, device approval, and insurance coverage decisions of artificial intelligence. *Korean J Radiol.* 2021;22(3):442-53.
<https://doi.org/10.3348/kjr.2020.0981>.

Beam AL, Kohane IS. Big data and machine learning in health care. *JAMA.* 2018;319(13):1317-8.
<https://doi.org/10.1001/jama.2017.18391>.

Yu KH, Beam AL, Kohane IS. Artificial intelligence in healthcare. *Nat Biomed Eng.* 2018;2(10):719-31.
<https://doi.org/10.1038/s41551-018-0305-z>.

Wiens J, Saria S, Sendak M, Ghassemi M, Liu VX, Doshi-Velez F, et al. Do no harm: a roadmap for responsible machine learning for health care. *Nat Med.* 2019;25(9):1337-40.
<https://doi.org/10.1038/s41591-019-0548-6>.

Ghassemi M, Naumann T, Schulam P, Beam AL, Chen IY, Ranganath R. A review of challenges and opportunities in machine learning for health. *AMIA Summits Transl Sci Proc.* 2020;2020:191-200.

Nguyen AN, Truran D, Kemp M, Koopman B, Conlan D, O'Dwyer J, et al. Computer-assisted diagnostic coding: effectiveness of an NLP-based approach using SNOMED CT to ICD-10 mappings. *AMIA Annu Symp Proc.* 2018;2018:807-16.

Lundberg SM, Erion G, Chen H, DeGrave A, Prutkin JM, Nair B, et al. From local explanations to global understanding with explainable AI for trees. *Nat Mach Intell.* 2020;2(1):56-67.
<https://doi.org/10.1038/s42256-019-0138-9>.

Lundberg SM, Nair B, Vavilala MS, Horibe M, Eisses MJ, Adams T, et al. Explainable machine-learning predictions for the prevention of hypoxaemia during surgery. *Nat Biomed Eng.* 2018;2(10):749-60.
<https://doi.org/10.1038/s41551-018-0304-0>.

Adadi A, Berrada M. Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). *IEEE Access.* 2018;6:52138-60.
<https://doi.org/10.1109/ACCESS.2018.2870052>.

Guidotti R, Monreale A, Ruggieri S, Turini F, Giannotti F, Pedreschi D. A survey of methods for explaining black box models. *ACM Comput Surv.* 2018;51(5):93.
<https://doi.org/10.1145/3236009>.

Arrieta AB, Díaz-Rodríguez N, Del Ser J, Benetot A, Tabik S, Barbado A, et al. Explainable Artificial Intelligence (XAI): concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf Fusion.* 2020;58:82-115.
<https://doi.org/10.1016/j.inffus.2019.12.012>.

Liu X, Rivera SC, Moher D, Calvert MJ, Denniston AK, Ashrafian H, et al. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. *Lancet Digit Health.* 2020;2(10):e537-e548.
[https://doi.org/10.1016/S2589-7500\(20\)30218-1](https://doi.org/10.1016/S2589-7500(20)30218-1).

Norgeot B, Quer G, Beaulieu-Jones BK, Torkamani A, Dias R, Gianfrancesco M, et al. Minimum information about clinical artificial intelligence modeling: the MI-CLAIM checklist. *Nat Med.* 2020;26(9):1320-4.
<https://doi.org/10.1038/s41591-020-1041-y>.

Markus AF, Kors JA, Rijnbeek PR. The role of explainability in creating trustworthy artificial intelligence for health care: a comprehensive survey of the terminology, design choices, and evaluation strategies. *J Biomed Inform.* 2021;113:103655.
<https://doi.org/10.1016/j.jbi.2020.103655>.

Tjoa E, Guan C. A survey on explainable artificial intelligence (XAI): toward medical XAI. *IEEE Trans Neural Netw Learn Syst.* 2021;32(11):4793-813.
<https://doi.org/10.1109/TNNLS.2020.3027314>.

Vellido A. The importance of interpretability and visualization in machine learning for applications in medicine and health care. *Neural Comput Appl.* 2020;32(24):18069-83.
<https://doi.org/10.1007/s00521-019-04051-w>.

Rajkomar A, Oren E, Chen K, Dai AM, Hajaj N, Hardt M, et al. Scalable and accurate deep learning with electronic health records. *NPJ Digit Med.* 2018;1(1):18.
<https://doi.org/10.1038/s41746-018-0029-1>.

Shickel B, Tighe PJ, Bihorac A, Rashidi P. Deep EHR: a survey of recent advances in deep learning techniques for electronic health record (EHR) analysis. *IEEE J Biomed Health Inform.* 2018;22(5):1589-604.
<https://doi.org/10.1109/JBHI.2017.2767063>.