

ORIGINAL RESEARCH

Open access

Interpretable Machine Learning Model for Predicting Prior Authorization Approval Delays Using Payer-Specific Rules, Clinical Documentation Features, Procedure Type, Medical Necessity Indicators, and Historical Approval Timelines

Nguyen Van Nam¹, Tran Thi Hoa^{1*}, Le Minh Duc²

Abstract

Prior authorization delays impede timely patient care and contribute to administrative pressure across clinical and revenue cycle workflows. These delays can affect scheduling, medication access, procedural planning, and patient confidence in the care process. Current authorization management tools are largely reactive and often focus on tracking request status after submission. They rarely predict which requests are likely to experience approval delays or explain the operational, clinical, or payer-specific reasons behind those delays. This article proposes an interpretable machine learning model for predicting the likelihood of prior authorization approval delays. The model is designed to provide transparent, request-level explanations that can guide pre-submission correction and authorization preparation. The proposed framework uses a gradient-boosted tree model trained conceptually on historical authorization requests. Inputs include payer-specific rules, clinical documentation features, procedure type, medical necessity indicators, and historical approval timelines, with SHAP used to attribute predicted delay risk to individual features. Conceptually, the model would output a delay probability and an explanation of the dominant drivers of that prediction. These drivers could include incomplete documentation, mismatch with payer medical necessity criteria, procedure categories associated with additional review, or payer-procedure combinations with historically slow turnaround. An interpretable prior authorization delay model could support earlier correction of incomplete requests, reduce administrative waste, and improve patient access. By aligning predictive analytics with transparent explanations, the framework could make authorization preparation more proactive and accountable.

Keywords Explainable artificial intelligence, SHAP, Revenue cycle analytics, Prior authorization, Payer rules, Medical necessity

*Correspondence:

Tran Thi Hoa
tran.hoa@outlook.com

¹ Department of Intelligent Healthcare Informatics, Faculty of Medicine, Hanoi Medical University, Hanoi, Vietnam

² Department of Clinical Data Engineering, Faculty of Engineering, Ho Chi Minh City University of Technology, Ho Chi Minh City, Vietnam

Introduction

Prior authorization has become a major source of administrative burden in healthcare because it requires clinicians, staff, patients, and payers to coordinate approval before selected medications, procedures, services, or transports can proceed. Studies of authorization requirements across medications, oncology care, ambulance transport, and specialty services show that delays can interrupt care pathways, create uncertainty for patients, and increase workload for clinical and administrative teams [1-6]. The burden is not limited to a single specialty or service category, because authorization rules affect drugs, imaging, procedures, outpatient therapy, and high-cost services in different ways [7-9]. These pressures make prior authorization not only a payer-management mechanism but also a practical operational problem for health systems seeking timely access, predictable scheduling, and reliable reimbursement [10, 11].

Current authorization workflows often depend on manual document gathering, payer portal submission, phone calls, fax exchanges, and staff follow-up after the request has already entered the payer review process. Provider surveys and workflow analyses indicate that authorization software can reduce some friction, but many tools remain focused on electronic submission and tracking rather than pre-submission prediction of delay risk [1, 12]. In practice, authorization specialists may know that certain payers, procedures, or missing documents increase delay risk, yet this knowledge is often tacit, distributed across staff experience, and not consistently encoded into decision support [10, 13]. As a result, requests may be submitted with preventable documentation gaps or insufficient medical necessity evidence, leading to avoidable rework and delayed approval [3, 14].

Historical authorization data, payer-specific policy requirements, procedure codes, clinical documentation patterns, and prior turnaround timelines provide a foundation for predictive modeling. Prior studies on prior authorization for surgery, medications, oncology care, HIV prevention, and specialty drugs show that authorization outcomes vary by payer, product, service category, and policy design, suggesting that structured features could capture meaningful adjudication patterns [7, 8, 15-17]. Machine learning and artificial intelligence have also been proposed for related revenue cycle tasks, including claim denial prediction and administrative decision support, demonstrating the feasibility of using operational, clinical, and billing data to anticipate downstream financial or access barriers [18-20]. However, because authorization decisions affect access to care, predictive models should not function as opaque black boxes, and explainability is essential for staff trust, auditability, and corrective action [21-23].

This article proposes an interpretable machine learning framework that could predict prior authorization approval delays before submission and explain the factors driving each prediction. The model would combine payer-specific rule flags, clinical documentation completeness, procedure type, medical necessity indicators, and historical approval timelines into a structured prediction pipeline. Rather than replacing authorization specialists or clinicians, the framework would support their preparation work by identifying requests likely to be delayed and by indicating why the risk is elevated. The central thesis is that explainable delay prediction can transform authorization management from reactive status tracking into proactive, evidence-guided request preparation while preserving human oversight and accountability.

Background

Prior authorization in healthcare

Prior authorization is a utilization management process in which payers require approval before covering selected services, medications, procedures, transports, or therapies. The process involves clinicians who determine clinical need, administrative staff who assemble and submit documentation, payers who adjudicate coverage, and patients whose care may depend on timely approval. Research across medication authorization, oncology care, outpatient antimicrobial therapy, nonemergent ambulance transport, and specialty procedures shows that delays can arise from missing documents, policy variation, payer review timelines, and the administrative burden of repeated communication [2, 4-6, 24]. These drivers make authorization delay prediction a multi-stakeholder problem that depends on clinical evidence, payer policy logic, procedure type, and historical payer behavior [1, 10].

Payer-specific rules and medical necessity determination

Payer-specific medical policies define which diagnoses, prior therapies, procedure indications, medication criteria, and supporting documents are needed before a request is approved. Studies of PCSK9 inhibitor authorization, HIV pre-exposure prophylaxis coverage, buprenorphine policy requirements, and health plan variation show that authorization rules differ across payers, regions, and service categories [7, 8, 16, 17, 25]. Medical necessity determination therefore reflects both clinical evidence and payer-defined coverage logic, which may not be transparent to patients or clinicians at the point of care [3, 14]. For prediction, these rules can be represented as structured features that flag whether the submitted request aligns with payer-specific criteria or contains missing elements likely to slow adjudication [15, 26].

Clinical documentation features and procedure type

Clinical documentation features are central to authorization because payers often require evidence of diagnosis, symptom severity, prior treatment, conservative management, functional limitation, or guideline-based indication. Procedure type further shapes review complexity because high-cost drugs, advanced imaging, surgery, radiation oncology, antimicrobial therapy, and specialty services may trigger different authorization pathways and evidence requirements [4, 6, 7, 24]. Prior work on lumbar stenosis surgery authorization and specialty medication approval suggests that clinical indications, service category, and payer requirements can jointly influence whether a request proceeds smoothly or encounters delay [15, 16]. An interpretable model should therefore treat documentation completeness and procedure type as interacting signals rather than isolated administrative fields [20, 27].

Historical approval timelines and payer behavior

Historical approval timelines can reveal payer-procedure combinations that repeatedly experience long review cycles, additional documentation requests, or inconsistent approval patterns. Evidence from authorization studies in specialty drugs, pre-exposure prophylaxis, buprenorphine, oncology care, and ambulance transport indicates that payer policies and administrative practices can create systematic differences in access and turnaround [5, 7-9, 17]. These patterns can be transformed into rolling historical features, such as prior approval speed, prior denial tendency, or recent policy-related delay signals for a given payer and procedure category. Such features would be expected to help a model distinguish a request delayed because of incomplete documentation from one delayed because of historically slow payer response for a particular service [10, 16].

Interpretable machine learning for revenue cycle

Interpretable machine learning is particularly important in revenue cycle applications because predictions may influence staff prioritization, documentation requests, appeal preparation, and patient access workflows. Artificial intelligence has been applied conceptually and operationally to claim denial management, health insurance claim deduction reduction, and broader administrative decision support, indicating that revenue cycle data can support predictive analytics when models are properly governed [18-20]. SHAP-based explanation methods can attribute tree-based model predictions to individual features, while arguments for inherently interpretable models emphasize that high-stakes healthcare decisions require transparency rather than unexplained black-box outputs [21, 22]. In the prior authorization context, explainability should clarify whether the predicted delay is driven by payer rules, documentation gaps, procedure category, medical necessity alignment, or historical payer behavior [23, 26].

Model Development Overview

High-level prediction framework

When a prior authorization request is initiated, the proposed framework would ingest structured data from the authorization request, payer policy library, EHR documentation, claims or procedure coding systems, and historical authorization logs. The model would output a conceptual probability that approval will be delayed beyond a locally defined

operational threshold, along with a ranked explanation of feature contributions. This design is consistent with prior work suggesting that machine learning could support prior authorization decision processes and that artificial intelligence must remain understandable when used in coverage-related workflows [15, 26–28]. The prediction would be generated before payer submission so that staff could correct missing evidence, attach medical necessity documentation, or prepare escalation materials when the delay risk is high [1, 13].

Figure 1 illustrates the proposed explainable prior authorization delay-prediction workflow, showing how payer rules, clinical documentation, procedure type, medical necessity indicators, and historical timelines are transformed into interpretable delay-risk outputs and pre-submission corrective actions.




Figure 1. Explainable Prior Authorization Delay-Prediction Workflow from Payer Rules and Clinical Documentation to Pre-Submission Correction and Human Oversight

Core input features

The core input feature set would include payer identity, payer policy rule flags, clinical documentation completeness, procedure or drug code category, medical necessity indicators, and historical approval timelines for similar payer-procedure combinations. Payer identity and policy rule features would capture coverage variability, while documentation features would represent whether required clinical evidence is present in notes, orders, referrals, and supporting records [7, 16, 25]. Procedure type features would distinguish services with different authorization burden, such as surgery, imaging, oncology care, specialty medication, outpatient antimicrobial therapy, or transport [4–6, 15, 24]. Historical timeline features would summarize prior turnaround behavior and could help identify payer-service combinations with predictable administrative delay [8, 10, 17].

Design principles

The framework should be designed for real-time use during authorization preparation rather than retrospective review after a delay has already occurred. It should produce explanations that authorization specialists, clinicians, and revenue cycle leaders can understand, including whether the predicted delay arises from missing documentation, weak medical necessity alignment, procedure complexity, or payer history [21, 26]. It should also support auditability by logging predictions, explanations, staff actions, and payer outcomes so that health systems can evaluate whether the tool improves preparation without reinforcing inequitable or opaque payer behavior [22, 23, 28]. The model's purpose is therefore operational guidance, not autonomous authorization decision-making or replacement of payer, clinician, or staff judgment [3, 14].

Data Sources and Feature Engineering

Extraction from prior authorization systems and payer portals

The primary data source would be historical prior authorization records, including request date, submission channel, payer, procedure or medication code, requested service, payer response, approval or denial status, resubmission events, and turnaround time. Payer portals and authorization software can provide structured response fields and timestamps, although workflow studies indicate that electronic tools may still coexist with manual follow-up and fragmented communication [1, 13]. These authorization records would be linked to EHR documentation, scheduling information, claims or billing data, and service-line identifiers to construct a complete operational view of the request. Because administrative cost and staff time are substantial in insurance-related activities, careful data linkage could help identify where delay prediction might reduce rework and improve workflow efficiency [11].

Constructing payer-specific rule features and medical necessity indicators

Payer-specific rule features would translate coverage policies into structured indicators that describe whether required diagnoses, prior treatments, severity documentation, imaging findings, medication histories, or letters of medical necessity are present. Medical necessity indicators could be derived from structured diagnosis and procedure codes, referral details, order metadata, and natural language processing of clinical notes, with human review used for governance where documentation meaning is ambiguous. Studies of medication and service-specific authorization requirements show that policy criteria can vary substantially across payers and regions, supporting the need for dynamic rule encoding rather than a single universal authorization checklist [7, 8, 17, 25]. These features would allow the model to distinguish requests that appear clinically complete from those likely to trigger payer requests for additional information [15, 16].

Procedure type and historical timeline features

Procedure type features would categorize requests by service family, such as surgical procedures, diagnostic imaging, infusion therapy, specialty drugs, radiation oncology services, transport, or home-based care support. Historical timeline features would summarize recent payer response patterns for similar request categories, including whether a payer-procedure pair has tended to require prolonged review, additional documentation, or repeated clarification. Prior authorization studies across spine surgery, oncology, outpatient antimicrobial therapy, specialty pharmacy, and ambulance transport suggest that service type and payer practices can jointly shape approval speed and administrative burden [4-6, 15, 24]. These engineered features would make the model sensitive not only to the current request’s clinical content but also to the operational history of comparable requests [10, 16].

Table 1 maps each predictor domain to its corresponding feature-engineering strategy, SHAP-style explanation, and practical pre-submission action for authorization staff.

Table 1. Predictor-to-Explanation Mapping for an Explainable Prior Authorization Delay-Prediction Model

Predictor domain	Example data elements	Feature-engineering approach	Expected contribution to delay prediction	SHAP-style explanation visible to users	Practical pre-submission action
Payer-specific rule alignment	Payer ID, plan type, coverage policy, site-of-care requirement, prior therapy requirement, required attachments	Encode payer policy requirements as binary, ordinal, or categorical rule-alignment flags	Identifies whether the request matches known payer coverage expectations or is likely to trigger additional review	“Delay risk increased because required payer-specific criteria are not fully documented.”	Complete payer-specific checklist before submission and attach required evidence.

Clinical documentation completeness	Progress notes, referral notes, operative planning notes, medication history, symptom documentation, conservative management evidence	Score presence, completeness, and location of required clinical elements using structured checks and NLP-assisted review	Detects whether the clinical record contains enough evidence for payer adjudication	“Delay risk increased because clinical documentation does not clearly show prior treatment history or severity.”	Ask clinician or documentation team to add missing clinical justification.
Procedure type and service category	CPT codes, HCPCS codes, drug codes, imaging category, surgical category, infusion therapy, transport, oncology service	Group procedures into authorization-relevant categories and encode expected review complexity	Captures inherent authorization burden associated with high-cost, complex, or policy-sensitive services	“Delay risk increased because this service category commonly requires additional payer review.”	Prepare procedure-specific documentation bundle and confirm service-line requirements.
Medical necessity indicators	Diagnosis codes, diagnosis-to-procedure match, disease severity, guideline-based criteria, failed prior therapy, functional limitation	Construct structured necessity indicators and composite documentation-support scores	Assesses whether the request is clinically aligned with payer-defined necessity criteria	“Delay risk increased because medical necessity evidence is incomplete or weakly aligned with the requested service.”	Add diagnosis rationale, severity evidence, prior therapy documentation, or letter of medical necessity.
Historical approval timeline	Prior requests by payer, procedure category, service line, submission channel, response date, approval date, denial date, additional information requests	Create rolling historical features for payer-procedure turnaround, delay frequency, and request-for-information patterns	Captures payer behavior and recent operational patterns that may affect approval speed	“Delay risk increased because similar requests to this payer have historically taken longer to approve.”	Submit earlier, prioritize staff follow-up, prepare escalation documentation, and warn scheduling teams.
Submission-channel and workflow context	Portal submission, fax, phone follow-up, staff queue status, missing portal fields, resubmission history	Encode submission pathway and operational friction points as workflow features	Identifies administrative pathways associated with avoidable delay or incomplete status capture	“Delay risk increased because the request pathway has historically required manual follow-up.”	Route request to the most reliable submission channel and assign follow-up responsibility.

Missingness and data-quality indicators	Missing payer response date, incomplete documentation fields, absent prior authorization history, scanned documents unavailable for extraction	Treat missingness as an interpretable feature rather than only an imputation problem	Distinguishes true clinical absence from administrative data gaps that may impair review	"Delay risk increased because required information could not be verified in the available record."	Verify source documents manually and correct data capture before payer submission.
---	--	--	--	--	--

Interpretable Machine Learning Architecture

Model choice and rationale

A gradient-boosted tree model would be appropriate for this conceptual framework because authorization data are likely to include mixed categorical, ordinal, binary, and continuous features with missingness and nonlinear interactions. Such a model could represent interactions among payer, procedure type, documentation completeness, medical necessity alignment, and historical payer behavior without requiring a fully manual rule hierarchy. Tree-based models are also compatible with SHAP explanations, which can translate complex model behavior into request-level and institution-level feature attribution [21]. Because prior authorization predictions may affect care access and administrative prioritization, the model should be paired with interpretability safeguards and compared with simpler interpretable alternatives such as logistic regression where appropriate [22, 23].

Input feature vector and preprocessing

The model input vector would encode payer identity, procedure or drug category, policy rule alignment, documentation completeness, medical necessity support, prior payer turnaround behavior, and operational context at the time of request preparation. Categorical fields such as payer and procedure code could be encoded using transparent categorical strategies, while continuous measures such as documentation completeness or historical turnaround summaries would be normalized or grouped into clinically interpretable ranges. Missing data would be treated as potentially informative, because absence of documentation, unavailable payer history, or incomplete portal response fields may itself signal authorization risk [1, 12]. Preprocessing should therefore preserve the distinction between clinically absent evidence, administratively missing data, and genuinely unavailable historical information [20, 21].

Output: delay probability and explanation

The model output would be a calibrated conceptual probability that a request will experience approval delay beyond a predefined operational threshold, accompanied by a local explanation that identifies the most influential features for that request. A SHAP waterfall or ranked contribution display could show whether delay risk is driven by incomplete clinical documentation, payer-specific policy mismatch, a high-burden procedure category, weak medical necessity indicators, or historically slow payer response [21]. Global explanations could help revenue cycle leaders identify recurring institutional bottlenecks, such as procedure categories that repeatedly lack required documentation or payers with consistently prolonged response patterns [10, 18]. The output should be used as decision support for pre-submission correction and prioritization, not as an automated determination of whether a patient should receive care or whether a payer should approve coverage [26, 28].

Incorporating Payer-Specific Rules and Medical Necessity Indicators

Encoding payer coverage policies

Payer coverage policies would be encoded as structured features that represent whether a request satisfies the payer's stated requirements for diagnosis, prior therapy, severity, site of care, service category, and supporting documentation. This policy library would need to account for variation across commercial plans, public payers, and regional benefit structures, because authorization requirements differ substantially across medication classes, preventive therapies, specialty services, and procedures [7, 8, 17, 25]. Each rule could be represented as a binary, ordinal, or categorical alignment feature, allowing the model to distinguish requests that appear complete from those likely to trigger additional payer review. Because prior authorization reform has been framed as necessary for better patient care, rule encoding should support timely preparation rather than create another opaque administrative checkpoint [29].

Medical necessity scoring from clinical documentation

Medical necessity scoring would translate clinical evidence into interpretable signals indicating whether the request is supported by diagnosis codes, symptom descriptions, prior treatment history, disease severity, functional limitation, imaging findings, or guideline-aligned rationale. Structured EHR fields could capture diagnosis-to-procedure matching, while natural language processing could identify supporting evidence in clinical notes, referral text, letters of medical necessity, and specialist recommendations. Prior authorization research in spine surgery, specialty medications, and payer policy variation suggests that clinical indication and documentation quality can materially influence approval behavior and adjudication speed [7, 15, 16]. The resulting score should remain transparent enough for staff to see which required elements are present, incomplete, or missing before submission [20, 26].

Interaction between procedure type and payer behavior

The model would be designed to capture interactions between procedure type and payer behavior because some services may be delayed primarily due to complexity, while others may be delayed because of payer-specific review patterns. High-cost drugs, radiation oncology services, outpatient antimicrobial therapy, nonemergent ambulance transport, and specialty procedures have different authorization burdens and evidence requirements, making service category a meaningful predictor of delay risk [4-6, 24]. Payer behavior could interact with these categories, such that one payer may process a procedure efficiently while another may require additional review for the same clinical scenario. Modeling these interactions could help authorization teams identify whether the risk is driven by the clinical-service profile, the payer's administrative pattern, or both [10, 16].

The role of historical timelines

Historical timelines would allow the model to represent payer responsiveness as a dynamic operational feature rather than a static payer label. Rolling summaries of prior turnaround, additional information requests, resubmissions, and final payer responses could reveal temporal changes, such as a payer recently applying stricter documentation requirements or a procedure category becoming more difficult to authorize. Evidence from studies of specialty drug authorization, HIV pre-exposure prophylaxis policy variation, and buprenorphine prior authorization suggests that payer requirements and access barriers can shift across plans, regions, and time [8, 9, 17, 25]. These temporal features would need regular updating and governance so that the model reflects current authorization behavior rather than outdated payer patterns [1, 12].

Explainability Methods for Authorization Delay Predictions

Global SHAP analysis for payer and procedure insights

Global SHAP analysis would summarize which features most strongly influence predicted authorization delay risk across the institution. In this framework, global explanations could show whether missing documentation, weak medical necessity alignment, payer identity, procedure category, or historical payer turnaround are the dominant drivers of delay predictions. SHAP methods are well suited for tree-based models because they can connect local feature contributions to broader model behavior, making institutional patterns more interpretable to revenue cycle leaders and clinical operations teams [21]. Such global analysis could support targeted workflow redesign, including improved documentation templates, payer-specific checklists, and service-line education [18, 19].

Local explanation for the authorization specialist

Local explanations would translate a single request's predicted delay risk into actionable information for the authorization specialist preparing the submission. For example, the explanation might indicate that the request is at elevated risk because the documentation lacks a required prior therapy history, the procedure falls under a payer-specific policy rule, and similar requests to the same payer have historically taken longer to approve. This request-level transparency is important because authorization teams need to know what can be corrected before submission, not merely that a delay is possible [13, 26]. In high-stakes healthcare administration, interpretable outputs are preferable to unexplained black-box predictions because staff must be able to justify workflow actions and avoid inappropriate automation [22, 23].

Counterfactual explanations for pre-submission correction

Counterfactual explanations would describe practical pre-submission changes that could reduce predicted delay risk, such as attaching a letter of medical necessity, adding prior treatment documentation, correcting diagnosis-to-procedure alignment, or including payer-required severity evidence. These explanations should avoid unsupported claims of guaranteed approval and instead frame corrections as actions that would be expected to improve request completeness. This approach aligns with the operational purpose of prior authorization decision support, which is to strengthen documentation and reduce preventable rework before payer review [1, 3, 14]. Counterfactual guidance would be most useful when it is grounded in payer-specific rules and displayed in language that authorization specialists and clinicians can act on immediately [7, 25].

Audit trail and transparency in authorization performance

Every prediction, explanation, staff action, submission update, payer response, and final turnaround outcome should be logged to create an audit trail for model governance and operational review. Auditability would allow health systems to examine whether the model is improving documentation quality, whether certain payers or procedures remain associated with persistent delay, and whether explanations are being used appropriately by staff. Transparency is especially important because artificial intelligence used in insurance-related workflows can raise concerns about fairness, accountability, and coverage decision opacity [23, 28]. An audit trail would also allow revenue cycle leaders to evaluate whether administrative interventions reduce preventable delays without shifting burden onto clinicians or patients [10, 11].

Clinical Integration and Pre-Authorization Workflow

Integration into prior authorization software

The model would be integrated into prior authorization software at the request preparation stage, before the authorization is submitted to the payer. When a specialist enters or imports the request, the system would retrieve relevant payer rules, scan available documentation, identify the procedure category, summarize historical payer-procedure timelines, and display a delay risk explanation. Prior work on electronic prior authorization tools and clinical workflow integration suggests that technology must fit staff processes rather than simply digitize existing administrative burden [1, 13]. The model should therefore provide concise, actionable recommendations, such as missing document alerts, medical necessity checklist gaps, or payer-specific attachments needed for submission [20, 26].

Closed-loop feedback from outcomes

After the payer responds, the approval date, denial status, request for additional information, resubmission, or final turnaround time would be captured and linked back to the original prediction. This closed-loop outcome capture would allow the model to be periodically recalibrated as payer policies, procedure mix, documentation practices, and operational workflows change. Studies showing policy variation and authorization burden across medications, specialty services, oncology care, and public programs indicate that authorization environments are not static, making ongoing updating essential [6-9]. However, updates should be governed through validation, audit review, and human oversight so that the system adapts to changing payer behavior without reinforcing biased or inappropriate administrative patterns [22, 28].

Evaluation Strategy

Predictive performance metrics

The model should be evaluated conceptually using delay classification, calibration, and reliability measures that indicate whether predicted delay risk aligns with observed authorization outcomes. Because authorization delay may be unevenly distributed across payers, procedures, and service lines, evaluation should consider discrimination, calibration, and performance across clinically meaningful subgroups rather than relying on a single aggregate statistic. Prior authorization studies across surgery, specialty drugs, preventive therapy, transport, and oncology show that payer and service categories can produce heterogeneous authorization experiences, supporting subgroup-aware evaluation [5-7, 15, 17]. Such evaluation should be framed prospectively and operationally, without claiming performance until real-world validation has been completed [26, 27].

Explanation quality and user acceptance

Explanation quality should be evaluated by whether authorization specialists, clinicians, and revenue cycle leaders understand the model's reasoning and can use it to improve request preparation. Staff-facing evaluation could examine whether explanations are clear, whether they identify correctable documentation gaps, whether users trust the model appropriately, and whether explanations reduce unnecessary escalation or rework. Research on interpretable healthcare artificial intelligence emphasizes that explanation is valuable only when it supports human understanding, accountability, and safe decision-making [21-23]. For this framework, explanation assessment should therefore focus on usefulness in the authorization workflow rather than technical interpretability alone [1, 13].

Operational impact

Operational evaluation should examine whether the model could reduce preventable authorization delays, improve request completeness, decrease avoidable denials, and reduce staff time spent on corrections or repeated payer communication. These outcomes are relevant because prior authorization contributes to administrative cost, provider burden, patient frustration, and operational inefficiency across multiple service categories [3, 10, 11, 24]. Revenue cycle studies using artificial intelligence for claim denial management and administrative decision support suggest that predictive analytics can be operationally meaningful when tied to workflow action and governance [18-20]. In this framework, impact should be assessed through prospective pilots that compare authorization preparation processes before and after implementation without overstating causal effects before formal evaluation [26, 28].

Table 2 summarizes the governance, implementation, monitoring, and failure-mode safeguards required to deploy the explainable prior authorization delay model responsibly in revenue cycle operations.

Table 2. Governance, Implementation, Monitoring, and Practical Action Framework for Explainable Prior Authorization Delay Prediction

Implementation domain	Operational purpose	Key design requirement	Potential failure mode	Governance or monitoring safeguard	Practical implementation action
Pre-submission workflow integration	Embed prediction during authorization preparation before payer submission	Trigger model scoring when a new authorization request is created or imported	Prediction occurs too late to support correction	Monitor timing of prediction relative to submission time	Place model output inside the authorization work queue before final submission.
Human oversight	Ensure the model supports staff and clinician judgment rather than replacing it	Display explanations as recommendations, not automated coverage decisions	Staff may over-rely on the model or ignore clinical nuance	Require human review for high-risk flags and correction recommendations	Assign authorization specialist ownership for reviewing model explanations.
Explanation usability	Make delay drivers understandable and actionable	Use plain-language explanation cards tied to specific missing evidence or payer rules	Explanations are technically accurate but not operationally useful	Conduct user testing with authorization staff and clinicians	Convert SHAP drivers into checklist-style correction prompts.
Payer-rule maintenance	Keep payer policy features current as coverage criteria change	Maintain a versioned payer policy library with date-stamped rule updates	Outdated rules generate misleading delay predictions	Audit payer-rule mappings and review changes at scheduled intervals	Create a revenue cycle policy-maintenance workflow with accountable owners.
Fairness and access monitoring	Detect systematic differences in delay prediction or workflow response by payer, procedure, service line, or patient subgroup	Evaluate prediction patterns and operational actions across relevant subgroups	Model reinforces unequal access patterns or normalizes payer-specific delay	Perform subgroup monitoring and governance review	Report delay-risk patterns to quality, compliance, and revenue cycle leadership.
Model calibration and validation	Ensure predicted delay risk aligns with observed authorization outcomes	Evaluate calibration and reliability across payer-procedure groups	Model appears useful globally but performs poorly for selected services or payers	Conduct local validation and subgroup calibration review	Recalibrate model before broad deployment and after major payer-policy changes.

Audit trail	Preserve accountability for predictions, explanations, staff actions, and outcomes	Log prediction time, input features, explanation output, user action, submission date, and payer response	Decisions become difficult to reconstruct after delays or disputes	Maintain searchable logs for quality improvement and compliance review	Build an authorization analytics dashboard for internal audit, not decorative figure display.
Operational impact evaluation	Determine whether the model improves workflow outcomes	Compare request completeness, rework, turnaround, escalation, and staff burden before and after implementation	Tool adds work without reducing delay or rework	Use prospective operational evaluation with predefined implementation metrics	Pilot in selected service lines before enterprise-wide deployment.
Failure-mode response	Identify and manage incorrect predictions or misleading explanations	Define escalation pathways for high-risk or uncertain cases	False reassurance causes incomplete submissions, or false alarms increase workload	Track false positives, false negatives, and staff override reasons	Add manual review rules for high-cost, urgent, or clinically sensitive requests.

Limitations

Data quality and payer rule dynamics

A major limitation is that prior authorization data may be incomplete, fragmented across payer portals, authorization platforms, EHR systems, scanned documents, phone notes, and staff-entered status updates. Payer policies also change frequently, and documentation completeness can be difficult to measure objectively when medical necessity evidence appears in unstructured clinical notes. These issues can create label noise, missing timestamps, inconsistent definitions of delay, and outdated policy-rule mappings that may weaken model reliability [1, 12, 25]. The framework would therefore require careful data governance, policy library maintenance, and periodic validation before being used to guide operational decisions [21, 22].

Generalizability across payers and procedures

A model developed in one health system may not generalize to another because payer mix, procedure mix, documentation practices, portal use, staffing models, and local escalation pathways differ. Authorization patterns for surgery, oncology care, specialty medications, preventive therapies, transport, and outpatient services may also vary by region and payer contract structure [5, 8, 15, 17, 24]. As a result, the model should be locally calibrated and externally evaluated before being applied across institutions or service lines. Generalizability should be treated as an empirical question, with careful monitoring to ensure that predictions support equitable and transparent workflow improvement rather than amplify administrative inequities [23, 28].

Conclusion

An interpretable machine learning model for predicting prior authorization approval delays could help health systems move from reactive tracking to proactive authorization preparation. By combining payer-specific rules, clinical

documentation features, procedure type, medical necessity indicators, and historical approval timelines, the model could identify requests likely to face avoidable delay before submission.

The central strength of the proposed framework is its integration of predictive modeling with transparent, request-level explanation. Rather than simply flagging a request as high risk, the model would indicate whether the concern arises from missing documentation, payer-rule mismatch, procedure complexity, weak medical necessity evidence, or historically slow payer response.

Important challenges remain before such a model could be deployed responsibly. Payer policies change rapidly, authorization data are often fragmented, documentation quality is difficult to measure, and prospective evaluation is needed to determine whether predictions and explanations improve workflow outcomes.

Future work should focus on collaborative pilots involving health systems, payers, clinicians, authorization specialists, informatics teams, and revenue cycle leaders. Such pilots could test whether explainable prior authorization prediction tools can be embedded into standard workflows while preserving human oversight, transparency, and patient-centered access to care.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 10 Dec 2024

Revised: 28 Jan 2025

Accepted: 27 Mar 2025

Published online: 20 July 2025

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

Salzbrenner SG, McAdam-Marx C, Lydiatt M, Holding B, Scheier LM, Hill PW. Perceptions of prior authorization by use of electronic prior authorization software: a survey of providers in the United States. *J Manag Care Spec Pharm.* 2022;28(10):1121-8.

Jones LK, Ladd IG, Gionfriddo MR, Gregor C, Evans MA, Graham J. Medications requiring prior authorization across health insurance plans. *Am J Health Syst Pharm.* 2020;77(8):644-8.

Gaines ME, Auleta AD, Berwick DM. Changing the game of prior authorization: the patient perspective. *JAMA.* 2020;323(8):705-6.

Bianchini ML, Kenney RM, Lentz R, Zervos M, Malhotra M, Davis SL. Discharge delays and costs associated with outpatient parenteral antimicrobial therapy for high-priced antibiotics. *Clin Infect Dis.* 2020;71(7):e88-e93.

Contreary K, Asher A, Coopersmith J. Evaluation of prior authorization in Medicare nonemergent ambulance transport. *JAMA Health Forum.* 2022;3(7):e222093.

Lin NU, Bichkoff H, Hassett MJ. Increasing burden of prior authorizations in the delivery of oncology care in the United States. *J Oncol Pract.* 2018;14(9):525-8.

Doshi JA, Puckett JT, Parmacek MS, Rader DJ. Prior authorization requirements for proprotein convertase subtilisin/kexin type 9 inhibitors across US private and public payers. *Circ Cardiovasc Qual Outcomes.* 2018;11(1):e003939.

McManus KA, Powers S, Killelea A, Tello-Trillo S, Rogawski McQuade E. Regional disparities in qualified health plans' prior authorization requirements for HIV pre-exposure prophylaxis in the United States. *JAMA Netw Open.* 2020;3(6):e207445.

Keshwani S, Maguire M, Goodin A, Lo-Ciganic WH, Wilson DL, Hincapie-Castillo JM. Buprenorphine use trends following removal of prior authorization policies for the treatment of opioid use disorder in 2 state Medicaid programs. *JAMA Health Forum.* 2022;3(6):e221757.

Sahni NR, Istvan B, Stafford C, Cutler D. Perceptions of prior authorization burden and solutions. *Health Aff Sch.* 2024;2(9):qxae096.

Tseng P, Kaplan RS, Richman BD, Shah MA, Schulman KA. Administrative costs associated with physician billing and insurance-related activities at an academic health care system. *JAMA.* 2018;319(7):691-7.

Bhattacharjee S, Murcko AC, Fair MK, Warholak TL. Medication prior authorization from the providers perspective: a prospective observational study. *Res Social Adm Pharm.* 2019;15(9):1138-44.

Chen WC, Carpenter C, Sidiqi B, Pattison AJ, Hwang J, Pappas D, et al. Integrating prior authorization into clinical workflows for care access and practitioner experience. *JAMA Netw Open.* 2025;8(12):e2549093.

Resneck JS Jr. Refocusing medication prior authorization on its intended purpose. *JAMA.* 2020;323(8):703-4.

De Barros A, Abel F, Kolisnyk S, Geraci GC, Hill F, Engrav M, et al. Determining prior authorization approval for lumbar stenosis surgery with machine learning. *Global Spine J.* 2024;14(6):1753-9.

Doshi JA, Li P, Puckett JT, Pettit AR, Raman S, Parmacek MS, et al. Trends and factors associated with insurer approval of proprotein convertase subtilisin/kexin type 9 inhibitor prescriptions. *Value Health.* 2020;23(2):209-16.

McManus KA, Fuller B, Killelea A, Strumpf A, Powers SD, Rogawski McQuade ET. Geographic variation in qualified health plan coverage and prior authorization requirements for HIV preexposure prophylaxis. *JAMA Netw Open.* 2023;6(11):e2342781.

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.

Pal S, Gaur M, Chaudhuri R, Kalaivanan R, Chetan KV, Praneeth BH, et al. Driving impact in claims denial management using artificial intelligence. In: International Conference on Advances in Computing and Data Sciences. Cham: Springer; 2022. p. 107-20.

Su SC, Huang CC, Gung RR, Hsiung LK, Gao ZW, Tsai CE. A decision support system with artificial intelligence and natural language processing to mitigate the deduction rate of health insurance claims. *Appl Sci.* 2021;11(24):11623.

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.

Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat Mach Intell.* 2019;1(5):206-15.

Wang F, Kaushal R, Khullar D. Should health care demand interpretable artificial intelligence or accept "black box" medicine? *Ann Intern Med.* 2020;172(1):59-60.

Bingham B, Chennupati S, Osmundson EC. Estimating the practice-level and national cost burden of treatment-related prior authorization for academic radiation oncology practices. *JCO Oncol Pract.* 2022;18(6):e974-e987.

Nguemeni Tiako MJ, Dolan A, Abrams M, Oyekanmi K, Meisel Z, Aronowitz SV. Thematic analysis of state Medicaid buprenorphine prior authorization requirements. *JAMA Netw Open.* 2023;6(6):e2318487.

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.

Choudhury A, Perumalla S. Using machine learning to minimize delays caused by prior authorization: a brief report. *Cogent Eng.* 2021;8(1):1944961.

Mello MM, Rose S. Denial-artificial intelligence tools and health insurance coverage decisions. *JAMA Health Forum.* 2024;5(3):e240622.

Miller AP, Shor R, Waites T, Wilson BH. Prior authorization reform for better patient care. *J Am Coll Cardiol.* 2018;71(17):1937-9.