Frontiers in Emerging Computer Science and Information Technology

ISSN 2978-3712

AI Framework for Fair and Auditable Prior Authorization

Rajasekhar Kondreddy [1]

[1] Independent Researcher, Atlanta GA, USA

Email: Shekark333@gmail.com

Received: 29 October 2026 | **Accepted:** 07 November 2026 | **Published:** 8 November 2026

Pages 17–30 | **Crossref DOI** 10.64917/fecsit/Vol3Iss11-02 | **Vol. 3, No. 11 (November 2026)**

Abstract

Prior authorization is a critical function in health insurance; however, traditional processes regularly cause operational delays, member dissatisfaction, inconsistent denials, provider frustration, delayed access to care, and increased regulatory scrutiny. Current authorization workflows are often overloaded by manual reviews, unclear policies, fragmented clinical documentation, limited decision transparency, and insufficient visibility into denial reasons. As health insurers adopt AI-enabled automation, these challenges become more complex if decisions are not explainable, evidence-based, auditable, and supported by human oversight. This study introduces a secure AI-driven decision-support framework for prior authorization that integrates explainable decision logic, policy-rule traceability, clinical-evidence mapping, bias monitoring, human review checkpoints, and appeal-ready explanations within a single governance model. Unlike automation focused mainly on speed and cost, the proposed framework prioritizes fair, transparent, and accountable decision-making across the authorization lifecycle. Each AI recommendation explains why it was made and links the decision to medical guidelines, policy rules, supporting evidence, prior reviewer actions, and member-specific details. The framework is expected to reduce authorization turnaround time, decrease avoidable denials, reduce appeal overturns, improve provider and member satisfaction, and make decisions more consistent and fairer across clinical and demographic groups.

Keywords: Prior authorization; Explainable AI; Clinical decision support; Policy-rule traceability; Algorithmic fairness; Human-in-the-loop review; Health insurance; Auditability; AI governance.

ISSN 2978-3712

How to Cite: The final bibliographic citation, including volume, issue, page numbers, and DOI (if applicable), will be assigned and made available upon official publication.

1. Introduction

Prior authorization (PA) is a process through which health plans review a requested drug, procedure, or service in advance to determine whether it will be covered. Historically, PA was introduced to limit unnecessary spending, promote appropriate care, and protect members from unsafe or ineffective services. In practice, growing policy complexity, disconnected clinical records, manual reviews, and inconsistent payer rules have made the process administratively heavy. The American Medical Association's 2024 national physician survey reported that the large majority of physician's experience PA-related care delays and believe the process negatively affects patient outcomes [1].

This study proposes a secure, AI-driven decision-support framework that combines clinical-evidence mapping, policy-rule traceability, human review checkpoints, bias monitoring, and appeal-ready explanations. Unlike traditional automation focused mainly on speed and cost, the framework links every recommendation to clinical guidelines, member-specific evidence, benefit rules, and reviewer actions.

The solution improves fairness by monitoring how approvals and denials vary across different patient groups, detecting unexplained differences, and routing uncertain or potentially biased cases for human review. The goal is to make prior authorization faster and more consistent while keeping clinical accountability, policy compliance, and member protection at the center of the process.

2. Literature Review

This literature review was developed by examining peer-reviewed studies, professional association reports, regulatory publications, and recent healthcare AI governance materials related to prior authorization, clinical decision support, explainable AI, algorithmic fairness, human oversight, interoperability, and AI governance. The review focused mainly on work published between 2019 and 2026, because this period captures the rapid growth of AI in healthcare, the emergence of algorithmic bias concerns, and recent federal action around electronic prior authorization. The search used terms such as prior authorization burden, AI in prior authorization, clinical decision support, explainable AI in healthcare, algorithmic bias healthcare, human-in-the-loop AI, FHIR prior authorization, and AI governance in healthcare.

Prior authorization was originally designed as a utilization management tool to confirm medical necessity and coverage before care is delivered. In practice, the process has become a major administrative and clinical burden. The 2024 AMA physician survey reported that 93% of physicians experienced PA-related care delays, 94% believed PA negatively affected patient outcomes, and 82% reported that patients sometimes abandoned recommended treatment because of PA problems [1]. Sahni et al. studied the issue from patient, provider, and payer perspectives and showed that the operational burden is carried heavily by providers and clinical staff, while patients experience the downstream effects as delay, confusion, or disrupted treatment [2]. These studies show that PA is not only a payer workflow

ISSN 2978-3712

issue; it is also a care-delivery problem that affects access, trust, and clinical continuity.

As PA volume and complexity increased, insurers and technology vendors began exploring automation and AI-enabled decision support. Prior work suggests that AI can help extract clinical information, match documentation against criteria, identify missing evidence, and reduce repetitive manual review. Lenert et al. argued that an AI-based approach to PA could be more aligned with expert clinical reasoning than rigid rule-only automation if it is designed around clinical judgment rather than denial acceleration [3]. However, recent AMA reporting shows growing concern that unregulated payer AI may increase denials through automated or batch-based decisions with limited human review [4]. This means automation alone is not a reform. AI must be implemented as a governed decision-support layer with transparency, evidence linkage, and clinical accountability.

Clinical decision support systems have long helped clinicians organize patient data, guidelines, and evidence-based recommendations. In PA, the challenge is more complex because a request must satisfy both clinical appropriateness and benefit coverage. A service may be medically reasonable but not covered under a specific policy, or it may be covered only when certain documentation is present. Therefore, AI for PA must connect the member's clinical condition, the applicable medical guideline, and the payer's benefit or medical-necessity rule. If AI produces only a score or yes/no output, reviewers cannot easily verify whether the right diagnosis, guideline, and policy requirement were used. This supports a framework where clinical evidence mapping and policy-rule mapping remain separate, traceable, and auditable.

Explainable AI has become important because healthcare decisions require trust, reasoning, and accountability. Abbas et al. reviewed 62 studies on explainable AI in clinical decision support and found that many methods are technically promising, but real-world usability, clinician trust, and explanation fidelity remain underdeveloped [5]. Rosenbacke et al. similarly found that explanations can improve trust only when they are clear, concise, and clinically meaningful; overly complex explanations may not help and can even reduce trust [6]. For PA reviewers, this means that feature-importance charts alone are not enough. Reviewers need simple explanations showing the guideline used, the policy rule applied, the supporting documentation found, and what evidence is missing.

Algorithmic fairness is critical because biased healthcare decisions can affect access to necessary care. Obermeyer et al. showed that a widely used healthcare algorithm underestimated the needs of Black patients because it used healthcare spending as a proxy for illness; historical spending differences caused unequal results even when the model appeared statistically valid [7]. In PA, a similar risk can occur if AI is trained on historical denial patterns, incomplete documentation, cost-based proxies, or reviewer behavior. Fairness therefore must be monitored not only through average accuracy, but also through approval rates, denial rates, requests for more information, appeal overturns, reviewer overrides, and turnaround time across relevant clinical and demographic groups.

ISSN 2978-3712

Healthcare AI literature often recommends a human-in-the-loop model, but human review is meaningful only when the reviewer has time, authority, and evidence to challenge the system. A reviewer who simply accepts AI output is not providing real clinical oversight. In PA, human review should be required for uncertain, adverse, borderline, documentation-incomplete, or potentially biased cases. The reviewer must see the evidence, understand the rationale, override the recommendation when appropriate, and document the reason. This makes human oversight an accountable checkpoint rather than a symbolic final signature.

AI governance frameworks emphasize that healthcare AI must be valid, safe, secure, transparent, explainable, privacy-preserving, and accountable. The NIST AI Risk Management Framework identifies these traits as central to trustworthy AI systems [8]. In PA, these principles must become operational controls. Each decision should be reconstructable later, showing the request details, clinical guideline used, policy rule applied, evidence found, missing documentation, AI recommendation, confidence level, reviewer action, and final outcome. Auditability is essential because regulators, providers, courts, and members increasingly expect health plans to explain how AI-supported coverage decisions are made.

Interoperability is another key part of PA reform. CMS released the Interoperability and Prior Authorization Final Rule, CMS-0057-F, in January 2024 to improve information exchange and require APIs supporting patient access, provider access, payer-to-payer exchange, and prior authorization workflows [9]. This creates a technical foundation for faster PA, but interoperability alone does not guarantee fairness, explanation, or clinical accountability. APIs can move data faster; they do not decide whether the AI decision is trustworthy. Therefore, electronic PA must be paired with evidence mapping, policy traceability, human review, bias monitoring, and audit logging.

The regulatory landscape is moving quickly because AI is now influencing decisions that affect access to care. CMS rules are pushing payers toward electronic PA, denial-reason transparency, and improved turnaround-time reporting [9]. At the same time, professional organizations such as the AMA have raised concerns that payer AI may operate without enough individualized review or clinical oversight [4]. Future PA systems must therefore be designed for defensibility, not just efficiency. Payers will increasingly need to show why a decision was made, what evidence was used, whether a qualified person reviewed the case, and whether similar cases are handled consistently.

The reviewed literature shows strong progress in separate areas. PA studies document burden and patient harm. Automation studies show that AI can reduce manual effort. Explainable AI research shows how recommendations can become more understandable. Fairness research warns that healthcare algorithms can reproduce inequity. Governance frameworks explain the need for accountability and risk management. CMS interoperability rules create the technical pathway for electronic PA. However, the literature does not yet provide a complete practical framework that combines all these elements into one PA operating model. Existing work usually focuses on one part of the problem: speed, automation, explainability, fairness, human review, or interoperability. The contribution of this paper is to bring these pieces

ISSN 2978-3712

together into a single framework where every AI recommendation is linked to clinical evidence, policy rules, reviewer action, fairness monitoring, and an appeal-ready explanation. This positions AI not as a replacement for clinical judgment, but as a governed decision-support system that improves speed, consistency, transparency, and fairness in prior authorization.

3. Proposed Framework

The proposed framework is designed as an end-to-end prior authorization decision-support model in which operational speed, clinical reasoning, policy compliance, fairness, and auditability are treated as one connected problem rather than separate technical tasks. As shown in Figure 1, the workflow begins with the provider request and then moves through secure intake, evidence and policy mapping, explainable AI decision support, human accountability, decision output, and continuous monitoring. The purpose of this design is not to replace clinical judgment with automation. Instead, it gives reviewers a structured and evidence-based decision environment in which each recommendation can be checked, explained, challenged, and improved over time.

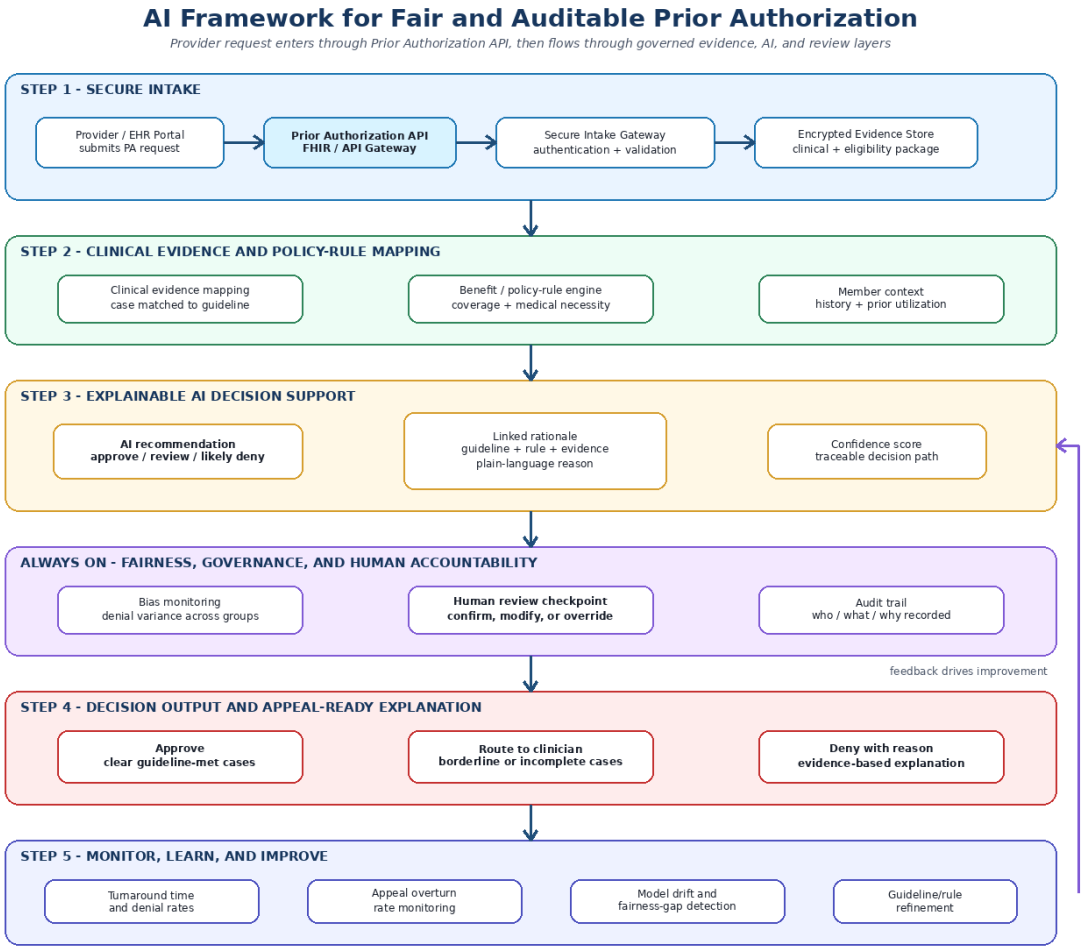


Figure 1. Redesigned explainable and auditable AI framework for fair prior authorization decision support.

ISSN 2978-3712

The first part of the framework is secure intake. This is the point where the provider or electronic health record system submits the authorization request through a Prior Authorization API. In traditional workflows, the request may arrive through several channels, such as portals, faxes, attachments, or manual uploads, which often causes missing documentation and repeated follow-up. In the proposed framework, the intake layer creates a single controlled entry path. The request is authenticated, the submitting provider is verified, required fields are checked, and the clinical documents are linked to the member's eligibility and benefit context. The result is an encrypted evidence package that becomes the trusted starting record for downstream review. This stage is important because an AI recommendation is only as reliable as the information it receives. If eligibility, diagnosis, clinical notes, or benefit rules are incomplete at the start, the later decision may be fast but still wrong. Secure intake therefore protects both the member and the payer by improving data quality before any recommendation is generated.

After intake, the framework organizes the submitted information into two related but distinct forms of knowledge: clinical evidence and policy rules. The clinical side asks whether the requested service is reasonable for the member's condition, diagnosis, history, and current documentation. The policy side asks whether the requested service meets the coverage and medical-necessity requirements of the member's specific plan. Keeping these two streams, separate is intentional. A service may be clinically appropriate but still require additional policy documentation; similarly, a policy rule may be satisfied only when specific evidence is present in the medical record. The mapping layer therefore matches the case to relevant clinical guidance, retrieves the applicable benefit and medical-necessity rule, and connects both to member-specific history such as prior care, utilization, previous authorizations, and supporting documents. This creates a transparent bridge between medicine and coverage policy, rather than allowing either side to remain hidden inside an opaque model.

The explainable AI decision-support engine uses the mapped evidence to generate a preliminary recommendation. The recommendation may be to approve the request, route it for clinician review, or identify it as likely not meeting requirements based on the available documentation. The key design principle is that the system does not return a bare decision label. Every recommendation is accompanied by a confidence level and a traceable rationale that identifies the clinical guideline considered, the policy rule applied, the evidence found in the file, and any information that appears missing or inconsistent. This is especially important in prior authorization because reviewers, providers, and members need to understand why a decision was made. A black-box result may reduce processing time, but it does not create trust or defensibility. The proposed framework therefore treats explanation as part of the decision itself, not as an after-the-fact report added later.

Human accountability is built into the workflow rather than placed outside it. The system may support a reviewer, but it does not remove responsibility from qualified personnel. Straightforward requests that clearly satisfy guideline and policy requirements may move quickly toward approval. However, cases that are uncertain, high-impact, documentation-

ISSN 2978-3712

incomplete, clinically complex, or potentially adverse are routed to a human reviewer. The reviewer can confirm, modify, or override the AI recommendation and must have access to the same evidence trail used by the system. This design avoids the weakness of symbolic human oversight, where a person merely accepts automated output without enough context. In this framework, human review is a meaningful checkpoint because the reviewer can see the reasoning, inspect the documents, compare the case against policy, and record the final judgment.

Fairness and governance operate continuously across the entire workflow. The framework monitors decision patterns such as approvals, denials, requests for additional information, appeal overturns, and reviewer overrides across relevant demographic and clinical groups. The purpose is not to assume unfairness in every decision, but to detect patterns that cannot be explained by clinical need, policy rules, or documentation quality. This is consistent with the broader principle that healthcare AI must be monitored after deployment, because bias can emerge from historical data, changing policies, shifting documentation practices, and model drift. Governance also includes a complete audit trail. For each authorization, the framework records what information was received, which guideline and policy rule were used, what evidence supported the recommendation, who reviewed the case, whether the recommendation was overridden, and why the final decision was made. This audit record supports internal quality review, regulatory examination, and appeal handling.

The decision-output stage translates the review into an action that can be understood by providers and members. When a request clearly meets the guideline and benefit criteria, the framework supports approval with minimal delay. When the evidence is incomplete or the case is borderline, the system routes it for clinician review or requests additional information rather than defaulting immediately to denial. When a denial is appropriate, the explanation is structured to be appeal-ready from the beginning. It states the relevant guideline, the applicable policy rule, the evidence that was found, and the specific missing or conflicting information. This approach improves procedural clarity because the provider does not have to guess why the request failed or what evidence is needed to contest the decision. It also helps payers demonstrate that the decision was based on documented criteria rather than unexplained automation.

The final stage is continuous monitoring and learning. Prior authorization rules, clinical practice, member populations, and documentation patterns change over time. A framework that performs well at implementation may become less reliable if it is not monitored. For this reason, the proposed model tracks operational and governance indicators such as turnaround time, avoidable denials, appeal-overturn rates, reviewer override rates, requests for additional information, and fairness-gap measures. These findings are fed back into policy refinement, guideline updates, reviewer training, and model recalibration. The feedback loop is not meant to let the model learn without control; rather, it provides evidence for controlled improvement under governance oversight. In this way, the framework becomes a learning authorization system that can improve consistency and efficiency while remaining clinically accountable.

ISSN 2978-3712

Overall, the proposed framework contributes a practical architecture for responsible AI use in prior authorization. Its novelty lies in the integration of secure intake, evidence mapping, policy traceability, explainable recommendation logic, human review, fairness monitoring, auditability, and appeal-ready communication into one workflow. The design recognizes that prior authorization is not simply a technical classification task. It is a decision process that affects access to care, provider workload, payer compliance, and member trust. By making every recommendation evidence-linked, reviewable, and monitorable, the framework supports faster processing without sacrificing transparency or fairness.

4. Methodology and Implementation Considerations

Implementing the proposed framework requires decisions across four practical dimensions: data and system integration, model architecture, explainability and governance tooling, and security, privacy, and compliance.

4.1 Data and System Integration

The framework is designed to ingest structured and unstructured inputs, including EHR extracts, fax-scanned notes through OCR, claims history, and benefit-plan configuration through a single secure intake layer, consistent with emerging interoperability requirements for electronic prior authorization [9].

4.2 Model Architecture

A hybrid architecture is recommended because prior authorization requires both flexible clinical-evidence interpretation and exact policy-rule enforcement. As illustrated in Figure 2, AI-based extraction is used to read and organize clinical evidence, while deterministic rule services apply benefit configuration, coverage policy, and medical-necessity criteria. The decision-orchestration layer combines these outputs into a recommendation, but the recommendation remains traceable to source evidence, policy rules, confidence scores, and reviewer action. This separation is important because policy logic must remain exact and auditable, while evidence extraction can benefit from machine-learning or language-model techniques when supported by human review.

ISSN 2978-3712

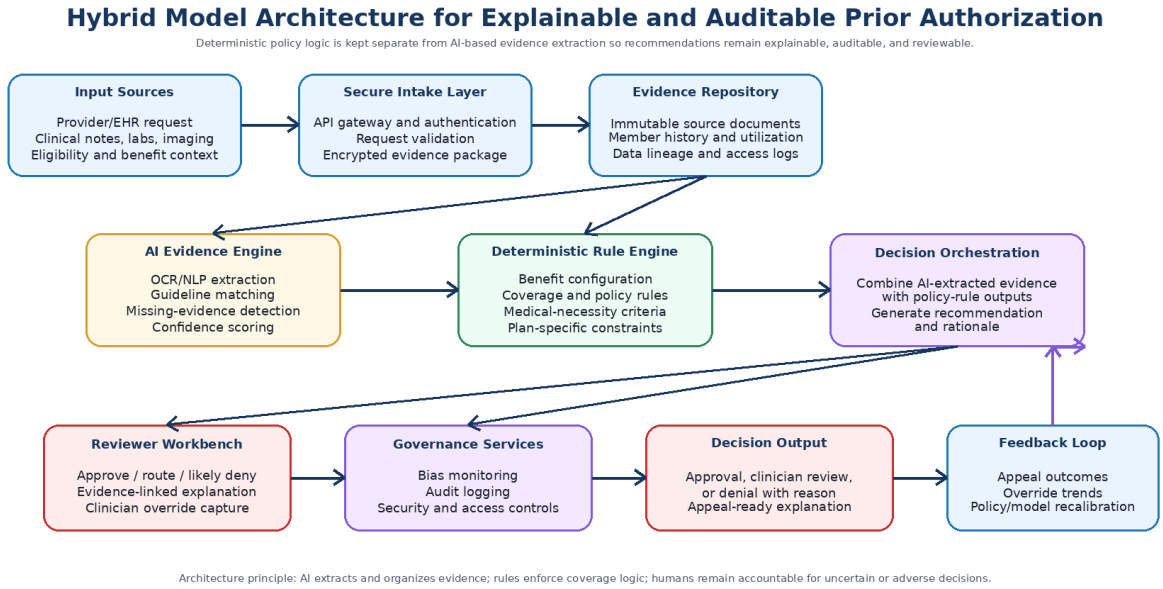


Figure 2. Hybrid model architecture for explainable and auditable prior authorization decision support.

4.3 Explainability and Governance Tooling

Rather than relying solely on statistical explanation methods, the framework attaches structured citations, including guideline identifier, policy-rule identifier, and evidence span, to every recommendation at generation time. A governance committee consisting of clinical, compliance, and technical stakeholders should own bias-monitoring thresholds and review-escalation criteria.

4.4 Security, Privacy, and Compliance

All data in transit and at rest should be encrypted, with role-based access control restricting evidence and member data visibility to authorized reviewers. Audit logs should be immutable and retained in accordance with applicable regulatory retention requirements.

5. Hypothesized Outcomes Thresholds

The expected outcomes in this study should be interpreted as hypothesis thresholds for future pilot evaluation rather than empirically demonstrated results. The thresholds are derived from the framework design, the literature review, and the qualitative comparison supplied in the configuration matrix, where the full framework is positioned as high in explainability, fairness, auditability, and efficiency compared with partial AI-only or rules-only configurations.

For evaluation, the traditional approach is treated as the baseline condition. It represents a workflow that commonly depends on manual document handling, fragmented intake, reviewer interpretation, and limited structured traceability. The proposed full framework is expected to outperform that baseline because it combines secure intake, policy-rule traceability, evidence-linked AI recommendations, human review, fairness monitoring, and audit logging. These

ISSN 2978-3712

expected values should be tested through retrospective validation, shadow-mode deployment, and controlled rollout before being presented as measured performance outcomes.

Table 1. Hypothesized outcome thresholds for traditional prior authorization versus the proposed full framework.

Evaluation Dimension	Traditional PA Baseline	Proposed Full Framework Hypothesis Threshold	Measurement Method	Status
Authorization turnaround time	Manual or semi-manual baseline; often delayed by fragmented intake and follow-up	20-30% relative reduction from baseline cycle time	Compare median request-to-decision time before and after pilot	Hypothesis threshold
Avoidable / documentation-driven denials	Baseline index = 100; denials may occur when evidence is missing or hard to trace	15-25% relative reduction in avoidable denials	Track denials later overturned or resolved after additional documentation	Hypothesis threshold
Appeal overturn rate	Baseline index = 100; overturns indicate weak rationale or incomplete initial review	10-20% relative reduction after evidence-linked explanations	Compare appeal outcomes and overturn reasons by service line	Hypothesis threshold
Requests for additional information	Baseline index = 100; repeated follow-up common when intake is incomplete	10-20% relative reduction through secure intake validation	Measure duplicate or missing-document follow-up requests	Hypothesis threshold

ISSN 2978-3712

Evaluation Dimension	Traditional PA Baseline	Proposed Full Framework Hypothesis Threshold	Measurement Method	Status
Explainability	Low to medium; rationale may be narrative, inconsistent, or not linked to evidence	High; at least 90% of recommendations should include guideline, policy rule, evidence span, and plain-language rationale	Audit completed rationales for required explanation elements	Hypothesis threshold
Auditability	Medium; records may exist but are not always reconstructable end to end	High; at least 95% of decisions should have complete audit trail	Review logs for evidence, rule, recommendation, reviewer, override, and final decision	Hypothesis threshold
Fairness monitoring	Low; subgroup variance often not monitored as part of routine PA workflow	High; monthly monitoring with governance review when unexplained denial variance exceeds 5-8 percentage points	Track approval, denial, RFI, override, and appeal metrics across relevant groups	Hypothesis threshold
Efficiency	Medium; human effort remains high even when some automation is present	High; efficiency gain expected without removing human accountability	Measure reviewer handling time, queue age, and auto-clearance of guideline-met cases	Hypothesis threshold

ISSN 2978-3712

Evaluation Dimension	Traditional PA Baseline	Proposed Full Framework Hypothesis Threshold	Measurement Method	Status
Human accountability	Variable; reviewer reasoning may not always be connected to evidence and policy trace	High; 100% of adverse, low-confidence, or borderline cases routed to qualified review	Check routing logs, reviewer notes, and override documentation	Hypothesis threshold

In this evaluation design, the attached configuration matrix is treated as a qualitative model-selection guide: AI-only configurations may improve speed but remain weak in explainability, fairness, and auditability; adding rules improves policy consistency; adding evidence improves traceability; adding human review improves accountability; and the full framework is hypothesized to achieve high performance across all four evaluation dimensions. The proposed values therefore define pilot success thresholds, not confirmed empirical outcomes.

6. Discussion

The primary contribution of this work is architectural rather than algorithmic. It does not propose a new machine-learning model; instead, it presents a governance-integrated system design in which explainability, evidence traceability, fairness monitoring, and human accountability are structural requirements rather than optional add-ons.

This distinguishes the framework from automation efforts that report efficiency and cost gains without an equally developed governance layer. It also operationalizes recommendations from the AI-in-healthcare literature, which call for explanations that are clinically meaningful rather than purely statistical, and from the fairness literature, which calls for continuous rather than point-in-time bias auditing.

The framework directly addresses the central risk that AI adoption in prior authorization could make existing transparency and equity problems worse rather than better if governance is treated as an afterthought.

7. Regulatory and Policy Landscape.

The proposed framework aligns with evolving U.S. healthcare interoperability, transparency, and AI-governance requirements. CMS Interoperability and Prior Authorization Final Rule (CMS-0057-F) promotes electronic prior authorization through FHIR-based Prior Authorization APIs, Provider Access APIs, and Payer-to-Payer APIs to improve data exchange, accelerate authorization workflows, and reduce administrative burden. The framework supports structured documentation

ISSN 2978-3712

requirements, denial-reason transparency, turnaround-time compliance, and public reporting of prior authorization performance metrics. HIPAA privacy and security controls are incorporated through encryption, role-based access controls, audit logging, and minimum-necessary data access principles. In addition, the architecture can accommodate state-specific prior authorization regulations, continuity-of-care mandates, disclosure requirements, and emerging healthcare AI governance obligations.

Regulatory compliance considerations are embedded throughout the framework lifecycle, ensuring that interoperability, explainability, transparency, privacy, security, and auditability requirements are treated as architectural controls rather than post-implementation compliance activities.

8. Limitations

The proposed framework is a conceptual and architecture-driven model that has not yet been validated through large-scale real-world deployment across multiple payers, provider organizations, and clinical specialties. Therefore, the performance improvements presented in this study should be interpreted as hypothesized outcomes rather than proven results. The framework's effectiveness depends on the quality, completeness, and interoperability of clinical documentation, claims data, and policy-rule repositories. Variations in documentation practices, legacy system limitations, and data quality issues may affect recommendation accuracy and explainability. In addition, fairness monitoring can identify potential disparities but cannot completely eliminate all sources of historical or operational bias. Successful implementation also requires ongoing governance, regulatory compliance management, and organizational investment in technology, workforce training, and continuous model oversight.

9. Conclusion and Future Work

Prior authorization is at an important inflection point. AI can meaningfully reduce delay and administrative burden, or it can compound the opacity and inconsistency that already frustrate providers and patients. This paper presents a framework in which explainable decision logic, clinical-evidence mapping, policy-rule traceability, human review checkpoints, bias monitoring, and appeal-ready explanations are combined within a single governance model.

Future work should pursue empirical validation through a phased pilot with real payer and provider data, cross-payer benchmarking of turnaround time and denial-variance metrics, and longitudinal study of appeal-overturn rates following framework adoption. Additional research should evaluate specialty-specific adaptation, reviewer workload impact, provider experience, member understanding of explanations, and the operational cost of maintaining continuous fairness monitoring.

References

1. American Medical Association, 2024 AMA Prior Authorization Physician Survey, 2024.
2. N. R. Sahni, B. Istvan, C. Stafford, and D. Cutler, “Perceptions of prior authorization burden and solutions,” Health Affairs Scholar, vol. 2, no. 9, 2024.

ISSN 2978-3712

3. L. A. Lenert, S. Lane, and R. Wehbe, “Could an artificial intelligence approach to prior authorization be more human?” Journal of the American Medical Informatics Association, vol. 30, no. 5, pp. 989-994, 2023.
4. American Medical Association, “How AI is leading to more prior authorization denials,” 2025.
5. Q. Abbas, W. Jeong, and S. W. Lee, “Explainable AI in clinical decision support systems: A meta-analysis of methods, applications, and usability challenges,” Healthcare, vol. 13, no. 17, 2025.
6. R. Rosenbacke, A. Melhus, M. McKee, and D. Stuckler, “How explainable artificial intelligence can increase or decrease clinicians’ trust in AI applications in health care: Systematic review,” JMIR AI, vol. 3, 2024.
7. Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan, “Dissecting racial bias in an algorithm used to manage the health of populations,” Science, vol. 366, no. 6464, pp. 447-453, 2019.
8. National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), 2023.
9. Centers for Medicare & Medicaid Services, CMS Interoperability and Prior Authorization Final Rule (CMS-0057-F), 2024.
10. Murphy, J., et al. (2026). Adverse effects of health plan prior authorization on clinical effectiveness and patient outcomes. *The American Journal of Medicine, 139*(1), 24–32.e1.
11. Jain, R. (2024). AI-driven personalization in usage-based insurance. *The American Journal of Engineering and Technology, 6*(1), 25–32.
12. Jain, R. (2024). Voice AI risk signaling using home-assistant devices. *The American Journal of Engineering and Technology, 6*(5), 43–48.
13. Chin, M. H., et al. (2023). Guiding principles to address algorithm bias and racial and ethnic disparities. *JAMA Network Open, 6*(12), e2345050.