

TRUST-AWARE AUTONOMOUS PRIOR AUTHORIZATION ORCHESTRATION FRAMEWORK USING FHIR, AI, AND MULTI-AGENT HEALTHCARE SYSTEMS

Ramani Eedara

Kent State University, 800 E Summit St, Kent, OH 44240. cedararamani@gmail.com

Abstract: Prior authorization is essential to healthcare administration as a process to ensure that treatments are both clinically necessary and cost-effective prior to approval by the healthcare insurance company. In most cases, however, one or more of these data exchange and/or manual efforts coincide, increasing the burden of data management and the time needed to deliver patients' care. As HL7 Fast Healthcare Interoperability Resources (FHIR), an innovative healthcare framework that allows electronic healthcare data to be shared and exchanged much faster and in significantly richer detail, becomes increasingly embraced, recent regulatory developments like the CMS-0057-F Final Rule and the work of artificial intelligence (AI) provide exciting new avenues to modernize prior authorization for the benefit of patients. This survey would be keen to collect and discuss the latest published research related to interoperability in healthcare (IEE) systems, FHIR format, AI-powered clinical decision support systems, explainable AI, multi-agent systems in healthcare, and the potential and benefits of using these tools to enhance the PA work process. It also identifies key areas that should be tackled regarding issues of trust, transparency, auditability, regulatory control and human oversight, and gaps in the field. Following the evidence review, the paper presents the proposed Trust-Aware Autonomous Prior Authorization Framework (TAPAF), which integrates interoperability considerations, accountabilities, governance, and AI agents to enable efficient, explainable, and patient-centric prior authorization processes.

Keywords: Prior Authorization; HL7 FHIR; Multi-Agent Healthcare Systems; Explainable Artificial Intelligence; Healthcare Interoperability

1. INTRODUCTION

Prior authorization (PA) is a recent component of healthcare spending, as a healthcare payer decides if a proposed medical necessity and coverage are met before payment from the payer is approved for proposed medical services, diagnostic testing, pharmaceuticals, or specific treatments [1]. The objective of PA is to: facilitate an evidence-based practice-based approach to healthcare delivery; leverage resources and prevent waste; prevent inconsistencies across healthcare delivery in the clinical decision-making process; and prevent gaps in healthcare delivery payments consistent with healthcare evidence-based clinical practice guidelines [2]. Prior authorization has evolved into a strategic tool to cope with a rising number of patients, the rising complexity of treatment, and rising healthcare costs, now becoming a part of healthcare management and financial sustainability throughout the world [3]. However, this use of “traditional” PA workflows is still mostly manual, with numerous papers and/or electronic forms, phone calls, fax messages, and multiple disparate information systems for any single PA request for authorizations to MRI, lab, physician notes, medicaid coverage, etc. required [4]. These hard tasks result in much wasted time in admin and impair timeliness for diagnosis, increased work for clinicians and clinician burnout, as well as lost patient satisfaction. A considerable amount of time spent during the week on authorization-related activities has been demonstrated in many studies, and in many instances, this time is deducted from patient attention and clinical time spent by HCP [5]. While the healthcare data space has evolved and moved into digital form in recent years, so



has the chaos in the healthcare app space with the emergence and adoption of interoperable EHRs (electronic health records), cloud-based clinical information systems (CIS), and healthcare data exchange frameworks that seek to eliminate inefficiencies in healthcare communication [6]. Among these advances is the new interoperability standard Health Level Seven (HL7) Fast Healthcare Interoperability Resources (FHIR), which is the current main interoperability standard because of its modular resource structure, RESTful application programming interfaces, extensibility, and ability to include current healthcare software architectures [7]. To represent and exchange representations of various clinical entities, FHIR offers a set of specific resources that can be used by any health care provider, any health care insurer, any laboratory, and any third-party application to exchange representations of structured clinical information in near-simultaneous fashion [8]. Meanwhile, the advancement of artificial intelligence, machine learning, natural language processing (NLP), and foundation models has been very impressive, with a multitude of applications, including summarizing physician notes, automating administrative tasks, forecasting treatment outcomes, deriving clinical evidence from unstructured medical records, and even using the information to drive evidence-informed treatment and care [9]. More recently, the paradigm of multi-agent systems has emerged as a promising approach for distributed healthcare automation. In this model, autonomous software agents are assigned specialized roles to retrieve clinical information, interpret healthcare policies, assess eligibility and medical necessity, validate documentation, coordinate workflows, and collaboratively achieve complex healthcare objectives [10]. In addition, the technologies offer a stepping stone to building a reliable, autonomous prior authorization system that has the potential to make the process more efficient and reduce the risk of accountability, regulatory violations, or compromising patient safety through a combination of explainable AI, explainable decision logic, full audit trails, governance policies, and clinician oversight.

Although healthcare is advancing toward interoperability and intelligent automation, fully autonomous and trustworthy patient administration systems remain limited by heterogeneous implementation approaches, fragmented data ecosystems, and evolving regulatory requirements. The vast majority of current electronic prior authorization systems are simply digitizing the paper authorization/file submission process with little to no change in the process or impact on intelligent clinical orchestration [3]. Although there are various connected electronic health records (EHRs) and various types of laboratory information systems, radiology images are stored within radiology archives, and medical care is processed within the insurance databases, along with some clinical documentation repositories, there's little integration between networks, making it tough for medical record claims to be processed correctly. There are many Interoperability standards that can be used, but patient data across claims used for medical necessity is spread out across different health care systems, such as EHR, Clinical documentation repositories, Insurance systems and databases, Pharmaceutical software, laboratory information systems, etc. While the HL7 FHIR standard has enhanced the capacity to share structured healthcare-related information, there are also several problems to overcome to get the best out of health-related data exchange: Semantic interoperability issues, inconsistencies in implementation, some poorly represented resource types in FHIR, and varying levels of adoption [8]. In addition, payer organizations may also have their own coverage policies, different copayment requirements, and even rapidly changing clinical policy guidelines, creating yet more complexity for health care providers attempting to meet a multitude of payers' documentation requirements for authorizations [2]. A series of regulatory initiatives have been underway for the past several years that encourage the industry to move towards standardized ePrior Authorization (ePA) using FHIR-based application programming interfaces, standardized documentation exchanges, quicker turnaround times, improved transparency along the entire ePA life cycle, etc. But most of the current AI-based health care solutions are not built as a self-contained orchestration system to enable the use of AI or as a suite of features to customize for a health care setting; they are synonymous with regulatory compliance, and regulatory compliance is certainly not the same as 'intelligent automation' or 'trustworthy decision-making'. To date, AI solutions that span healthcare challenges have tackled problems one by one and not yet reached an E2E orchestration model of AI [9]. In addition, the adoption of machine learning and large language models adds new challenges on algorithmic bias, hallucinated recommendations, limited interpretability, privacy preservation, cybersecurity, accountability and ethical governance, when the proposed recommendations by AI are impacting clinical or reimbursement decisions [11-13]. The framework is not intended to replace clinicians in interactions with clinical stakeholders, patients, regulators, or insurers. Instead, it is designed to

support clinical decision-making by providing transparent, explainable, auditable, and reliable outputs that address the information needs of stakeholders while preserving the clinician's central role in patient care. A preferable solution is multi-agent architectures, because they can communicate with each other's limited healthcare-specific functions via standardized protocols; therefore, they allow modularity, scalability, resiliency, and adaptability with regard to workflow management. With prior authorization being at the heart of the subject, however, there seem to be few studies investigating the topic in a comprehensive, holistic manner, namely, from an explainability, cooperation with autonomous agents, trust interlocked with the prior authorization, governance over AI, observance of rules and regulations, and human-in-the-loop perspectives. The majority of the research focused on one particular technology, hence the need for systematic research on this theme and the synthesis of this gap [14, 15].

To overcome these problems, researchers from the respective streams in the healthcare interoperability literature, the AI literature, the literature on Explainable AI (XAI), and the Multi-Agent Systems literature collected and integrated existing research to investigate how such a combination of these systems can get the job done together, complementing the document-centric mode of operation used in administration, with an intelligent, interoperable and trustworthy component of the healthcare workflow, in the meantime, as much as "prior authorization." This review differs from those conducted previously, which tended to concentrate on the diffusion of some single facet of FHIR use (e.g., implementation, use of AI in clinical decision support, or use of FHIR in healthcare processes), by taking an interdisciplinary approach that focuses on the principles of trust, governance, transparency, and compliance with regulation as central to designing autonomous healthcare systems. In this paper, the following aims have been structured:

1. Discuss the evolution and changes in prior authorization requirements, how EPA systems have evolved, the progress of HCR/FHIR standards, and new regulations on electronic PA systems.
2. Discuss the advancements of A.I., X.A.I., and a multi-agent healthcare framework that brings intelligence and evidence to the healthcare space to support clinical decisions, automate clinical workflow, and ensure medical necessity.
3. Analyze current work in terms of interoperability, level of automation, explainability, trust, audits, regulatory compliance/human in the loop decision-making, and identify similarities, strengths and weaknesses.
4. Identify critical research gaps related to the loss of transience and fragmentation, fewer trust mechanisms, weak governance, and no transparency, as well as agents in an ecosystem of health services.
5. Develop the **Trust-Aware Autonomous Prior Authorization Framework (TAPAF)** by integrating FHIR-enabled interoperability, specialized AI agents, explainable AI, governance mechanisms, and human oversight.

2. LITERATURE REVIEW

Over the last twenty years, the field of electronic authorization, interoperability, AI studies, AI explainability, multi-agent systems and medical governance has diversified and rapidly expanded. Structured literature review methodology was adopted before synthesizing the studies available in the literature to guarantee that the literature review was comprehensive, relevant, and was conducted in a methodological way. The method sets out the process for searching, data sources, selection of studies and the inclusion criteria for assessing the most relevant publications found. Thematic synthesis was conducted based on the conceptual analysis of the studies selected; the stages aimed in the literature work were then organized according to the main research fields of the studies as discussed in the next sections. This systematic approach will be used to scientifically ground the work related to the analysis of the current work, and advance the proposed TAPAF.

2.1 Research Methodology

The current survey results come from a structured literature review, which was aimed at identifying, assessing, and synthesizing literature on a range of topics including: ePrior Authorization (ePA), Interoperability in the

Healthcare field, Artificial Intelligence (AI), Explainable Artificial Intelligence (XAI), Multi-Agent Systems (MAS), and Healthcare Governance and regulatory compliance. A comprehensive literature search was conducted using IEEE Xplore, ACM Digital Library, SpringerLink, ScienceDirect, PubMed, Wiley Online Library and Google Scholar databases to cover a plethora of literature on healthcare informatics and computer science. The keywords used were: electronic prior authorization, prior authorization, HL7 FHIR, healthcare interoperability, artificial intelligence in healthcare, clinical decision support, explainable AI, multi-agent systems, healthcare automation, trustworthy AI, healthcare governance and FHIR APIs, and these terms were then used with Boolean operators like AND or OR. Inglorious studies were included, and studies published in the past five years (2020-25) were prioritized to consider the advancements in AI-powered interoperability in health care and intelligent prior authorization systems.

As a consequence of this, a selection of relevant and quality literature was identified for review in accordance with pre-defined inclusion and exclusion criteria related to the publications retrieved.

Inclusion Criteria

- Conference Proceedings, peer-reviewed Journal articles.
- Healthcare interoperability and prior authorization from an international standpoint, regulatory and authoritative technical reports.
- Expertise & study on the topic of electronic prior authorization (ePA), interoperability in health service or on HL7 FHIR-based data exchange.
- Study on artificial intelligence, explainable AI, clinical decision support systems, multi-agent systems in healthcare.
- Reporting on decision making, healthcare governance, regulatory compliance, privacy, security, trust or in-the-loop decision making.

Exclusion Criteria

This was the case with one or more publications that were picked up in several databases.

Publisher's notes, letters to the editor, book reviews and short letters (not technical).

- Non-English publications.
- Research typically not related to either healthcare-interoperability or intelligent clinical decision support or to prior authorization.
- Information published in publications lacking adequate methodological information and/or scientific evidence.
- Articles included only non-healthcare-specific domains that weren't relevant to the aims of this survey.

The number of publications retrieved in the first search in the selected databases: ~320. After removing duplicate studies and screening of title–abstracts, 138 studies were shortlisted for full-text evaluation. These studies were judged against a set of inclusion and exclusion criteria, and then 92 publications were selected for detailed analysis and synthesis for use as scientific foundations and guidance for the development of the proposed TAPAF for this survey.

2.2 Related Work

The introduction of Prior Authorization Support, Coverage Requirements Discovery, and Documentation Templates and Rules specifications allows the seamless exchange of clinical and administrative information between a provider and a payer, greatly improving the interoperability of providers and payers. At the same time, the development of automated extraction of clinical evidence, clinical necessity review, and intelligent clinical decision support using natural language processing and machine learning has helped minimize manual workloads in authorization processes [16]. Other recently explored approaches are distributed healthcare coordination obtained by multi-agent systems, which have proven to be more scalable than single-agent systems because of their distributed

efforts, adaptability to the process of organizing healthcare workflows, and integration with access to heterogeneous healthcare systems [17, 18]. The other one is trustworthy and explainable AI, which focuses on transparency, auditability, fairness, and human oversight to become the cornerstone of research with regard to the use of AI in critical well-being and healthcare scenarios [19]. However, most of the existing approaches focus on interoperability, AI support for decision-making, explainability, or autonomous agents isolated from the entire AI system or from a physical system [20]. Furthermore, more research has been done on interoperability with FHIR [21] and interoperability with FHIR as end-to-end orchestration of prior authorization [22] than in other areas like trust-aware interoperability, autonomous AI agents, explainability, governance [23], and regulatory compliance. This gap is an impetus to carry out a comprehensive review as detailed in this paper.

3.3 Table I. Summary of Related Literature

Reference	Research Focus	Methodology / Technology	Major Contributions	Limitations / Research Gap
[13]	SMART on FHIR interoperability	SMART on FHIR, OAuth 2.0, RESTful APIs	Introduced a standardized API-based framework for interoperable healthcare applications, enabling secure clinical data exchange and simplifying integration across EHR systems.	Focuses on application interoperability without addressing prior authorization workflows, AI-assisted decision-making, or autonomous orchestration.
[14]	FHIR-based Prior Authorization Support (PAS)	HL7 FHIR, X12 Integration, PAS Implementation Guide	Standardized electronic prior authorization transactions between providers and payers, improving interoperability and reducing manual data exchange during authorization requests.	Provides interoperability specifications but lacks intelligent workflow automation, explainability, and AI-driven medical necessity assessment.
[15]	Coverage Requirements Discovery (CRD) and Documentation Templates & Rules (DTR)	CDS Hooks, HL7 FHIR, Clinical Decision Support	Automated retrieval of payer documentation requirements and standardized supporting clinical documentation, thereby improving workflow efficiency and reducing administrative burden.	Limited to documentation support and does not include autonomous decision-making, trust management, or integrated AI agents.
[16]	Artificial intelligence for healthcare	Deep Learning, Machine Learning, Clinical Decision Support	Demonstrated how AI can improve clinical prediction, medical documentation, disease diagnosis, and healthcare decision support across multiple clinical domains.	Limited discussion on explainability, interoperability, regulatory compliance, and prior authorization applications.
[17]	Autonomous multi-agent systems	Distributed AI, Autonomous Agents, Agent Coordination	Established foundational principles for distributed intelligent agents capable of collaboration, task delegation, planning, and autonomous decision-	Generic framework without healthcare-specific interoperability, regulatory requirements, or payer-provider coordination.

			making in complex environments.	
[18]	Multi-agent healthcare workflow automation	Intelligent Agents, Healthcare Workflow Management	Demonstrated improved coordination, adaptive task allocation, and distributed clinical workflow management using collaborative software agents.	Limited integration with FHIR standards, explainable AI, healthcare governance, and prior authorization processes.
[19]	Trustworthy AI risk management	NIST AI Risk Management Framework, AI Governance	Defined core principles for trustworthy AI, including transparency, fairness, accountability, privacy, reliability, and continuous risk management throughout the AI lifecycle.	General governance framework without implementation guidance for healthcare-specific authorization workflows.
[20]	Ethical AI in healthcare	WHO AI Ethics Framework	Established recommendations for responsible AI adoption emphasizing patient safety, explainability, fairness, accountability, human oversight, and ethical governance.	Policy-oriented guidance lacking architectural frameworks for intelligent healthcare automation.
[21]	Electronic prior authorization regulation	CMS-0057-F Final Rule, FHIR APIs	Mandated standardized FHIR-enabled electronic prior authorization, improved interoperability, shortened response timelines, and enhanced transparency between providers and payers.	Regulatory initiative focuses on compliance rather than autonomous AI orchestration or intelligent clinical reasoning.
[22]	Foundation models for intelligent systems	Large Language Models, Foundation Models	Reviewed capabilities, opportunities, and risks of foundation models for reasoning, summarization, and knowledge-intensive AI applications across diverse domains.	Does not specifically address healthcare prior authorization, FHIR interoperability, explainability, or multi-agent clinical workflows.

3. EVOLUTION OF PRIOR AUTHORIZATION

The process has undergone a dramatic shift from a mostly manual administrative task to a more interoperable, automated, and digitized healthcare workflow. In the early days, payers found PA to be a means of trimming expenses and deciding on the medical necessity and coverage eligibility of the treatments that were being prescribed, much of which was paper-based, facsimile, and phone call dependent. A clinician would be overwhelmed sifting through mountains of patient data to document patient demographics, previous medical history, imaging, physician notes, previous documentation, and more, and then have their payment information pulled by a population health claimant at the payer end to ensure coverage was provided for that patient per the payer's standard claims criteria. This can entail a significant amount of time and effort, and could be challenging due to transcription mistakes, transcription delays, and multiple documentation interpretations of policies, as well as interpretation timing delays. When EHRs

started to become ubiquitous, various paper-based workflows shifted to new EHR/online prior authorization, and clinical evidence shifted from paper to electronic; all prior authorization requests and authorizations were now delivered electronically. While EPA has numerous benefits regarding paper savings and process efficiencies, many of the EPA systems were still considered "digital facades" of a paper process, rather than being intelligently designed as a decision support platform. In recent years, digital infrastructure has been developing rapidly, but those digital infrastructure providers often were unnecessarily doubling down on the problems of lots of data sources to support disparate and changing payers' requirements and multiple charges of documentation time and delays in authorization responses. Healthcare interoperability standards have been moving at a rapid pace, particularly HL7 FHIR, which is taking strides in helping providers, payers, labs, pharmacies, and CIS communicate via APIs and inter-provider interfaces. The transfer allowed for a shift toward automating access to clinical evidence and retrieval, toward policy standardization for assessment, and toward intelligent integration with workflows, which resulted in more transparent, efficient, and patient-friendly prior authorization processes.

Figure 1. Evolution of Prior Authorization Toward Intelligent Autonomous Orchestration

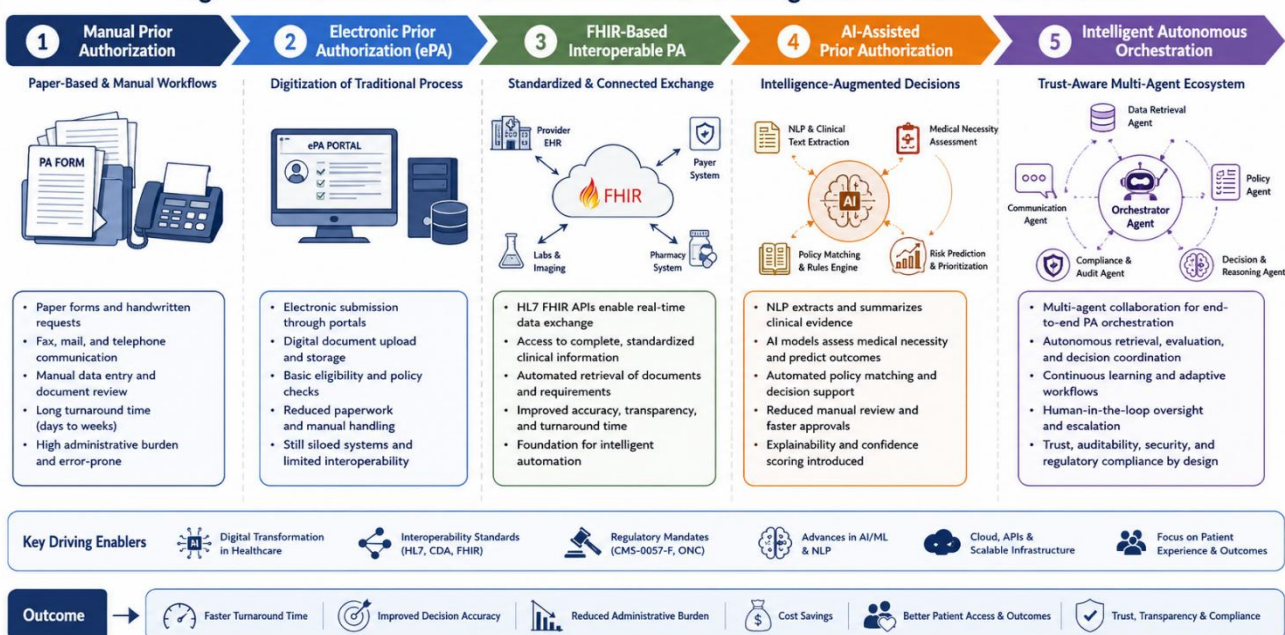


Figure 1. Evolution of Prior Authorization Toward Intelligent Autonomous Orchestration

Currently, the PA process is underway to be digitized, and with the assistance of AI, explainable decision-making, and distributed autonomous systems, this is changing to become intelligent orchestration. In present-day health practices, FHIR has grown to a large extent, as it allows for private and standardized exchange of a patient's health information among various types of input storage infrastructures involved in patient healthcare, from EHRs to payers, to client networking, computerization, to laboratory information systems, and so on [23-25]. Meanwhile, advancements in machine learning, NLP technology, and LLM technology have made it possible to automatically extract medical evidence from unstructured physician notes, predict medical necessity, summarize patient history, and match the various policies to make the administrative burden even lighter and generate authorizations with greater timeliness. More recently, the concept of multi-agent healthcare systems has started to offer a potential path in the creation of systems in which individual agents are specifically tasked with completing their own operations, such as data retrieval, compliance, and clinical evidence checks; producing explanations; and managing workflows to meet complex authorization constraints of a healthcare Payer. Unlike conventional workflow automation systems, these intelligent agents have the capacity to flexibly change their action plan based on evolving consumer needs or requirements, interact with endpoints in different healthcare infrastructures, and supply understandable statements when it comes to the types of recommendations they draw. However, as AI systems become more autonomous,

ensuring their explainability, fairness, auditability, cybersecurity, regulatory compliance, and constant human oversight throughout the authorization process are essential for having proper AI governance [26-28]. Technologies, such as trusted AI, connectionless sharing of healthcare data, and autonomous systems, were helping to facilitate the evolution from interoperability-enabled tools or focus and automation, towards seamless healthcare administration. The evolution will give the conceptual background to an autonomous Prior Authorization Orchestration Framework described in subsequent sections with the ability to operate in a trust-aware way.

4. HEALTHCARE INTEROPERABILITY FOUNDATIONS

As the healthcare information world has become digital, seamless and secure clinical and administrative data sharing is achievable only when presented in a standardized and efficient way, facilitated by an interoperable system in today's highly interdependent and heterogeneous healthcare information ecosystem. In the face of all these different systems, including Electronic Health Records and optimization systems, hospitals must communicate with each other in order to enable better patient care, reduce redundant paperwork, and facilitate care coordination. In order to improve the health of the patient, reduce the burden of administration and enhance efficiency, healthcare organizations are looking for an easy and effective way of communication, especially because of all the systems available, such as Electronic Health Records, Laboratory Information Systems, Pharmacy Information Systems, Radiology Information Systems, Insurance Information Systems, and finally Clinical Decision Support Systems. In a siloed healthcare environment with systems that are interfaced to each other with their own language, patient data is incomplete, clinically unproven data is repeated, decision time is delayed, and payer/Provider communication is inefficient. Developed by international standards organizations, these will produce a range of different standards for interoperability, including data models and messaging protocols, amongst other health data exchanges between systems, and a common terminology. One of the standards, Health Level Seven, allows for communication between the next generation of clinical systems through message- and document-based communications. The introduction of the Clinical Document Architecture (CDA) brought yet another capability to store clinical documents in a standard manner and to allow semantic interoperability via an XML-based approach to documents. However, they were soon followed by the introduction of the FHIR project, which developed yet another version of a more modern, resource-oriented representation of healthcare information, following the highly extensible healthcare information models, web technologies, and a RESTful approach augmented by representation of healthcare information in the JSON and XML formats [29]. The addition of FHIR has assisted in driving forward interoperability initiatives throughout the healthcare industry, permitting healthcare application integration, cloud healthcare solutions, clinical workflow standardization, and real-time data exchange. Moreover, in the electronic Prior Authorization space, the seamless and smooth transfer of patient information, guidelines, lab reports, previous drug usage, and clinical history is just as critical to enhancing efficiency and quality of Prior Authorizations. Therefore, besides being a technical must, interoperability represents a strategic lever for the ability to leverage AI in clinical workflow for automation, multiple agent interaction, clinical decision support tools that provide explanations, and regulatory compliance of regulatory systems in healthcare next-generation ecosystems (NXEs) [30].

4.1 Health Level Seven (HL7)

Health Level Seven (HL7) is an international protocol that has increasingly been adopted in the field of information exchange for healthcare delivery and has significantly contributed to the evolution of healthcare interoperability in the digital age. Established with the aim of facilitating the exchange of health care information across the different types of health care information systems, HL7 has created standardized messaging formats, data models, and vocabulary specifications as well as implementation guidelines that would enable hospital, laboratory, pharmacist, and insurance and/or health care application systems to work in harmony and securely. [27] In particular, for sending out clinical information such as admission and discharge information, order of medication or lab results, or appointment reminders, the older versions, more precisely, Version 2 (V2), were used de facto as the standard. Later, a more structured information model was introduced- the Reference Information Model (RIM), which is used to achieve greater semantic consistency across systems in the health industry. The HL7 V3 was more difficult to

implement and thus was not as successful as the HL7 V2. This preparation, however, helped pave the way towards universal interoperability in healthcare communication and clinical information exchange through a standardized communication protocol and shared information. Many legacy healthcare-related applications are continuing to utilize these standards and will be forming the conceptual framework for interoperability standards in the coming years. In electronic prior authorization, HL7 standards help ensure consistent communication between provider organizations and payer systems, enabling the smooth transfer of administrative and clinical information across different healthcare systems, ensuring that the data is consistent, secure, and accurate.

4.2 Clinical Document Architecture (CDA)

HL7 developed CDA to better align the meaning and structure of clinical documents exchanged within the healthcare environment. CDA is not decodable from a message; it is just a representation of a CDR, usually in the form of a 'discharge summary', 'consultation', 'diagnostic message', 'referral letter', 'operative note', continuity of care document', etc., and represents it in a hierarchical structure in XML [28]. Each CDA document is the set of metadata plus a single clinical document, allowing healthcare collection members to provide a complete clinical context, allowing for automated processing of the documents, and allowing for the ongoing exchange of clinical information. The standardized document structure and document record allow us to ensure the clinical information is available so it is easy to understand from any health system. CDA offers both a machine-readable and a human-readable version of healthcare information, making it widely popular for a variety of applications, such as health information exchanges (HIE), public health reporting, clinical quality measurement, and integration with electronic medical records (EMR). CDA remains valuable for document standardization and the exchange of structured clinical documents. However, its primary focus on document-level workflows makes it less suitable for real-time, resource-level data exchange required by modern healthcare applications. Retrieving individual clinical data elements from CDA documents can be inefficient and time-consuming, particularly in cloud-based, API-driven healthcare environments. In contrast, FHIR adopts a modular, resource-oriented architecture that enables flexible, real-time data exchange while maintaining compatibility with CDA where necessary. This combination of modularity, interoperability, and seamless integration allows FHIR to address the demands of distributed healthcare ecosystems, complementing and, in many scenarios, gradually replacing CDA as the preferred standard for healthcare interoperability [29].

4.3 Fast Healthcare Interoperability Resources (FHIR)

FHIR is the most recent version of the HL7 interoperability specification standard and has become the standard to adopt in modern healthcare interoperability use cases in just a few years. Each of the FHIR resources, like Patient, Practitioner, Observation, Condition, Encounter, Medication Request, Coverage, Claim, and Document Reference, is represented as a separate, independent but related resource to which access and use are made using secure Web standard tools, RESTful Web services, JSON, or XML. It is this architecture that simplifies and makes application integration much easier and convenient, and enables scalable, cloud-native, mobile healthcare environments. Furthermore, FHIR introduces numerous best practices of modern software development that will aid in building safe, modern, and pleasant-to-use health care programs that are interoperable, clinically consistent, and of high quality, as well as using an established web stack. Moreover, tools such as “SMART on FHIR” (SHC) and other HL7 Da Vinci projects, namely “Coverage Requirements Discovery (CRD)”, “Documentation Templates and Rules (DTR)” and “Prior Authorization Support (PAS)” are currently adding teeth to FHIR in its specific applications to certain aspects of the healthcare workflows, including ePA, clinical decision support and payer-provider interoperability [30-32]. They enable the automatic extraction of clinical information, a consistent assessment of policies, safe sharing and exchange of information, and intelligent coordination of workflows in disparate healthcare environments. FHIR has indeed become one of the most impactful among the interoperability standards in today's digital healthcare landscape, underpinning many uses including AI-powered healthcare automation, multi-agent clinical systems, and even trust-aware prior authorization orchestration.

5. REGULATORY LANDSCAPE

The adoption of digitalization now taking place in healthcare has elevated the importance of regulation and standards for interoperability, data privacy and security, and the appropriate and informed healthcare decisions. Securing and safeguarding health information sharing in healthcare and maintaining healthcare exchange and interoperability standards and/or patient privacy have been the basis for governments and standards organizations' broad regulations. Worrying is not just about the security of patient information; it is also about ensuring that clinical information is easily accessible, timely, and traceable, addressing questions of accountability and decisions impacted by AI, and complying with healthcare norms and regulations surrounding AI tools and usage. In the United States, HIPAA establishes the minimum standards that will protect ePHI via administrative, physical, and technical measures. The Health Information Technology for Economic and Clinical Health (HITECH) Act later encouraged the use of EHRs all across the country by offering incentives for meaningful use of electronic health technologies and tightened privacy and breach notification obligations [32, 33]. Other policies like the Office of the National Coordinator (ONC) interoperability standards, the 21st Century Cures Act, and more recent policies have also been introduced to curb information blocking and to standardize the exchange of healthcare information using enhanced technologies including HL7 FHIR. Beyond these healthcare regulatory requirements, numerous other rules of engagement such as the NIST AI Risk Management Framework and the WHO guidance on ethical AI highlight key concepts of fairness, explainability, transparency, human oversight, and accountability when it comes to instances where AI has an influence on healthcare decisions [34]. Together, all of these will enable the effective use of AI technologies, the effective acceptance of AI, and the effective construction of a “trust-aware” platform that can support autonomous, personalized, and custom-built health experiences and help to meet the increasing demands for digital health, setting some standards for the “trust-able” deployment of autonomous, mutual previous authorization (APA) systems [35].

Table II. Major Regulatory Frameworks Supporting Digital Healthcare and Prior Authorization

Reference	Regulation / Framework	Primary Objective	Key Provisions	Relevance to Prior Authorization and AI
[31]	Health Insurance Portability and Accountability Act (HIPAA)	Protects the confidentiality, integrity, and availability of electronic Protected Health Information (ePHI).	Establishes the Privacy Rule, Security Rule, Breach Notification Rule, access controls, encryption requirements, audit logs, and secure healthcare data exchange.	Provides the security and privacy foundation for AI-enabled prior authorization systems handling sensitive patient information across interoperable healthcare networks.
[32]	Health Information Technology for Economic and Clinical Health (HITECH) Act	Accelerate nationwide adoption of Electronic Health Records (EHRs) and strengthen healthcare interoperability.	Introduced Meaningful Use incentives, enhanced HIPAA enforcement, strengthened breach reporting, promoted electronic health information exchange, and encouraged digital healthcare transformation.	Facilitates availability of structured electronic clinical data required for automated prior authorization and AI-assisted clinical decision support.
[33]	21st Century Cures Act & ONC	Promote patient data accessibility and eliminate information blocking	Mandates standardized APIs, encourages FHIR adoption, enhances patient access to	Enables seamless retrieval of clinical evidence and supports interoperable AI-driven prior

	Interoperability Rule	across healthcare organizations.	health records, and improves interoperability among healthcare stakeholders.	authorization workflows using standardized healthcare APIs.
[34]	CMS-0057-F Final Rule	Modernizes electronic prior authorization through standardized interoperability and improved payer-provider communication.	Mandates FHIR-based electronic prior authorization APIs, requires faster payer response timelines, standardized documentation exchange, denial reason reporting, and public API availability.	Serves as the principal regulatory driver for intelligent prior authorization systems by establishing standardized interoperable workflows across healthcare ecosystems.
[35]	NIST AI Risk Management Framework (AI RMF) & WHO Guidance on Ethics and Governance of AI for Health	Promote trustworthy, transparent, ethical, and accountable AI deployment in healthcare.	Defines principles including explainability, fairness, transparency, privacy, accountability, risk management, human oversight, continuous monitoring, and lifecycle governance.	Establishes governance principles essential for trust-aware autonomous prior authorization systems employing AI agents and explainable clinical decision support.

6. ARTIFICIAL INTELLIGENCE IN PRIOR AUTHORIZATION

Over the last few years, Artificial Intelligence (AI) has been reshaping the nature of prior authorization (PA) from a mostly administrative and document-heavy process to one that is more intelligent and based on real data, and enhances operational efficiency, clinical accuracy, and decision consistency. In normal PA processes, the healthcare providers would be responsible for gathering patient information from various sources, including documentation from a physician, lab reports, imaging reports, medications, instructions for care, etc., and then submit an authorization claim to the payers. Multiple steps may require a substantial amount of material, can involve prolonged steps with potential human errors, and may include long wait times that may impact patient care [36]. Machine learning (ML), NLP, Deep Learning, and Large Language Models (LLMs) allow for better processing of structured and unstructured clinical data in the hands of artificial systems. Medical technologies can leverage AI to automatically pull out pertinent medical proof from EHRs, summarize lengthy physician contacts, find missing clinical data, guarantee the accuracy of diagnostic codes, and assess medical necessity for payers' coverage policies. Providers will also be able to use predictive analytics to foresee requests for authorization that could be granted or rejected prior to being made, allowing them to help streamline a slew of incomplete documentation problems before they occur. Modern AI systems can also perform intelligent classification of documents, semantic search, retrieval of clinical evidence, and automated matching of policies, which decreases manual work and enhances consistency of the authorization process. Moreover, HL7 FHIR allows AI models to draw on the most relevant and accurate data to make informed recommendations by enabling real-time data exchange from multiple distributed clinical systems. Also, AI-powered methods give interpretable reasons for the recommendations generated, boosting the confidence of the clinician and ensuring regulatory compliance. While each of these AI capabilities individually produces value, the blend of all together is enhanced workflow and reduced administrative workload, reduced processing time, and an increased level of provider-payer collaboration; that's the basis of the next-generation intelligent Prior Authorization AI systems Figure 2 [37].

Prior utilization systems with AI, with visualization of AI interactions with the FHIR data retrieval logic, medical necessity check, explainable AI, and clinical review.




Figure 2. AI-Enabled Prior Authorization Workflow

AI in prior authorization is moving beyond a 'predictive' solution and towards intelligent and collaborative intelligence. More recent solutions utilize more industry-specific AI elements specifically designed for certain tasks, such as understanding healthcare payers' policies, understanding clinical evidence, validating documentation, and coaching the healthcare provider on healthcare-specific AI-recommended alternatives, using this additional industry-specific AI to enable the generative AI to appreciate the business setting it operates in. These systems are designed to work alongside clinicians, act as assistants to help with key skills, know how to process vast amounts of clinical information quickly, have a degree of intelligence and transparency in their operation, and be held accountable for decisions made. Explainable AI, confidence scoring capabilities, and audit logging and human-in-the-loop review capabilities help ensure automated recommendations are interpretable, clinically justified, and meaningful, even in high-stakes healthcare environments. With the continuation of the development of interoperability standards, such as HL7 FHIR, AI systems will increasingly be aware of context, adaptive, and collaborative, promoting seamless communication and collaboration across various segments of the health care system, including health care providers, health information technology systems, and insurance companies. The innovations have paved the way for new, intuitive prior authorizations automated by intelligence, regulation of adherence, handling ethics, and the centrality of outcomes in healthcare.

7. MULTI-AGENT HEALTHCARE SYSTEMS

To deal with complex workflows in healthcare systems, ranging from clinical to administrative, Multi-Agent Systems (MAS) have become an alluring system for intelligent, distributed, adaptive applications, known as Intelligent Healthcare Applications (IHA). A multi-agent system consists of some autonomous software agents operating in an autonomous manner and operating in an individually acting manner; that is, they can act independently to sense the environment, make decisions, communicate with other agents, execute the desired task, and operate with some degree

of autonomy [38, 39]. As opposed to most centralized software architectures, MAS assigns responsibilities between the agents, each of which is specialized in some particular task, making it more scalable, fault-tolerant, flexible, and responsive in the dynamic and changing healthcare domain. Agents may be customized to perform specific tasks in the evolving healthcare climate, like collecting patient information, gathering clinical details, interpreting a policy's rules, verifying eligibility, reviewing medical necessity, arranging appointments, allowing claims processing, and coordinating task workflow. They communicate and interoperate with each other via standardized communication protocols as well as standard healthcare interoperability standards to ensure smooth integration [40]. Further, the application of FHIR by HL7 makes it easier for agents to fetch, share, and modify the clinical data in real-time via RESTful APIs, thereby enabling them to obtain standardized access to healthcare resources. In line with this, emerging AI technology, particularly with the implementation of LLMs, has given healthcare agents a new arsenal to enhance their reasoning capabilities, grasp context, provide summaries, locate specific information in clinical text, and lead agents through adaptive decision-making processes. Furthermore, explainable AI, confidence measurement, audit tracking, and human-in-the-loop validation and assurance validate and assure the behavior of the autonomous agents within the framework of ethics and regulations [41]. When referred for prior authorization, MAS can split the process into multiple parts and have dedicated agents to document clinical information, coordinate with payers' coverage guidelines, establish compliance, explain the results to providers, and orchestrate the process. In terms of technology enablers, a next-generation AI-based prior authorization system can be achieved by employing multi-agent systems, which can handle distributed intelligence and thus greatly reduce the administrative burden, improve interoperability and scalability, and ensure trustworthy health care automation [42].

Table III. Functional Roles of Multi-Agent Systems in AI-Enabled Prior Authorization

Agent Type	Primary Responsibility	Input Data	Output / Action	Benefits
Clinical Data Retrieval Agent	Retrieves structured and unstructured patient information from EHRs and connected clinical systems.	Patient demographics, diagnoses, laboratory reports, physician notes, imaging records, medications	Standardized clinical dataset using HL7 FHIR resources for downstream processing.	Eliminates manual data collection, improves interoperability, and accelerates evidence retrieval [39].
Medical Evidence Agent	Extracts clinically relevant evidence supporting the requested treatment or procedure using AI-based document understanding.	Clinical notes, discharge summaries, laboratory reports, medical history	Structured clinical evidence linked to authorization request.	Improves completeness of documentation and reduces missing supporting evidence [40].
Policy Evaluation Agent	Interprets payer-specific coverage policies and eligibility criteria for the requested healthcare service.	Insurance policies, clinical guidelines, authorization rules	Policy compliance report and eligibility assessment.	Standardizes policy interpretation while minimizing manual review inconsistencies [40].
Medical Necessity Assessment Agent	Evaluates whether the requested intervention satisfies evidence-based clinical and payer requirements using AI models.	Clinical evidence, diagnosis codes, payer policies	Medical necessity recommendation with confidence score.	Supports faster and more consistent authorization decisions while improving clinical accuracy [41].

Workflow Orchestration Agent	Coordinates interactions among multiple specialized agents and manages task sequencing throughout the authorization lifecycle.	Outputs from collaborating agents, workflow status	Optimized authorization workflow and task scheduling.	Enhances scalability, coordination, fault tolerance, and process efficiency in distributed healthcare environments [42].
Explainability Agent	Generates human-readable explanations describing how AI recommendations were derived.	AI predictions, supporting evidence, policy references	Transparent decision rationale and supporting clinical evidence.	Improves clinician trust, regulatory compliance, and explainability of AI-assisted decisions [41].
Compliance and Audit Agent	Continuously monitors regulatory compliance, records system activities, and maintains immutable audit logs.	Workflow events, authorization decisions, system logs	Compliance reports, audit trail, traceability records.	Supports HIPAA, CMS, and organizational governance requirements while improving accountability [42].
Human Oversight Agent	Enables clinician or utilization reviewer intervention for high-risk, ambiguous, or exceptional authorization requests.	AI recommendations, confidence scores, explanation reports	Final authorization decision or workflow escalation.	Ensures human-in-the-loop governance, reduces AI risks, and maintains clinical accountability [42].

8. EXPLAINABLE AI FOR MEDICAL DECISION SUPPORT

AI has dramatically improved clinical decision support capabilities, which include automated diagnosis, predictive analytics, recommendations for treatment, and optimization of administrative workflow. As machine learning and deep learning algorithms become more sophisticated, however, the lack of transparency, interpretability, and, especially, clinical accountability has raised concerns about their use in critical treatment situations in healthcare. For many advanced AI models, predictive or recommended activities are "black-box" models that are difficult for clinicians to decipher why the model makes these predictions or recommendations given the information provided. A lack of transparency can lead to a decrease in physician trust, complicating regulatory acceptance and validation of AI-assisted medical decisions. Explainable Artificial Intelligence helps to overcome these challenges by introducing interpretable reasoning to allow healthcare professionals to understand how each provider's factors for a patient, clinical evidence, laboratory results, imaging reports, or payer policies contribute to the AI-generated results and insights. Feature importance analysis, SHAP, LIME, attention visualization, rule-based reasoning, and counterfactual explanations enable the clinician to validate their recommendations, as well as allowing them to find any biases or inconsistencies. Explainable AI can assist with clear medical necessity evaluations, policy alignment, and validation of medical evidence, providing supportive evidence as well as confidence levels and decision justifications. Moreover, the use of explainability along with audit logging, uncertainty estimation, and human-in-the-loop review builds trust, helps satisfy regulatory requirements, and enables the use of AI as a decision support tool and not a replacement for clinical expertise.

Table IV. Explainable AI Techniques for Medical Decision Support

XAI Technique	Working Principle	Healthcare Applications	Advantages	Limitations
Feature Importance Analysis	Quantifies the contribution of individual input variables toward model predictions.	Disease prediction, risk stratification, clinical outcome analysis, prior authorization.	Simple to understand, highlights influential clinical features, improves clinician confidence.	Does not fully explain complex feature interactions or nonlinear relationships.
SHAP (Shapley Additive Explanations)	Computes the contribution of each feature using cooperative game theory to explain individual predictions.	Medical necessity assessment, treatment recommendation, diagnostic prediction, clinical decision support.	Highly interpretable, consistent explanations, supports both local and global model interpretation.	Computationally expensive for large datasets and complex deep learning models.
LIME (Local Interpretable Model-Agnostic Explanations)	Builds an interpretable surrogate model to explain predictions around a specific data instance.	Patient-level diagnosis explanation, authorization decision support, clinical risk assessment.	Model-independent, easy to visualize, provides localized explanations.	Local explanations may not accurately represent overall model behavior.
Attention Visualization	Highlights portions of clinical text or medical images that contribute most to model decisions.	Clinical Natural Language Processing (NLP), radiology, pathology, medical document summarization.	Improves transparency of transformer-based models and facilitates evidence verification.	Attention weights do not always correspond to true causal reasoning.
Rule-Based Explanation Systems	Generates human-readable explanations using predefined clinical rules and expert knowledge.	Clinical guideline adherence, payer policy interpretation, medical necessity validation.	Highly transparent, clinically interpretable, supports regulatory compliance.	Difficult to maintain as clinical guidelines and payer policies evolve.
Counterfactual Explanations	Identifies minimal changes required to alter a model's prediction or recommendation.	Treatment optimization, prior authorization appeals, personalized clinical recommendations.	Provides actionable insights and supports clinician understanding of decision boundaries.	May generate clinically unrealistic or impractical scenarios if not properly constrained.
Confidence Scoring & Uncertainty Estimation	Quantifies the certainty associated with AI-generated predictions.	High-risk diagnosis, authorization recommendations, clinical decision support.	Helps clinicians identify uncertain cases requiring manual review and human oversight.	Confidence scores may not always correlate with prediction accuracy.
Audit Trail and Decision Logging	Records AI inputs, reasoning process, intermediate outputs, and final	Regulatory compliance, clinical governance, prior authorization	Improves accountability, reproducibility, and compliance with	Increases storage requirements and requires robust governance

	recommendations for traceability.	review, healthcare auditing.	healthcare regulations.	mechanisms for log management.
--	-----------------------------------	------------------------------	-------------------------	--------------------------------

9TRUST-AWARE AUTONOMOUS PRIOR AUTHORIZATION FRAMEWORK (TAPAF)

This paper extends the literature review on healthcare interoperability frameworks, AI technologies, explainable AI technologies, multi-agent systems, and regulatory frameworks to develop a Trust-Aware Autonomous Prior Authorization Framework that tackles current shortcomings of electronic authorization systems. Unlike conventional frameworks that primarily focus on the digitalization of manual processes, the proposed framework is built upon a layered architecture that integrates HL7 FHIR-based interoperability, intelligent business agents, and semantically enriched decision support. It adopts a human-in-the-loop approach in which AI agents augment clinical decision-making by providing transparent, explainable, and context-aware recommendations while ensuring that final decisions remain under the authority of qualified human professionals rather than being replaced by autonomous AI systems. In essence, an FHIR-enabled interoperability layer integrates and combines readily available clinical information, pervasively collected by the Electronic Health Records (EHR), the Laboratory Information System (LIS), and other payer systems, clinical repositories, and pharmacology and imaging technologies such as the Radiology Information System (RIS), all accessible by secure APIs, a standardized, real-time, and uniform way of accessing patient information. Subject matter users above this layer will execute jobs specific to their field of service, including clinical data retrieval, medical evidence extraction, application of payer policies, performing documentation validation, medical necessity reviews on documentation, orchestrating workflows, performing compliance reviews, and creating explanations for clinical work done through autonomous agents. Making distributed intelligence without increasing workflow and administration complexity is about an interaction of these collaborative agents through the same interface. These agents then go through a ‘trust layer’ in the middle, which consists of AI explainability components, risk detection tools, confidence score, policy review, audit logging, ethics checks, and harmonization. It is important to note that the framework doesn't permit any decision without human input; in fact, it establishes a necessary step of review by clinicians or by utilization management professionals, within which they shall review the AI-generated recommendations, supportive evidence, and confidence scores prior to making the final decision on authorization. Feedback on the outcomes of the authorizations is then incorporated into a learning and governance module to further facilitate improvements to the models of the AI; to better inform and improve the quality of payer policy; to help with ongoing compliance with policies and regulations; and to further guide quality improvement decisions. Therefore, the proposed solution will need to establish an intelligent, interoperable, scalable, and trusted environment to offer transparent, accountable, and patient-centric prior authorization systems, while respecting evolving healthcare regulations and ethics regarding AI adoption.

Figure 3. Proposed Conceptual Trust-Aware Autonomous Prior Authorization Framework

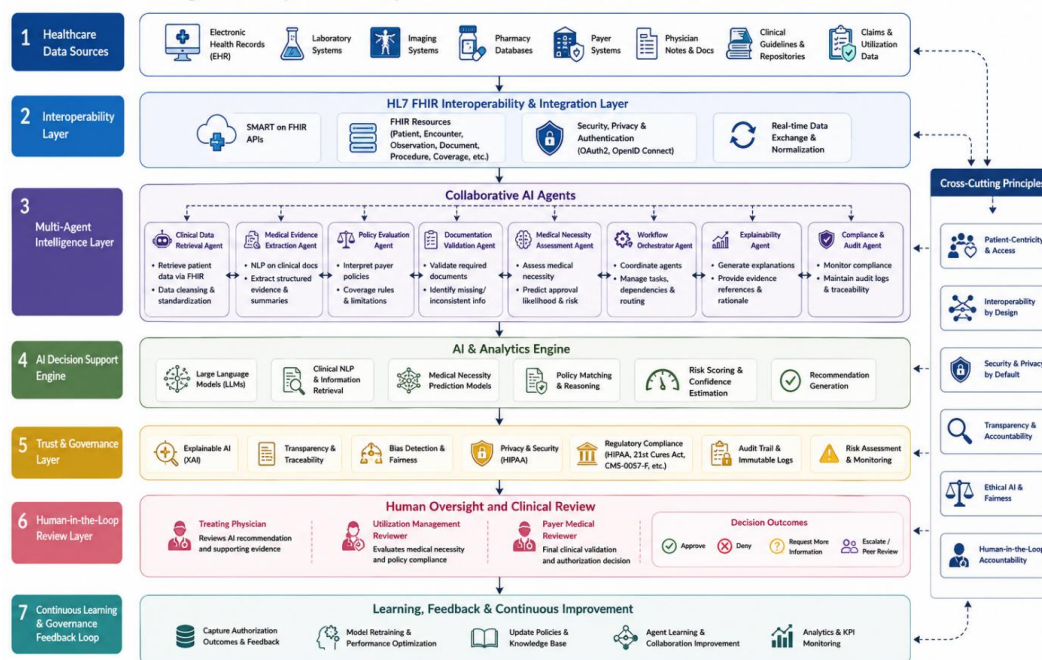


Figure 3. Proposed Trust-Aware Autonomous Prior Authorization Framework (TAPAF)

The proposed framework has seven interwoven layers dedicated to different parts of the orchestration process for prior authorization. The Healthcare Data Sources Layer is the gateway that collects structured and unstructured data from EHRs, labs, provider documentation, evidence-based clinical guidance, imaging repositories, pharmacy databases, and payer databases. The HL7 FHIR Interoperability Layer makes all these various bits of information well-formed so that they can be exchanged securely via API with other diverse clinical systems. The Multi-Agent Intelligence Layer's numerous specialized autonomous agents automatically process the standardized data files, independently conduct clinical evidence retrieval, document validation, interpretation of payer policies, medical necessity assessment, compliance checks, process synchronization, and explanation generation while communicating with each other and sharing intermediate results to enhance the quality of decision-making. The outputs are consolidated into a single file and analyzed in the AI Decision Support Engine with the help of large language models, natural language processing, predictive analytics, and policy matching algorithms for generating authentic and evidence-based authorization recommendations, along with confidence scores. All recommendations generated by the AI Decision Support Engine are routed through the Trust and Governance Layer for policy validation, regulatory compliance, and governance verification before being presented to human reviewers. The Human-in-the-Loop Review Layer enables medical professionals or utilization management reviewers to validate the outputs of the AI system, to make exceptions where needed, and to ensure clinical safety. Finally, the Continuous Learning and Governance Layer collects results of authorizations, regulatory changes, audit results, and user feedback, which will be used to optimize AI models, enhance agent-to-agent interactions, upgrade policy knowledge, and further boost the overall capabilities, reliability, and trustworthiness of the existing prior authorization ecosystem. The multi-layered design comprises intelligent automation combined with human skills, fitting for future AI-driven healthcare settings.

11. CONCLUSION

One of the most time-consuming administrative tasks in healthcare is prior authorization, which can lead to delays in treatment, cost issues, and administrative burdens. The aim of this review was to explore how prior authorization systems have transformed from manual processes to interoperable, AI-powered, and increasingly automated workflows, and the central role played by health information interoperability standards like HL7, CDA,

and FHIR in facilitating standardized clinical data sharing. The paper summarized the recent strides in artificial intelligence, explainable AI, and multi-agent healthcare systems, illustrating their applications in enhancing clinical evidence retrieval, medical necessity determination, orchestrating workflow, and decision support, while fostering transparency and accountability. Furthermore, the regulations with regard to HIPAA, HITECH, the 21st Century Cures Act, the CMS-0057-F frameworks, and ethical and regulatory issues around the implementation of AI have been correlated and emphasized the importance of interoperability, patient privacy, ethical challenges of AI implementation, and regulatory compliance in today's healthcare landscape. On the basis of synthesized literature, the present review suggested a TAPAF that encompasses all these features in a single orchestration framework. While prior approaches have mainly focused on one or some components of prior authorization, the proposed approach highlights the elements of collaborative intelligence, transparency, trust, and continuous governance as integral design principles for next-generation prior authorization systems. Despite ongoing technical, ethical, and regulatory hurdles, the integration of interoperable healthcare standards, the trustworthiness of AI, and distributed multi-agent systems offers great promise for addressing and revolutionizing the current challenges and opportunities associated with prior authorization.

REFERENCES

1. Busis, N. A., Khokhar, B., & Callaghan, B. C. (2024). Streamlining prior authorization to improve care. *JAMA Neurology*, 81(1), 5-6.
2. Hu, J. C., Hutchings, K., Jalali, A., Kapadia, S. N., Bao, Y., & Underhill, K. (2025). State Laws Prohibiting Prior Authorization for Medications for Opioid Use Disorder in Private Insurance, 2015–2023. *Health Affairs (Project Hope)*, 44(11), 1369.
3. Mello, M. M., Trotsyuk, A. A., Mahamadou, A. J. D., & Char, D. (2026). The AI Arms Race in Health Insurance Utilization Review: Promises of Efficiency and Risks of Supercharged Flaws. Article examines the artificial intelligence arms race in health insurance utilization review. *Health Affairs*, 45(1), 6-13.
4. Lavoie-Gagne, O., Woo, J. J., Williams, R. J., Nwachukwu, B. U., Kunze, K. N., & Ramkumar, P. N. (2025). Artificial intelligence as a tool to mitigate administrative burden, optimize billing, reduce insurance-and credentialing-related expenses, and improve quality assurance within health care systems. *Arthroscopy*, 41(8), 3270-3275.
5. Murphy, J., Beauchamp, N., Sun, K. J., Lau, B. D., Wilson, R. F., Lobner, K., ... & Johnson, P. T. (2025). Adverse Effects of Health Plan Prior Authorization on Clinical Effectiveness and Patient Outcomes: A Systematic Review. *The American journal of medicine*.
6. Smalley, W. E., Griffin, M. R., Fought, R. L., Sullivan, L., & Ray, W. A. (1995). Effect of a prior-authorization requirement on the use of nonsteroidal anti-inflammatory drugs by Medicaid patients. *New England Journal of Medicine*, 332(24), 1612-1617.
7. Duda, S. N., Kennedy, N., Conway, D., Cheng, A. C., Nguyen, V., Zayas-Cabán, T., & Harris, P. A. (2022). HL7 FHIR-based tools and initiatives to support clinical research: a scoping review. *Journal of the American Medical Informatics Association*, 29(9), 1642-1653.
8. de Freitas, I. C. (2024). *Promoting the Integration of FHIR into an existing Health Information System* (Master's thesis, Universidade do Porto (Portugal)).
9. Rajpurkar, P., Chen, E., Banerjee, O., & Topol, E. J. (2022). AI in health and medicine. *Nature Medicine*, 28(1), 31-38.
10. Wooldridge, M. (2009). *An introduction to multiagent systems*. John Wiley & sons.
11. RISK, C. (2021). Risk management framework. *CORPORATE REPORT*.
12. World Health Organization. (2021). Ethics and governance of artificial intelligence for health: WHO guidance. World Health Organization.
13. Mandel, J. C., Kreda, D. A., Mandl, K. D., Kohane, I. S., & Ramoni, R. B. (2016). SMART on FHIR: a standards-based, interoperable apps platform for electronic health records. *Journal of the American Medical Informatics Association*, 23(5), 899-908.
14. Athri, A., & EVP, S. Architecting the Future of Prior Authorization.
15. MacKinnon, N., & Kumar, R. (2001). Prior authorization programs: a critical review of the literature. *Journal of Managed Care Pharmacy*, 7(4), 297-303.
16. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., ... & Dean, J. (2019). A guide to deep learning in healthcare. *Nature Medicine*, 25(1), 24-29.
17. Mahmoud, M. S. (2020). Multiagent systems: introduction and coordination control. CRC Press.
18. Iqbal, S., Altaf, W., Aslam, M., Mahmood, W., & Khan, M. U. G. (2016). Application of intelligent agents in health-care. *Artificial Intelligence Review*, 46(1), 83-112.
19. AI, N. (2024). Artificial intelligence risk management framework: Generative artificial intelligence profile. *NIST Trustworthy and Responsible AI, Gaithersburg, MD, USA*.

20. World Health Organization. (2023). *Regulatory considerations on artificial intelligence for health*. World Health Organization.
21. Kumari, P. (2021). How Information Blocking And The 21st Century Cures Act Impact Interoperability In Health Information Systems (Doctoral dissertation, Penn State Harrisburg, Capital College).
22. Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., ... & Liang, P. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
23. PORTAL, C., & BLOG, R. O. The Blueprint for Prior Authorization in Interventional Pain Management: 2026 Edition.
24. Chen, W. C., Carpenter, C., Sidiqi, B., Pattison, A. J., Hwang, J., Pappas, D., & Potters, L. (2025). Integrating prior authorization into clinical workflows for care access and practitioner experience. *JAMA Network Open*, 8(12), e2549093.
25. Saberi, M. A., McHeick, H., & Adda, M. (2025). From data silos to health records without borders: a systematic survey on patient-centered data interoperability. *Information*, 16(2), 106.
26. Dixon, B. E., Holmgren, A. J., Adler-Milstein, J., & Grannis, S. J. (2022). Health information exchange and interoperability. In *Clinical Informatics study guide: text and review* (pp. 203-219). Cham: Springer International Publishing.
27. Benson, T., & Grieve, G. (2016). *Principles of health interoperability* (pp. 293-324). Cham, Switzerland: Springer.
28. Dolin, R. H., Alschuler, L., Boyer, S., Beebe, C., Behlen, F. M., Biron, P. V., & Shabo, A. (2006). HL7 clinical document architecture, release 2. *Journal of the American Medical Informatics Association*, 13(1), 30-39.
29. Bender, D., & Sartipi, K. (2013, June). HL7 FHIR: An Agile and RESTful approach to healthcare information exchange. In *Proceedings of the 26th IEEE International Symposium on Computer-Based Medical Systems* (pp. 326-331). IEEE.
30. Mandl, K. D., Mandel, J. C., Murphy, S. N., Bernstam, E. V., Ramoni, R. L., Kreda, D. A., ... & Kohane, I. S. (2012). The SMART Platform: early experience enabling substitutable applications for electronic health records. *Journal of the American Medical Informatics Association*, 19(4), 597-603.
31. Centers for Disease Control and Prevention. (2003). HIPAA privacy rule and public health: guidance from CDC and the US Department of Health and Human Services. *MMWR. Morbidity and Mortality Weekly Report*, 52, 1.
32. Burde, H. (2011). The HITECH Act: an overview. *AMA Journal of Ethics*, 13(3), 172-175.
33. Kyle, M. A., & Song, Z. (2023). The consequences and future of prior-authorization reform. *The New England journal of medicine*, 389(4), 291.
34. Johnson, P. T., Lau, B., Conway, S. J., Riley, L., & Hill, P. (2022). Evidence-based prior authorization in the electronic health record: a better path forward. *Journal of the American College of Radiology*, 19(5), 652-654.
35. Wang, Y., Song, Y., Wang, Y., YU, L., & Wang, J. (2024). Ethics and governance of artificial intelligence for health: guidance on large multi-modal models. *Chinese Medical Ethics*, 1001-1022.
36. Topol, E. (2019). *Deep medicine: how artificial intelligence can make healthcare human again*. Hachette UK.
37. Obermeyer, Z., & Emanuel, E. J. (2016). Predicting the future—big data, machine learning, and clinical medicine. *The New England journal of medicine*, 375(13), 1216.
38. Singhal, K., Azizi, S., Tu, T., Mahdavi, S. S., Wei, J., Chung, H. W., ... & Natarajan, V. (2023). Publisher correction: Large language models encode clinical knowledge. *Nature*, 620(7973), E19.
39. Wooldridge, M., & Jennings, N. R. (1995). Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2), 115-152.
40. Abd Elaziz, M., Al-qaness, M. A., Dahou, A., Al-Betar, M. A., Mohamed, M. M., El-Shinawi, M., ... & Ewees, A. A. (2024). Digital twins in healthcare: Applications, technologies, simulations, and future trends. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 14(6), e1559.
41. Gunning, D., & Aha, D. (2019). DARPA's explainable artificial intelligence (XAI) program. *AI magazine*, 40(2), 44-58.
42. Russell, S., & Norvig, P. (2020). Artificial intelligence: a modern approach, 4/E.