Conceptual approaches to data transmission for AI-assisted patient assessment*

Yaroslav Kotov[1,†], Evhenia Yavorska[1,*,†], Bohdan Yavorskiy[1,†], Oksana Dozorska[1,†] and Vasyl Yatskiv[2,†]

[1] Ternopil Ivan Puluj National Technical University, Ruska str., 56 46001 Ternopil, Ukraine
[2] West Ukrainian National University, Lvivska str. 11, Ternopil, 46009, Ukraine

Abstract

Artificial intelligence (AI) is commonly utilized to enhance clinical decision-making by processing patient data and health information and generating evaluative findings. Nevertheless, communication of the information that patients provide to AI systems, as well as its conveyance to medical practitioners, calls for adherence to high standards to facilitate data integrity, security, and interoperability. This study proposes a conceptual model for AI-enhanced patient evaluation, specifying how subjective patient information can be gathered and transmitted through the HL7 FHIR standard in conjunction with associated health data workflows. The suggested strategy outlines an end-to-end data exchange—from patient input via a safe, standardized delivery pipeline to an AI model and ultimately to the medical practitioner—emphasizing how interoperability frameworks (like FHIR) and effective security practices (encryption, authentication, access control) facilitate unbreakable integration. We discuss recent related research in both health data exchange and AI, suggest an approach based on structured data exchange (e.g., FHIR resources for observations and symptoms), and consider the practical implications, challenges, and likely results of putting such a system into practice. The findings highlight the value of standard data formats and security as keys to unleashing AI's potential for patient evaluation. We finish with directions for the future, such as the requirement for pilot launches and additional research on data transmission optimization and ensuring trust in AI-supported clinical processes.

Keywords

AI-assisted patient assessment, data transmission, interoperability, HL7 FHIR, healthcare data standards, telemedicine

1. Introduction

In modern medicine, artificial intelligence programs stand to augment patient assessment by analyzing vast amounts of clinical information and providing decision support. Realizing AI-supported care involves effective communication of patient data—frequently subjective symptoms or patient-reported information—to the AI system and back to providers in an understandable form. Telehealth and e-health technologies today enable patients to transmit symptoms or health metrics via smartphones and wearables, which can be imported into the EHR alongside clinical data. AI algorithms can then mine this integrated dataset for patterns or risk factors to aid diagnosis and triage.

Although promising, facilitating effective, secure, and interoperable exchange is still challenging: patient-generated data are heterogeneous (text reports, survey responses, device readings) and may originate outside clinical settings, while proprietary or ad hoc exchanges negate integration with clinical and AI platforms. Although standards like HL7 FHIR exist, uptake within AI environments has been patchy. Robust frameworks are therefore needed to standardize data exchange between clinicians, AI, and patients, preserving semantic integrity and in adherence to

* CITI'2025: 3rd International Workshop on Computer Information Technologies in Industry 4.0, June 11–12, 2025, Ternopil, Ukraine
[1*] Corresponding author.
† These authors contributed equally.
kotov20010731@gmail.com (Ya. Kotov); yavorska_eb@yahoo.com (E. Yavorska); yavorskiy_b@tntu.edu.ua (B. Yavorskiy); oksana4elka@gmail.com (O. Dozorska); vy@wunu.edu.ua (V. Yatskiv)
0000-0001-6341-1710 (E. Yavorska); 0000-0003-4215-1176 (B. Yavorskiy); 0000-0001-7053-863X (O. Dozorska); 0000-0001-9778-6625 (V. Yatskiv)

privacy laws (e.g., HIPAA, GDPR). HL7 FHIR—a set of modular resources (Patient, Observation, Condition) represented in JSON/XML and accessed via a RESTful API—allows for both technical and semantic interoperability. Its growing uptake enables heterogeneous systems to be incorporated in AI-driven workflows, with standardised terminologies (SNOMED CT, LOINC, ICD) enabling data to be machine-readable and clinically meaningful.

This paper offers a theoretical framework for data communication in the scenario of AI-assisted patient evaluation, with HL7 FHIR serving as a means to structure and safely convey patient-generated data to an AI system, which subsequently provides actionable guidance to healthcare providers.

2. Related works

Researchers and practitioners have recognized the importance of interoperability and data standards in applying AI to clinical care. A recent scoping review by Balch et al. [1] surveyed 39 machine learning–enabled clinical information systems and found that many leverage FHIR to standardize data exchange for clinical decision support and analytics. These systems range from clinical decision support tools embedded in EHRs to standalone analytic platforms, illustrating FHIR's versatility in various AI applications. The review noted that while FHIR provides a useful framework for data exchange, implementations vary, and best practices for applying FHIR in AI contexts are still being refined. It also reinforced that common communication standards like FHIR help maintain semantic integrity of data as it moves between actors, a crucial factor for trustworthy AI outputs.

Interoperability pipelines for preparing healthcare data for AI have been proposed to tackle the challenges of heterogeneous data sources. Williams et al. [2] introduced a FHIR-based Data Harmonization Pipeline (FHIR-DHP) that transforms raw hospital records into a harmonized, AI-friendly format. Their workflow involved querying hospital databases, mapping data to FHIR resources, validating syntax, and exporting to an AI-ready dataset. The pipeline demonstrated that using a common FHIR data model can facilitate scalability and multi-institutional data sharing for machine learning, addressing the issue that nonstandardized data formats and lack of semantic interoperability often impede the use of big data in AI. Similarly, Namli et al. [3] developed a formal data preparation pipeline with FHIR profiling to define needed information (phenotypes, features) for AI models. By leveraging FHIR resources to represent complex electronic health records, they improved traceability and reproducibility in AI model training. These works focus primarily on offline data preparation and model training, but their principles of using FHIR for consistent data representation are directly relevant to real-time patient assessment scenarios.

Specific architectures have also been proposed for integrating patient self-reported data with clinical systems. Mukhiya et al. [5] presented a cloud-based self-reporting e-health system architecture built on a Service-Oriented Architecture and HL7 FHIR to allow patients to directly interact with EHRs for mental health monitoring. Their prototype used FHIR (and the SMART on FHIR framework) to enable secure exchange of questionnaire data from a mobile app to healthcare providers. The use of FHIR ensured that patient responses were standardized and accessible to clinicians for assessment, all within a secure environment. Another notable example is the CAPABLE project's multi-agent system for remote cancer patient monitoring (Lanzola et al. [4]). In CAPABLE, patients' PROs and sensor data are shared on a common platform using FHIR to achieve semantic interoperability among various agent-based services. The system provides coaching to patients and decision support to doctors, with a central Case Manager orchestrating data flow whenever new information is added. Their architecture underscores how FHIR's standardized data models enable different components (agents, AI algorithms, EHR systems) to correctly interpret and act on exchanged data.

HL7 International's FHIR Release 4 specification [6] and follow-up guidance on API security have been key enablers for these efforts. The National Institute of Standards and Technology (NIST) has highlighted the role of FHIR APIs in enabling AI access to health data [7], while HL7's own response statement "Playing with FHIR" emphasizes that security vulnerabilities lie in implementations rather than the standard itself [8].

These related works collectively show a trend toward adopting health data standards in AI systems to ensure integration and reliability. However, many existing solutions address either data interoperability or AI analysis in isolation. This work contributes a holistic conceptual approach that spans the entire journey of patient-provided information: from the point of collection (patient interface) to the point of use (physician's decision), with AI processing in between. By building on the successes of prior FHIR-based systems and incorporating comprehensive security and workflow considerations, a blueprint is provided for AI-assisted patient assessment that healthcare institutions can trust and implement. This approach is distinguished by its end-to-end perspective and emphasis on using international standards (FHIR and related protocols) to ensure that AI integration into clinical practice is both seamless and safe.

3. Proposed methodology

3.1. Architecture overview

We propose a modular architecture that manages the flow of patient-sourced data to an AI engine and back to the clinician, underpinned by HL7 FHIR for data representation and exchange. Figure 1 outlines the key components and steps.

Alt-text: The diagram illustrates a patient using a mobile health application to input subjective data (e.g., symptoms, medical history, or responses to a questionnaire). The data is immediately structured into HL7 FHIR resources (such as a QuestionnaireResponse for survey answers or an Observation for symptom metrics) and sent via a secure FHIR API to a healthcare server or cloud platform. An AI-assisted analysis module retrieves the patient's FHIR-formatted data, applies machine learning models to assess the patient's condition (for example, predicting a risk score or suggesting possible diagnoses), and generates a result. This AI result is then encapsulated in a FHIR resource (e.g., a DiagnosticReport or ClinicalImpression) and made available to the physician through their clinical interface (which could be an EHR system or decision support dashboard). The physician reviews the AI-provided assessment alongside the original patient data, all in a unified, standardized format. The entire process ensures that data fidelity is maintained from end to end and that each actor in the workflow (patient, AI, provider) exchanges information in a common language.

3.2. Data collection and FHIR encoding

The process begins at the patient side, where subjective information is captured. This could occur through a dedicated smartphone app, a web portal, or even telehealth sessions. To ensure completeness and structure, the system may present the patient with a standardized questionnaire (for instance, a symptom checker or a PROM – Patient-Reported Outcome Measure). Each piece of information the patient provides is converted to a corresponding FHIR resource. For example, if a patient reports having chest pain and shortness of breath, the application might create FHIR Observation resources with coded terms for these symptoms (using standard vocabularies like SNOMED CT for "chest pain" and "dyspnea"). If the patient fills out a questionnaire, a FHIR QuestionnaireResponse resource is generated, linking each answer to a defined question code. By structuring data at the source, we facilitate downstream processing: the AI model does not need to guess or interpret free-text inputs, and the semantics of each data element are explicit. Multiple data types can be handled – quantitative readings (like home-measured blood pressure or glucometer values) become Observation resources, while narrative explanations can be captured as annotated notes or as part of a FHIR Communication resource. Each resource is tagged with the patient's identifier and timestamp, and optionally metadata (such as data provenance) can be included to log that "patient self-reported via app" at a certain time.

3.3. Guaranteed data communication

After structuring the patient's data to meet FHIR standards, it is communicated securely to the healthcare/artificial intelligence system via a secure internet connection. HL7 FHIR is an interoperability standard, not a security standard; therefore, we incorporate proven security mechanisms to safeguard data during transmission. The interaction between the patient application and the server is through HTTPS, with TLS encryption to ensure against unauthorized interception of health data packets across the network. All interaction with the FHIR API is subject to an authentication and authorization process. Advanced signal transmitting techniques and modeling of phased array antennas for urban data transmission, can further enhance the reliability and security of patient data communication in telehealth settings [9,10]. These technologies support robust connectivity between patient devices and healthcare servers, minimizing data loss or interception. The method applied utilizes the SMART on FHIR framework for authentication; the patient application is in this instance a SMART client, which acquires an OAuth 2.0 access token upon the patient having authorized data sharing. This flow leverages OpenID Connect to facilitate user authentication and OAuth 2.0 for the delegated authorization use case, allowing the patient to securely grant the application permission to write their data into their medical record or artificial intelligence service without revealing their credentials. On the server side, an authorization server is responsible for token validation and verifying that the data write operation coming in is authorized (e.g., that the patient is writing into their own record). Access control rules are applied to ensure that only the AI module and the respective clinical personnel can access this patient's information. Security labels and consent resources within FHIR can be utilized to convey the patient's consent and any privacy restrictions (e.g., labeling some data as being restricted). By adhering to FHIR security guidelines, the system uses proven techniques: all data exchanges are encrypted, clients are verified, access is controlled, and an audit log is kept for regulatory compliance.

3.4. AI processing module

Once the patient data is stored in the FHIR repository of the system (which is typically within an EHR or an independent data lake), the AI module is triggered. The module may either be on-premises in the IT setup of the hospital or cloud-based. The system requests pertinent FHIR resources concerning the patient via the FHIR REST API; it may request all Observation resources on Patient X for the past 24 hours or retrieve the newest QuestionnaireResponse for a given survey. As the input data is highly structured, minimal preprocessing is necessary—semantic tags render it straightforward for the model to understand which datapoints map onto specific clinical concepts. The AI model itself can be a prediction model (e.g., a sepsis risk stratification model) or diagnostic reasoning model (e.g., a machine learning classifier or ensemble). If additional data is needed (e.g., previous diagnoses or medications), those can be pulled as well using FHIR API calls. For robust and reliable AI inference, especially in longitudinal patient monitoring, the stability and convergence of neural network models are critical. Recurrent neural networks with discrete delays, which are suitable for processing temporal health data, have been analyzed for exponential stability using Lyapunov functionals and Kertesz method-based estimations [11-13]. Ensuring such stability in AI models enhances their reliability in clinical decision-making, particularly when processing time-dependent patient-reported outcomes or sensor data streams. Further research has demonstrated conditions for global asymptotic stability in lattice-based differential models with delay, which can inform the design of distributed AI systems processing spatially or hierarchically structured health data [14-16]. These mathematical frameworks support the development of trustworthy AI components in clinical pipelines by ensuring predictable model behavior under variable input conditions. Data security is provided during AI processing: the AI connects securely through secure API calls as a valid client and does not retain patient data outside the processing scope.

3.5. Result generation and return to physician

Following diligent analysis, the artificial intelligence module generates an output that must be communicated to the clinician in a comprehensible format. This may take the shape of a risk score, a priority level, or a textual recommendation. We embed the AI output within a FHIR resource—a DiagnosticReport or ClinicalImpression typically—with a summary of findings (and pointers to the input data) and metadata specifying its origin from an AI algorithm. The physician sees the AI output through their normal workflow: when they open the patient's record in the EHR, the new report appears alongside lab reports and clinician notes. This is feasible because most modern EHRs can support FHIR or be extended by SMART on FHIR applications. All access is logged to audit it, and the physician can indicate agreement or disagreement, so that feedback loops can improve future models.

3.6. Standards and protocols utilized

1. HL7 FHIR is the shared language for data exchange, using Version 4.0.1 for the definition of resources to ensure broad compatibility.
2. SMART on FHIR (OAuth 2.0 + OpenID Connect) for application and user authentication and authorization.
3. Legacy integration by using translation services, which convert HL7 v2 messages to FHIR resources, providing a standard format for data.
4. Encryption of data-at-rest in databases and logs, managed in line with hospital policies.
5. Use standard vocabularies (SNOMED CT for problems/symptoms, LOINC for laboratory tests) within FHIR resources to provide semantic correctness.
6. Provenance and AuditEvent resources in FHIR for recording creation, access, and modification events—allowing for full traceability and accountability.
7. Modular, extensible framework for adding new data types (e.g., ImagingStudy for medical imaging or waveform data) to include in the pipeline by reference (with the actual image data being processed by DICOM standards).

This pairing of protocols and standards creates a secure, standardized, and transparent pipeline —structure via FHIR models, security via encryption and OAuth-based access controls, and clarity via semantic coding—a solid foundation for AI-enabled patient care workflows without disruption.

4. Results and discussion

Implementing the above conceptual approach can yield significant benefits in clinical practice, but it also comes with practical considerations and challenges. In this section, we discuss the implications of adopting a FHIR-centered AI data transmission pipeline, examine potential outcomes in terms of care quality and efficiency, and highlight the challenges and limitations that must be addressed.

4.1. Interoperability and data quality benefits

One of the key advantages of the proposed approach is that there is improved interoperability. By having one standard (FHIR) for all data exchanges, our system enables different components and stakeholders to work together out of the box. Patient apps, hospital IT systems, and AI services—perhaps provided by different suppliers—can communicate with one another since they share the FHIR API and data models. The facilitation of interoperability enhances the integration activities associated with it while concurrently enhancing the quality of data: every individual datum is encoded using standardized codes and fields, thereby minimizing ambiguity. Instead of depending on an unstructured raw-text symptom description that may be variable or need natural language processing for interpretation, for instance, an AI system is presented with a coded and structured entry regarding symptoms. This greater clarity renders artificial intelligence output more trustworthy, since the AI is confident that the input data conforms to anticipated schemas. Standardization also promotes consistency; as Lanzola et al. point out, the application of open standards such as FHIR and established terminologies can contribute to improved quality of data gathered as well as more trustworthy recommendations. On a practical level, this implies that a clinician will more readily believe and endorse an AI-generated evaluation if it originates from well-organized data and is presented in a recognizable format. Stepwise accumulation of standardized data might also bring advantages to health organizations by forming an entire, interoperable dataset for secondary purposes like retrospective analyses, quality improvement, or additional training of AI algorithms.

4.2. Clinical workflow integration

Our approach is designed to integrate smoothly into existing clinical workflows and be minimally disruptive. By presenting artificial intelligence results as standard clinical data (e.g., a DiagnosticReport), the platform presents AI insights via the same portals clinicians use to view lab results or radiology reports. This integration solves one of the most common challenges in implementing AI—lack of physician adoption and "workflow fit." If the AI requires separate logins or produces data in a vacuum, busy clinicians will ignore it. In this instance, however, the physician can view the patient's self-reported symptoms and the AI risk assessment at the same time in the EHR. This can facilitate faster clinical decision-making, for instance, by highlighting high-risk cases. Imagine a possible result: a triage system based on our pipeline flags a patient's incoming symptom report with a high risk score for cardiac arrest. The on-call doctor is immediately alerted through the EHR with the AI's suggestion to prioritize this patient, perhaps averting an adverse event. In this scenario, we see enhanced efficiency (the AI winnows out and prioritizes cases) and potentially better patient outcomes (sooner intervention). In presenting AI suggestions in context, our design tries to keep the doctor involved and up to speed—the physician doesn't merely view the AI's result, but the information it was based on, and all in a structured form he can audit.

4.3. Patient engagement and empowerment

A second beneficial implication relates to patient empowerment throughout the process of assessment. By inputting their own data and observing its immediate application in care-related decisions, this interaction has the potential to enhance both their involvement and satisfaction levels. Patients may sense that their contribution—their subjective complaints and concerns—is listened to and officially documented. Additionally, standardized education or feedback can be communicated back to the patient. For example, once the AI has analyzed, the patient's app could show a message: "Your symptoms suggest a likelihood of X; your physician has been informed and will follow up." This completes the loop, assuring the patient and possibly enhancing compliance with suggestions. Care would clearly be taken, naturally, in what is communicated to patients to avoid causing concern from unedited AI responses—assuming the doctor would review any information prior to forwarding to patients. Nonetheless, a pipeline format allows for such controlled back-and-forth communication.

4.4. Security and privacy concerns

While the method has strong security mechanisms, there has to be continuous vigilance in implementation. One of the issues of debate is balancing data access and privacy. Our approach applies secure authentication and authorization, but when data are sent to third-party AI services (specifically cloud-based), problems arise: are there data leaks or data abuses? HL7 claims the chief vulnerabilities do not lie in FHIR but in how implementers secure their servers and APIs. Thus, healthcare providers who utilize this framework must guarantee safe deployment of the FHIR server through adequate access controls, routine security audits, and adherence to established standards. A benefit is that utilization of a standardized API simplifies the implementation of centralized security. Patient consent can also be more fine-grained; e.g., a patient can consent to release some symptom information to an AI triage service but not their full medical record, and the system can enforce that degree of granularity. Data anonymization or pseudonymization may be incorporated for model-training steps (if AI models get better over time with more data, then the use of de-identified pooled data from numerous patients should be contemplated, with governance oversight). Highlighting that FHIR's security infrastructure can be made compatible with any protocol that is needed (e.g., JWT tokens, VPNs for internal access) assures that taking on this modern pipeline doesn't involve a trade-off in terms of privacy. Yet, ongoing software updates to prevent newly arising threats—such as keeping libraries updated to neutralize cyber weaknesses and tracking anomalous API usage patterns—are an unavoidable component of maintaining an AI-based system in the healthcare sector.

4.5. Challenges and limitations:

Despite its potential, the theoretical model faces many challenges:

- Data Heterogeneity and Mapping: Although FHIR supports a wide range of resource types, patient input in real-world settings may be unstructured. It is problematic to ensure that all kinds of information provided by patients are correctly articulated under the FHIR schema. Certain information will continue to exist in semi-structured or narrative form. Secondary use of natural language processing may be required to complete some FHIR elements from unstructured text (e.g., mapping the description "pain is crushing and 5 minutes in duration" to a structured representation). There is a possibility that erroneous mapping of data can lead to an inaccurate interpretation by the AI. This involves extensive designing of questionnaires and continuous testing of the pipeline with actual users to incorporate the various ways patients explain their symptoms.
- Technical Integration and Legacy Systems: Not all hospitals are equal in terms of IT maturity. Not all systems natively support FHIR. Our pipeline might require an integration layer or middleware that mediates between legacy formats (HL7 v2 messages, databases, etc.) and FHIR. This adds complexity and possible points of failure. However, industry trends are on the side of FHIR adoption, and HL7 v2-to-FHIR conversion tools exist. These types of projects would be assisted by institutional support for interoperability initiatives. Having IT departments involved early to create FHIR servers or take advantage of existing ones (most EHR vendors now offer FHIR API endpoints) is crucial. Additionally, the AI component needs to be integrated; this is easier to achieve for proprietary models but needs to be thoughtfully considered for third-party AI services so that they are FHIR standard compatible for data exchange or a proper conversion mechanism is in place. Some learning and development investment might be required to facilitate communication between all the components in fluent FHIR language.
- Performance and Scalability: As the volume of data or number of patients grows, performance must be maintained. Real-time or near-real-time analysis can be demanding, especially if AI models are computationally intensive. The architecture should be tested for high-throughput scenarios (e.g., hundreds of symptom reports per hour during a public health crisis). Cloud infrastructure can be leveraged to auto-scale resources, and the stateless nature of FHIR RESTful services aligns well with scalable deployment (multiple servers behind a load balancer). Studies suggest cloud-based FHIR implementations can scale to large workloads, and cloud providers offer managed FHIR services that are HIPAA-compliant. Still, careful tuning is needed: indexing of FHIR data for quick search, caching of frequent queries, and possibly asynchronous processing for AI if immediate response is not critical. Ensuring that predictions are delivered in a timely manner is vital; delays in data processing could negate the benefits (e.g., if an urgent risk prediction comes too late). Thus,

part of the results of our conceptual study is the recognition that system performance testing and optimization are an integral part of implementation.

- Clinical Validation and Trust in AI: Even with perfect data transmission, the usefulness of the system depends on the AI's accuracy and the clinicians' trust in it. One challenge is ensuring the AI model is well-trained and validated on the kinds of data it will encounter (including patient self-reported data, which might differ from clinician-recorded data). If the model has biases or high false-positive/negative rates, clinicians might ignore its output, or worse, it could lead to erroneous clinical decisions. Therefore, an important aspect of deployment (though beyond the data pipeline itself) is rigorous evaluation of the AI algorithm using clinical trial or pilot study methodologies. Another strategy to enhance trust is providing explainability: the pipeline can support this by linking AI outputs to the input data. For example, if the AI flags "high risk of stroke," it might also provide a rationale like "based on patient's reported dizziness, arm weakness, and history of hypertension," which could be shown to the physician. With data standardized, extracting such rationale becomes more feasible (the AI knows which specific observations led to a conclusion). Engaging clinicians in the design is crucial so that the final output format suits their needs (some may want a simple score, others a more detailed report). Addressing these human factors is just as important as the technical robustness of the data pipeline.

4.6. Potential effects

If implemented successfully, the conceptual strategy will have a number of positive effects. Healthcare delivery would become more anticipatory and personalized – since patient data streams to AI in real time, complications would be detected prior to the next clinic visit. For instance, a patient with a chronic illness may log symptoms daily; the AI may detect a dangerous trend and signal the care team to intervene before hospitalization is necessary. The orderly data collection may also yield insights regarding population health: de-identified aggregate FHIR data may reveal symptom trends, treatment outcomes, and AI performance across many patients, guiding public health or quality improvement efforts. Another outcome is the enrichment of documentation: patient information provided is captured in an orderly fashion, thus avoiding transcription mistakes and saving clinicians time, since they do not need to re-key information already input by the patient. Furthermore, since our platform is designed upon interoperability standards, it can potentially act as a basis for further development—e.g., the incorporation of other AI modules (one imaging-based, one genomics-based, etc.) that can be readily plugged into the same data pipeline, thereby increasing the range of AI-supported evaluation in a modular way.

In conclusion, the deliberations underscore that an HL7 FHIR-based data transmission conceptual framework can serve a useful purpose in AI-driven patient evaluation by bringing the right data to the right algorithm and, conversely, to the clinician in a dependable way. This method has the potential for interoperability, security, and efficiency, which is in tandem with global best practices in health information technology. Yet, it is necessary to take careful consideration of the implementation challenges involved, particularly with regards to security strengthening, system effectiveness, and clinical acceptance. By acknowledging and planning for these hurdles, healthcare organizations are able to utilize this approach to improve patient outcomes and provider workflows in the setting of AI-enhanced medicine.

Conclusion

This article provides a holistic conceptual model of patient data transmission to artificial intelligence systems, and result dissemination to physicians in a scenario with AI-assisted patient evaluation. Method emphasizes the application of HL7 FHIR standards to guarantee that the data is structured, compatible, and securely transmitted at each step. Starting from the patient's own feedback, we outlined how data can be standardized into FHIR resources, encrypted and authenticated using OAuth 2-based mechanisms, and flow into clinical processes so that the doctor's decision-making process is enhanced by insights produced by AI. Demonstrated, through argument and reference to current literature, that a standardized pipeline can enhance the data availability and comprehensibility to AI models and thereby the reliability of AI predictions, and yet be interfaced with current healthcare information systems with minimal disruption. The utility of this approach lies in gap bridging: it closes the technical gap between patient-generated data and clinical AI algorithms, and the communication gap between AI output and clinician use. Through the utilization of internationally accepted standards and protocols, the framework aligns with global health data interoperability initiatives and can be adapted for diverse healthcare environments. It addresses privacy and security issues of fundamental significance, understanding that patient trust is paramount in digital health technologies.

Future research directions

The architecture is theoretical in nature and needs to be validated in real-world settings. In future studies, we plan to conduct pilot studies to test the overall system with actual patients, healthcare providers, and artificial intelligence models. Those pilots mentioned above are aimed at collecting data on user experience for both patients and providers, system efficacy, and clinical outcomes. For instance, a pilot can be the utilization of an artificial intelligence triage tool within a primary care practice, where patients complete an electronic symptom survey prior to a visit, and the AI has the capability to offer an initial evaluation to the physician. Some metrics such as diagnostic concordance, time efficiency, and patient satisfaction will be assessed. Furthermore, more work is justified to polish the communication of AI explanations; adding explainable AI methods to the FHIR-based findings could help enhance transparency and trust. Also mentioned the need for adjustment to changing standards: HL7 FHIR is actively maintained, and upcoming versions or implementation guides (e.g., dedicated profiles for patient-reported data) may extend the possibilities of our pipeline. Another way for future research is to take into account the possibility of interoperability with other emerging health data standards, for example, using FHIR in the openEHR approach or using HL7 CDA for specific document-type information, to ensure greater compatibility across all systems. In summary, the development of an open and standardized data transmission framework is a critical step towards the safe and effective adoption of artificial intelligence in patient management. By tackling interoperability and security with a conceptual framework, this research provides the foundation for AI-supported patient assessment tools that are not just powerful and effective but also reliable and integrated into healthcare practice. The following development will necessitate an iterative development process, stringent validation, and continuous collaboration among technologists, healthcare professionals, and patients. This strategy ultimately guarantees that technology addresses the requirements of healthcare and results in improved health outcomes.

Declaration on Generative AI

The authors have not employed any Generative AI tools.

References

[1] A scalable and transparent data pipeline for AI-enabled health data ecosystems / T. Namli et al. Frontiers in Medicine. 2024. Vol. 11. URL: https://doi.org/10.3389/fmed.2024.1393123.

[2] FHIR-DHP: A Standardized Clinical Data Harmonisation Pipeline for scalable AI application deployment (Preprint) / E. Williams et al. JMIR Medical Informatics. 2022. URL: https://doi.org/10.2196/43847.

[3] Machine learning-enabled clinical information systems using fast healthcare interoperability data standards: a scoping review (Preprint) / J. A. Balch et al. JMIR Medical Informatics. 2023. URL: https://doi.org/10.2196/48297.

[4] The Case Manager: An Agent Controlling the Activation of Knowledge Sources in a FHIR-based Distributed Reasoning Environment / G. Lanzola et al. Applied Clinical Informatics. 2023. URL: https://doi.org/10.1055/a-2113-4443.
[5] Mukhiya S.K., Rabbi F., Pun K.I., Lamo Y. An architectural design for self-reporting e-health systems [Electronic resource] / S.K. Mukhiya, F. Rabbi, K.I. Pun, Y. Lamo // Proc. of the 2019 IEEE/ACM 1st International Workshop on Software Engineering for Healthcare (SEH). — Montreal, Canada, 2019. — P. 1–8. — URL: https://ieeexplore.ieee.org/document/8823897.
[6] HL7 International. HL7 FHIR Release 4 (Version 4.0.1): Fast Healthcare Interoperability Resources [Electronic resource]. — Ann Arbor, MI: HL7 International, 2019. — URL: https://www.hl7.org/fhir/r4/.
[7] National Institute of Standards and Technology (NIST). Supporting AI in Healthcare – Health IT Standards and Testing Tools [Electronic resource]. — Gaithersburg, MD: NIST, 2025. — URL: https://www.nist.gov/itl/ssd/systems-interoperability-group/health-it-testing-infrastructure/health-it-testing-1.
[8] HL7 International. Playing with FHIR: Hacking and Securing FHIR API Implementations – HL7 Response Statement [Electronic resource] / HL7 International. — Ann Arbor, MI: HL7 International, 2021. — URL: https://blog.hl7.org/statement-to-the-global-community-from-hl7-international-on-the-paper-playing-with-fhir-hacking-and-securing-fhir-apis.
[9] Palianytsia Y., Dunets V., Khvostivska L. (2023). Modeling of Phased Array Antenna for Data Transmission in Urban Environment. Proceedings of the 3rd International Workshop on Information Technologies: Theoretical and Applied Problems. 370-381.
[10] Palaniza, Y., Martseniuk, A., Yaskiv, V. (2024). Active Phased Array Antenna With Parallel Feeder Excitation At 3.7 Ghz. Advances in Electrical and Electronic Engineering, 22(2), 127-133. https://doi.org/10.15598/aeee.v22i2.5290
[11] Nakonechnyi O., Martsenyuk V., Sverstiuk A. On Application of Kertesz Method for Exponential Estimation of Neural Network Model with Discrete Delays. (2020) Mechanisms and Machine Science, 70, 165 – 176. https://doi.org/10.1007/978-3-030-13321-4_14.
[12] Martsenyuk V., Sverstiuk A. (2019). An exponential evaluation for recurrent neural network with discrete delays. System Research and Information Technologies, Vol.2, 83 – 93. https://doi.org/10.20535/SRIT.2308-8893.2019.2.07.
[13] Martsenyuk V., Andrushchak I., Sverstiuk A., Klos-Witkowska A. (2018). On investigation of stability and bifurcation of neural network with discrete and distributed delays. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 11127 LNCS, 300 – 313. https://doi.org/10.1007/978-3-319-99954-8_26.
[14] Martsenyuk, V., Soldatkin, O., Klos-Witkowska, A., Sverstiuk, A., & Berketa, K. (2024). Operational stability study of lactate biosensors: modeling, parameter identification, and stability analysis. In Frontiers in Bioengineering and Biotechnology (Vol. 12). Frontiers Media SA. https://doi.org/10.3389/fbioe.2024.1385459.
[15] Martsenyuk, V., Klos-Witkowska, A., & Sverstiuk, A. (2020). Stability Investigation of Biosensor Model Based on Finite Lattice Difference Equations. In Springer Proceedings in Mathematics & Statistics, 297–321. Springer International Publishing. https://doi.org/10.1007/978-3-030-35502-9_13.
[16] Martsenyuk, V., Sverstiuk, A., & Gvozdetska, I. S. (2019). Using Differential Equations with Time Delay on a Hexagonal Lattice for Modeling Immunosensors. In Cybernetics and Systems Analysis. Vol. 55, Issue 4, 625–637. Springer Science and Business Media LLC. https://doi.org/10.1007/s10559-019-00171-2.