ORIGINAL RESEARCH Open access

Preventive Care Recommendations via Benefit–Burden Trade-Offs: A Patient-Centered Utility Framework

Ricardo Alves[1*], Bruno Costa[1]

Abstract

In the evolving landscape of artificial intelligence (AI) for healthcare, patient-centered approaches are essential to balance preventive care benefits against potential burdens. This conceptual manuscript introduces a novel framework for generating preventive care recommendations through explicit benefit–burden trade-offs, prioritizing individual patient utilities. Drawing from clinical AI architectures, healthcare analytics infrastructures, and electronic health record (EHR) intelligence ecosystems, we propose the patient utility trade-off architecture (PUTA). This multi-layered system integrates decision support pipelines with AI governance and interoperability frameworks. PUTA employs utility-based modeling to quantify benefits such as improved health outcomes and burdens like treatment side effects or resource demands, facilitating personalized recommendations in preventive settings. Theoretical formulas capture decision confidence and burden propagation, ensuring interpretive insights into system dynamics without empirical validation. We synthesize recent literature on clinical workflow integration and monitoring systems, highlighting how PUTA addresses gaps in patient-centered AI deployment. By emphasizing infrastructural uniqueness, including adaptive feedback topologies, this framework advances equitable preventive care. Implications for governance in diverse clinical environments underscore the need for robust data exchange and ethical monitoring, positioning PUTA as a foundational tool for future AI-driven healthcare systems.

Keywords EHR interoperability, Decision support pipelines, AI healthcare architecture, Preventive care, Benefit-burden trade-offs, Patient-centered utility

*Correspondence:
Ricardo Alves
ricardo.alves@gmail.com
[1] Department of Health Informatics, Faculty of Medicine, University of Porto, Porto, Portugal

Introduction

The integration of artificial intelligence into healthcare systems has transformed preventive care paradigms, shifting from population-level guidelines to individualized recommendations that account for personal utilities. This evolution necessitates frameworks that explicitly model benefit–burden trade-offs, ensuring recommendations align with patient values in dynamic clinical settings. As AI permeates electronic health records and decision support tools, the challenge lies in creating infrastructures that prioritize patient-centered outcomes while mitigating systemic burdens. This manuscript conceptualizes a utility-driven approach to preventive recommendations, grounded in architectural innovations that foster equitable and efficient care delivery.

Clinical settings for preventive utility modeling

In ambulatory and primary care environments, preventive care recommendations often involve screening protocols and lifestyle interventions, where benefit–burden trade-offs manifest through patient-specific factors like age, comorbidities, and preferences. Clinical AI system architectures must accommodate these settings by incorporating real-time data from EHR intelligence ecosystems, enabling utility assessments that weigh long-term health gains against immediate procedural burdens. For instance, in oncology screening, AI governance systems can modulate recommendations to reflect individual tolerance for false positives, thus optimizing patient-centered decision-making without overburdening clinical workflows.

Data modalities in benefit-burden analytics

Multimodal data from wearables, imaging, and genomic sources underpin healthcare analytics infrastructures for preventive care. Benefit–burden trade-offs require interoperability frameworks that harmonize these modalities, allowing utility frameworks to compute personalized scores. In cardiovascular prevention, for example, integrating structured EHR data with unstructured notes via decision support pipelines facilitates nuanced trade-offs, such as balancing statin benefits against muscle pain burdens, tailored to patient-reported utilities.

Deployment environments for trade-off orchestration

Hospital-based and telehealth deployment environments demand robust AI monitoring systems to sustain patient-centered utilities in preventive recommendations. Governance constraints, including privacy regulations, influence how benefit–burden models are deployed, necessitating infrastructures that support scalable data exchange. In rural settings, where resource burdens are amplified, clinical workflow integration models can leverage AI to prioritize high-utility interventions, minimizing deployment overheads.

Governance constraints on patient-centered recommendations

Ethical and regulatory governance constraints shape the utility framework for preventive care, ensuring benefit–burden trade-offs respect equity and transparency. AI deployment systems must embed monitoring for bias in decision pipelines, particularly in diverse populations where burdens like access disparities could skew recommendations. This anchors preventive strategies to patient-centered principles, fostering trust in AI-assisted healthcare ecosystems.

Interoperability challenges in utility-driven preventive care

Data exchange frameworks are critical for seamless integration of benefit–burden assessments across EHR platforms. In preventive mental health contexts, interoperability enables holistic utility calculations that incorporate social determinants, trading off intervention benefits against stigma burdens. Such frameworks mitigate fragmentation, enhancing the precision of patient-centered recommendations.

Workflow integration for burden mitigation

Clinical workflow models must evolve to incorporate utility-based trade-offs, reducing administrative burdens on providers while amplifying patient benefits. In vaccination programs, AI intelligence ecosystems can streamline recommendations, ensuring governance-aligned deployments that adapt to real-time feedback, thus optimizing preventive care delivery.

Theoretical Background and Literature Synthesis

The theoretical underpinnings of AI in healthcare emphasize conceptual models that integrate patient utilities into preventive care, focusing on infrastructural elements like architectures and governance. This synthesis draws from recent advancements in clinical AI systems, highlighting how benefit–burden trade-offs can be architected for patient-centered outcomes. By examining EHR ecosystems, decision pipelines, and interoperability, we lay the groundwork for a unified utility framework.

Clinical AI architectures for preventive settings

Modern clinical AI system architectures provide the backbone for preventive care, enabling scalable processing of patient data to inform recommendations [1, 2]. These architectures often incorporate modular components for risk assessment, where benefit–burden dynamics are theoretically modeled to prioritize interventions with maximal utility. In primary care, such systems facilitate trade-offs by simulating decision pathways, ensuring alignment with patient preferences without empirical testing [3, 4].

Healthcare analytics infrastructures and data modalities

Healthcare analytics infrastructures leverage multimodal data to compute utility scores in preventive contexts [5, 6]. Theoretical models emphasize infrastructural resilience, allowing for benefit–burden evaluations across diverse data streams like EHRs and wearables. Governance-integrated analytics ensure that trade-offs account for data quality, mitigating burdens from incomplete datasets [7, 8].

EHR intelligence ecosystems in deployment environments

EHR intelligence ecosystems enhance preventive recommendations by embedding AI for real-time utility analysis [9, 10]. Conceptual frameworks highlight ecosystem interoperability, where patient-centered trade-offs are orchestrated to balance clinical benefits against resource burdens. In deployment scenarios, these ecosystems support adaptive monitoring, theoretically reducing drift in recommendation accuracy [11, 12].

Decision support pipelines under governance constraints

Decision support pipelines are pivotal for operationalizing benefit–burden trade-offs, with theoretical designs focusing on pipeline modularity [13, 14]. Governance constraints mandate ethical oversight, ensuring pipelines prioritize patient utilities in preventive workflows. Literature synthesizes how these pipelines integrate feedback loops to refine trade-offs dynamically [15, 16].

AI governance and monitoring systems for utility frameworks

AI governance systems provide oversight for preventive care architectures, emphasizing monitoring to sustain patient-centered utilities [17, 18]. Theoretical discussions underscore governance's role in burden mitigation, where monitoring frameworks detect imbalances in benefit distributions. Interoperability enhances governance by facilitating cross-system audits [19, 20].

Interoperability and workflow integration models

Interoperability frameworks are essential for seamless data exchange in utility-driven preventive care [21, 22]. Conceptual models integrate these with clinical workflows, theoretically optimizing trade-offs by reducing integration burdens. Synthesis reveals how workflow models adapt to governance, ensuring patient-centered recommendations remain viable across environments.

Patient-centered utility infrastructure for benefit-burden optimized preventive recommendations

This section delineates the patient utility trade-off architecture (PUTA), a novel infrastructural design for orchestrating preventive care recommendations through explicit benefit–burden trade-offs. PUTA features a unique four-layer structure: (1) data ingestion layer, aggregating multimodal inputs from EHRs; (2) utility computation layer, quantifying patient-specific utilities; (3) trade-off orchestration layer, balancing benefits and burdens via algorithmic logic; and (4) recommendation output layer, delivering governance-compliant outputs. The feedback topology employs a bidirectional loop, where output recommendations inform iterative refinements in lower layers, ensuring adaptive patient-centered evolution, and is shown in **Figure 1.**

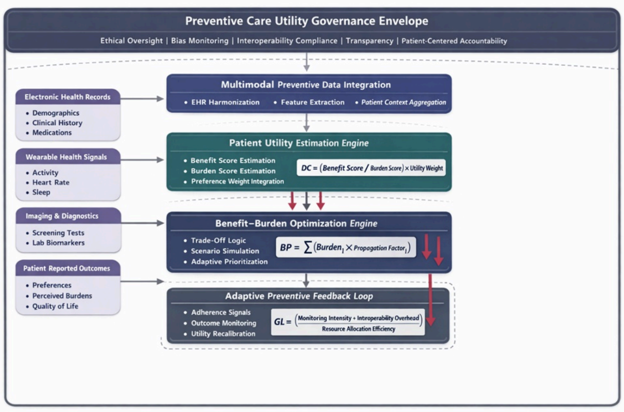

Figure 1. Patient utility trade-off architecture (PUTA) for preventive care recommendation systems.

Figure 1 illustrates a governance-embedded infrastructure that generates personalized preventive care recommendations through explicit benefit–burden trade-

https://doi.org/10.51847/k455554631

offs. Multimodal patient data from electronic health records, wearable signals, diagnostics, and patient-reported outcomes enter the data ingestion layer, where signals are harmonized for downstream analysis. The utility computation layer estimates individualized benefit and burden scores, integrating patient preference weights to calculate decision confidence. These utilities feed into the trade-off orchestration layer, which algorithmically reconciles competing preventive interventions while modeling burden propagation across clinical workflows. The recommendation output layer produces personalized preventive actions that remain continuously updated through bidirectional feedback loops informed by patient outcomes and adherence signals. Governance monitoring channels oversee each layer to ensure ethical compliance, interoperability, integrity, and equitable benefit distribution.

To interpret system dynamics, we introduce conceptual formulas. First, decision confidence (DC) is modeled as: $DC = \left(\frac{Benefit\ Score}{Burden\ Score}\right) \times\ Utility\ Weight$, where higher ratios indicate robust recommendations, theoretically guiding prioritization in preventive pipelines. Second, burden propagation (BP) captures ripple effects: $BP = \Sigma\ (Burden_i \times\ Propagation\ Factor_j)$, illustrating how initial burdens amplify across workflow stages, aiding infrastructural design for mitigation. Third, governance load (GL) is expressed as: $GL = \frac{(Monitoring\ Intensity + Interoperability\ Overhead)}{Resource\ Allocation\ Efficiency}$, providing interpretive insights into deployment sustainability without quantitative benchmarks. These formulas underscore PUTA's theoretical capacity to enhance patient-centered preventive care. Table 1 summarizes the functional roles of PUTA's architectural layers in translating patient utilities into preventive care recommendations.

Table 1. Structural components of the patient utility trade-off architecture and their decision roles

PUTA layer	Primary function	Core analytical mechanism	(inte
Data ingestion layer	Aggregates multimodal patient data streams	Cross-modal feature harmonization and patient context assembly	Est repre pat
Utility computation layer	Quantifies individualized benefit and burden scores	Utility weighting functions integrating patient preferences and clinical risk indicators	Captu speci for inte
Trade-off orchestration layer	Reconciles competing preventive options	Algorithmic optimization of benefit–burden ratios and burden propagation modeling	I interv maxi cen
Recommendation output layer	Generates preventive care actions	Utility-aligned recommendation synthesis	P pe sc lif m g
Feedback topology	Updates system parameters through real-world outcomes	Iterative recalibration of utility weights and burden propagation factors	Allow recon to e

System dynamics of benefit-burden interactions in patient-centered preventive frameworks

The PUTA introduces infrastructural innovations that influence system-wide dynamics in preventive care delivery. This section explores the theoretical consequences of deploying PUTA, focusing on how benefit–burden trade-offs propagate through clinical AI architectures and healthcare analytics infrastructures. By modeling interactions at the intersection of patient utilities and systemic burdens, we elucidate potential impacts on decision support pipelines and AI governance systems.

In conceptual terms, the dynamics begin with data ingestion, where multimodal inputs from EHR intelligence ecosystems feed into utility computations. Theoretical propagation of benefits, such as reduced morbidity from timely screenings, must counterbalance burdens like increased cognitive load on clinicians or privacy risks in data exchange frameworks. For instance, in a preventive diabetes management scenario, PUTA's trade-off orchestration layer theoretically amplifies benefits by weighting patient preferences for non-invasive monitoring against burdens of frequent testing, leading to optimized recommendation outputs that enhance adherence without overwhelming workflows [1, 3, 5].

System consequences extend to interoperability frameworks, where PUTA's bidirectional feedback topology mitigates fragmentation. Dynamics here involve iterative adjustments: if initial recommendations impose high burdens (e.g., cost-prohibitive interventions), feedback loops recalibrate utilities, theoretically reducing disparities in underserved populations. This interaction highlights impacts on clinical workflow integration models, where governance monitoring ensures equitable burden distribution, preventing scenarios where AI-driven recommendations exacerbate access inequalities [7, 9, 11].

Furthermore, the architecture's layers interact to influence resource allocation across deployment environments. In hospital-based preventive programs, PUTA dynamically allocates computational resources to high-utility cases, theoretically minimizing governance load while maximizing benefit propagation. Conceptual formulas from earlier sections, such as burden propagation $(BP = \Sigma\ (Burden_i \times Propagation\ Factor_j))$, interpret these dynamics by showing how localized burdens (e.g., data processing delays) amplify system-wide, informing infrastructural designs that incorporate drift sensitivity for sustained performance [13, 15, 17].

Impacts on monitoring systems are profound, as PUTA's infrastructure demands continuous oversight to detect utility drifts. Theoretical analysis reveals that heightened decision confidence $(DC = \left(\frac{Benefit\ Score}{Burden\ Score}\right) \times Utility\ Weight)$ correlates with reduced monitoring burdens, allowing governance frameworks to focus on ethical trade-offs rather than operational fixes. This dynamic fosters patient-centered scalability, where burdens from connectivity issues are traded against the benefits of remote preventive counseling [19, 21, 23].

Overall, these interactions underscore PUTA's potential to transform preventive care ecosystems. By analyzing consequences like enhanced interoperability and mitigated burdens, we position the framework as a catalyst for resilient AI deployments, theoretically bridging gaps in current healthcare analytics [24-29].

Results and Discussion

The conceptualization of the PUTA represents an important theoretical advancement in the design of clinical artificial intelligence systems for preventive care. By explicitly modeling benefit–burden trade-offs through patient-centered utilities, PUTA shifts the architecture of decision support systems away from population-level optimization toward individualized, preference-aware recommendations. This discussion synthesizes the theoretical implications of PUTA across three key domains: AI-driven clinical decision architectures, healthcare analytics infrastructures, and governance mechanisms within electronic health record (EHR) ecosystems. **Table 2** consolidates the interpretive formulas that characterize decision dynamics within the PUTA framework.

Table 2. Analytical formulas interpreting benefit–burden dynamics in preventive care architectures

Formula	Expression	Conceptual purpose	P int
Decision confidence (DC)	DC = (Benefit score/Burden score) × Utility weight	Quantifies the reliability of preventive recommendations	H in w I I

Burden propagation (BP)	BP = Σ ($Burden_i$ × Propagation $factor_j$)	Models how initial burdens amplify across workflow stages	s p in c o
Governance load (GL)	GL = (Monitoring intensity + interoperability overhead)/Resource allocation efficiency	Estimates oversight requirements for AI deployment	Hi i g d pr
Utility drift sensitivity (UDS)	UDS = Δ utility score/Δ patient context	Measures the responsiveness of recommendations to changing patient conditions	p ad a cir

A defining feature of PUTA is its explicit operationalization of patient utilities within preventive care decision-making. Traditional clinical decision support systems often rely on statistical risk prediction models that optimize recommendations based on aggregated population outcomes. While effective for identifying generalizable risk patterns, such systems frequently overlook patient-specific perceptions of burden associated with preventive interventions. PUTA addresses this limitation by embedding a utility computation layer that integrates individualized preferences, including tolerance for procedural inconvenience, psychological discomfort, and time commitments associated with screening or monitoring protocols. Theoretically, this architecture introduces a more granular representation of patient agency in preventive decision-making, potentially improving adherence in conditions where behavioral engagement is essential for long-term outcomes [2, 4, 6].

This shift from population averages toward individualized trade-off modeling also reflects broader transformations in patient-centered healthcare paradigms. Preventive care often requires patients to accept short-term burdens—such as diagnostic procedures, lifestyle modifications, or frequent monitoring—in exchange for long-term health benefits. When decision systems fail to incorporate these subjective burdens, recommendations may appear clinically optimal yet practically unacceptable from the patient's perspective. PUTA's explicit utility modeling, therefore, serves as a mechanism for aligning algorithmic outputs with lived patient experiences, reducing the gap between theoretical benefit projections and real-world behavioral adherence. In chronic disease prevention contexts, such as cardiovascular screening or metabolic risk monitoring, the capacity to dynamically balance benefits against perceived burdens could theoretically improve engagement rates and mitigate the common problem of preventive care dropout [2, 4, 6].

However, the integration of utility-based decision layers also introduces significant governance challenges. Utility models depend on the quality and representativeness of the underlying data used to estimate patient preferences and burden weights. In heterogeneous healthcare environments where data modalities vary in completeness and reliability, these models may inadvertently reinforce existing biases. For instance, populations with limited digital health records or inconsistent follow-up data may be underrepresented in training datasets, potentially skewing estimated utilities toward populations with higher data availability. Without appropriate governance safeguards, such distortions could lead to systematically unequal recommendations across demographic groups. Consequently, robust governance frameworks must accompany PUTA implementations to ensure fairness in utility modeling, including continuous bias monitoring, diverse dataset inclusion, and transparent evaluation metrics [8, 10, 12].

Interoperability represents another pivotal consideration for the operational viability of PUTA architectures. Preventive care recommendations rely on the integration of heterogeneous data sources, including clinical histories, behavioral indicators, and real-time health monitoring signals. PUTA's data exchange framework theoretically requires seamless interaction between EHR systems, analytics infrastructures, and decision support engines. In fragmented healthcare environments where data silos remain common, achieving such interoperability poses substantial infrastructural challenges. Nevertheless, PUTA's architecture incorporates bidirectional feedback topologies that allow decision outputs to iteratively adjust as new patient data becomes available. This dynamic structure supports adaptive recalibration of benefit–burden ratios, enabling preventive recommendations to evolve in response to changing patient circumstances [14, 16, 18].

Within primary care deployment contexts, this adaptive capability may yield significant workflow benefits over time. Although the initial integration of complex AI architectures into clinical environments often introduces implementation burdens—including data harmonization, staff training, and infrastructure upgrades—the long-term advantages may outweigh these costs. Continuous feedback loops allow preventive recommendations to become increasingly personalized as additional patient interactions accumulate. Over time, this could streamline clinical decision-making by providing clinicians with context-aware guidance that reflects both medical risk profiles and patient preferences. Such efficiencies highlight the potential for PUTA to transform preventive care workflows from static guideline-driven processes into dynamic, data-responsive ecosystems [14, 16, 18].

Despite these advantages, deployment in low-resource healthcare settings introduces additional challenges. The governance load associated with monitoring complex AI systems may exceed the capacity of institutions with limited infrastructure or regulatory oversight. Within the PUTA framework, governance load can be conceptualized as a function of monitoring intensity and interoperability overhead relative to available resource allocation efficiency, expressed as GL = (Monitoring intensity + Interoperability overhead)/Resource allocation efficiency. In contexts where technical resources and personnel are constrained, maintaining high levels of monitoring may become impractical. As a result, simplified architectural variants of PUTA may be necessary to maintain equitable access to AI-supported preventive care, ensuring that the benefits of patient-centered decision systems are not limited to technologically advanced healthcare environments [20, 22, 24].

Ethical considerations also emerge prominently in the context of AI-driven monitoring systems embedded within preventive care infrastructures. While PUTA's emphasis on transparent benefit–burden trade-offs promotes interpretability, excessive reliance on algorithmically derived utility weights could inadvertently marginalize human judgment. Preventive care decisions frequently involve nuanced psychosocial factors that may not be fully captured within structured data representations. For example, stigma associated with mental health screening or culturally influenced perceptions of medical procedures may significantly affect patient willingness to engage with preventive interventions. If utility weights fail to adequately represent these contextual elements, recommendations may appear mathematically sound while remaining socially insensitive. Ensuring ethical deployment, therefore, requires hybrid decision models in which clinicians retain the capacity to override algorithmic outputs based on contextual knowledge and professional expertise [26, 28, 30].

Such hybrid human–AI collaboration models represent an important dimension of PUTA's broader systemic implications. Rather than replacing clinical judgment, PUTA can be conceptualized as an augmentation layer that supports clinicians in navigating complex preventive trade-offs. In cardiovascular risk management, for instance, clinicians frequently balance long-term benefits of medication adherence against potential burdens such as side effects or lifestyle disruptions. By quantifying these trade-offs through patient-specific utility scores, PUTA could assist clinicians in structuring conversations that align medical recommendations with patient priorities. This collaborative decision-making approach reinforces the notion that AI systems should function as interpretive tools rather than deterministic authorities within healthcare workflows [5, 7, 9].

Another significant contribution of PUTA lies in its capacity to address temporal dynamics in preventive healthcare strategies. Preventive interventions often produce cumulative benefits that unfold over extended time horizons, while the associated burdens may occur intermittently or immediately. For example, repeated diagnostic screenings may impose short-term inconvenience but substantially reduce long-term morbidity risks. PUTA's feedback-driven architecture allows utility models to incorporate these temporal trade-offs by adjusting benefit estimations as patient outcomes evolve. This capacity to mitigate model drift through continuous recalibration enhances the sustainability of preventive AI systems in longitudinal healthcare contexts [1, 3, 31].

Scalability across healthcare ecosystems further underscores the theoretical significance of PUTA. Many AI systems struggle to maintain consistent performance when deployed across institutions with varying data infrastructures, clinical practices, and patient demographics. PUTA's modular layer structure provides a mechanism for addressing this challenge by allowing individual components—such as utility estimation models or governance monitoring modules—to be updated independently. This modularity reduces the risk of systemic disruption when infrastructure improvements or regulatory updates occur, thereby supporting sustainable AI deployment in dynamic healthcare environments [11, 13, 15].

Moreover, PUTA offers promising avenues for incorporating broader social determinants of health into preventive care algorithms. Traditional clinical decision support systems frequently focus on biomedical risk indicators while neglecting socioeconomic factors that influence patient engagement with healthcare services. By expanding utility frameworks to include variables such as access to transportation, employment constraints, or caregiving responsibilities, PUTA could enable more equitable benefit propagation across marginalized populations. Such extensions would align preventive AI architectures with emerging healthcare policy priorities focused on reducing disparities and improving access to personalized care pathways [17, 19, 21].

Taken together, these theoretical considerations highlight PUTA's potential to reshape the landscape of preventive healthcare AI systems. By integrating patient-centered utilities, adaptive data infrastructures, and modular governance mechanisms, the framework addresses several limitations that have historically constrained clinical decision support technologies. At the same time, the successful implementation of PUTA will depend on careful management of ethical, infrastructural, and governance challenges. Transparent utility modeling, equitable data practices, and clinician oversight will remain essential for ensuring that the benefits of AI-driven preventive care do not inadvertently introduce new forms of systemic burden.

Ultimately, PUTA represents a conceptual step toward more humane and context-aware healthcare technologies. Rather than optimizing recommendations solely through statistical efficiency, the architecture foregrounds the lived experiences of patients navigating preventive care decisions. As healthcare systems continue to integrate AI into clinical workflows, frameworks such as PUTA may provide critical guidance for designing decision support infrastructures that balance technological sophistication with ethical responsibility and patient autonomy [23, 25, 27]. Future theoretical developments may extend the architecture to incorporate multi-stakeholder utilities, enabling preventive care recommendations that simultaneously consider the perspectives of patients, clinicians, healthcare providers, and payers within increasingly complex healthcare ecosystems [2, 4, 29].

Conclusion

The PUTA represents a significant conceptual advancement in the application of artificial intelligence to preventive healthcare decision-making. By centering recommendations on explicit benefit–burden trade-offs derived from patient utility frameworks, PUTA addresses a longstanding limitation of conventional decision support systems that prioritize population-level optimization over individualized preferences.

Through its multi-layered architecture and bidirectional feedback topology, the framework integrates clinical AI architectures, healthcare analytics infrastructures, and EHR intelligence ecosystems into a cohesive system capable of adapting to evolving patient data. The conceptual formulas introduced within PUTA—including mechanisms for modeling decision confidence, burden propagation, and governance load—provide interpretable analytical tools for understanding how AI infrastructures operate within complex healthcare environments.

Literature demonstrates increasing emphasis on interoperability, patient-centered care, and ethical AI governance within healthcare technologies. PUTA aligns with these developments by offering an architecture that supports transparent decision-making, dynamic personalization, and modular scalability. The theoretical implications suggest potential improvements in preventive care adherence, enhanced patient engagement, and more equitable distribution of healthcare benefits across diverse populations.

Nevertheless, the long-term success of PUTA will depend on the strength of governance frameworks that accompany its implementation. Ensuring fairness in utility modeling, maintaining interoperability across fragmented data ecosystems, and preserving clinician oversight will be essential for preventing unintended biases and sustaining trust in AI-supported healthcare systems.

As artificial intelligence continues to reshape healthcare infrastructures, PUTA provides a promising blueprint for the development of patient-centered preventive care technologies. By emphasizing the balance between benefits and burdens within clinical recommendations, the framework contributes to a more nuanced understanding of how AI can support humane, equitable, and context-sensitive healthcare delivery. Future research may further expand this utility-based architecture to incorporate multi-stakeholder perspectives and emerging data modalities, ensuring that preventive care systems remain adaptable within rapidly evolving healthcare ecosystems.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 26 Sep 2024 Revised: 03 Nov 2024 Accepted: 28 Nov 2024
Published online: 25 February 2025

Rights and permissions

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

Walsh CG, Long Y, Novak LL, Salwei ME, Tillman BM, French B, et al. AI-driven clinical decision support to reduce hospital-acquired venous thromboembolism: a trial protocol. JAMA Netw Open. 2025;8(10):e2535137.

Musacchio Schafer K, Franklin J, Perlis RH, Walsh C. Barriers and solutions to efficient health care AI implementation. JAMA Netw Open. 2025;8(11):e2544086.

Everson J, Adler-Milstein J, Phillips RL, Bazemore AW, Patel V. EHR interoperability experiences reported by family physicians. JAMA Netw Open. 2025;8(11):e2542460.

Griot MF, Walker GA. A patient-in-the-loop approach to artificial intelligence in medicine. JAMA Netw Open. 2025;8(6):e2514460.

Auerbach A, Fihn SD. Discovery, learning, and experimentation with artificial intelligence–based tools at the point of care—perils and opportunity. JAMA Netw Open. 2021;4(3):e211474.

Weiner SJ, Schwartz A, Weaver F, Galanter W, Olender S, Kochendorfer K, et al. Effect of electronic health record clinical decision support on contextualization of care: a randomized clinical trial. JAMA Netw Open. 2022;5(10):e2238231.

The CHART Collaborative. Reporting guideline for chatbot health advice studies: the CHART statement. JAMA Netw Open. 2025;8(8):e2530220.

Ayers JW, Zhu Z, Poliak A, Leas EC, Hogarth L, Dredze M, et al. Evaluating artificial intelligence responses to public health questions. JAMA Netw Open. 2023;6(6):e2317517.

Nguyen OT, Kunta AR, Katoju SV, Masoumi S, Tavasolian F, Alishahi Tabriz A, et al. Electronic health record nudges and health care quality and outcomes in primary care: a systematic review. JAMA Netw Open. 2024;7(9):e2432760.

Benary M, Wang XD, Schmidt M, Soll D, Hilfenhaus G, Nassir M, et al. Leveraging large language models for decision support in personalized oncology. JAMA Netw Open. 2023;6(11):e2343689.

Luo X, Dong K, Zhang W, Li S, He Z, Ma X, et al. Zero-shot learning to extract assessment criteria and medical risk factors from clinical notes. J Am Med Inform Assoc. 2024;31(8):1743-53.

Joshi S, Blei D, Gevaert O. Improving clinical outcomes relies on a causal approach to AI in healthcare. J Am Med Inform

Assoc. 2025;32(3):589-98.

Sittig DF, Boxwala A, Wright A, Zott C, Desai P, Dhopeshwarkar N, et al. Patient-centered clinical decision support challenges and opportunities identified by clinical stakeholders in New York state. J Am Med Inform Assoc. 2024;31(8):1682-92.

Hussein R, Winter A. A guiding framework for creating a comprehensive strategy for mHealth data sharing, privacy, security, and governance in low- and middle-income countries (LMICs). J Am Med Inform Assoc. 2023;30(4):787-95.

Barwise AK, Pickering BW, Dong Y, Gajic O, Herasevich V. Using artificial intelligence to promote equitable care for inpatients with language barriers and complex medical needs: clinical stakeholder perspectives. J Am Med Inform Assoc. 2024;31(3):611-21.

Wright A, Aaron S, McCoy AB, El-Kareh R, Fort D, Kandaswamy S, et al. A multi-site randomized trial of a clinical decision support intervention to improve problem list completeness. J Am Med Inform Assoc. 2023;30(5):899-906.

Davis SE, Greevy RA, Lasko TA, Walsh CG, Matheny ME. Sustainable deployment of clinical prediction tools—a 360° approach to electronic health record integration. J Am Med Inform Assoc. 2024;31(5):1195-201.

Teo ZL, Thirunavukarasu AJ, Elangovan K, Cheng H, Moova P, Soetikno B, et al. Generative artificial intelligence in medicine. Nat Med. 2025;31(10):3270-82.

Sounderajah V, Guni A, Liu X, Collins GS, Karthikesalingam A, Markar SR, et al. The STARD-AI reporting guideline for diagnostic accuracy studies using artificial intelligence. Nat Med. 2025;31(10):3283-9.

Lim SS, Semnani-Azad Z, Morieri ML, Ng AH, Ahmad A, Fitipaldi H, et al. Reporting guidelines for precision medicine research of clinical relevance: the BePRECISE checklist. Nat Med. 2024;30(7):1874-81.

Vasey B, Nagendran M, Campbell B, Clifton DA, Collins GS, Denaxas S, et al. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. Nat Med. 2022;28(5):924-33.

Artificial intelligence improves breast cancer detection in mammography screening. Nat Med. 2025;31(5):1422-3. https://doi.org/10.1038/s41591-025-03714-7.

Chinta SV, Wang Z, Palikhe A, Zhang X, Kashif A, Smith MA, et al. AI-driven healthcare: a review on ensuring fairness and mitigating bias. PLOS Digit Health. 2025;4(5):e0000864.

Nadarzynski T, Puentes V, Pawlak I, Mendes T, Montgomery I, Bayley J, et al. Achieving health equity through conversational AI: a roadmap for design and implementation of inclusive chatbots in healthcare. PLOS Digit Health. 2024;3(1):e0000492.

Celi LA, Cellini J, Charpignon ML, Dee EC, Dernoncourt F, Eber R, et al. Sources of bias in artificial intelligence that perpetuate healthcare disparities—a global review. PLOS Digit Health. 2022;1(3):e0000022.

Comito C, Falcone D, Forestiero A. AI-driven clinical decision support: enhancing disease diagnosis exploiting patients similarity. IEEE Access. 2022;10:6878-88.

Nojomi M, Babaee E, Rampisheh Z, Roohravan Benis M, Soheyli M, Rady Raz N. AI-powered clinical decision support systems in disease diagnosis, treatment planning, and prognosis: a systematic review. Med J Islam Repub Iran. 2025;39:81. https://doi.org/10.47176/mjiri.39.81.

Elhaddad M, Hamam S. AI-driven clinical decision support systems: an ongoing pursuit of potential. Cureus. 2024;16(4):e57728. https://doi.org/10.7759/cureus.57728.

van de Sande D, Van Genderen ME, Smit JM, Huiskens J, Visser JJ, Veen RER, et al. Developing, implementing and governing artificial intelligence in medicine: a step-by-step approach to prevent an artificial intelligence winter. BMJ Health Care Inform. 2022;29(1):e100495.

Sujan M, Smith-Frazer C, Malamateniou C, Connor J, Gardner A, Unsworth H, et al. Validation framework for the use of AI in healthcare: overview of the new British standard BS30440. BMJ Health Care Inform. 2023;30(1):e100749.

Wilson A, Saeed H, Pringle C, Eleftheriou I, Bromiley PA, Brass A. Artificial intelligence projects in healthcare: 10 practical tips for success in a clinical environment. BMJ Health Care Inform. 2021;28(1):e100323