International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 10 | Oct - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Human-Centred Design Thinking Approach for an AI-Based Healthcare Application

Prof.s.AbikayilAarthi
Assistant professor,Department of computer science and Engineering,
Kings College of Engineering,punalkulam,TN,India,
aarthi.cse@kingsengg.edu.in

L.Saaikirthiga
Dept of computer science and engineering,
Kings College of Engineering,Punalkulam,TN,
saaikirthiga1304@gmail.com

R.Vigneshwari
Dept of computer science and engineering,
Kings College of Engineering,Punalkulam,TN,
vigneshwariravichandran2005@gmail.com

A.Sandhiya
Dept of computer science and engineering,
Kings College of Engineering,Punalkulam,TN,
sandhiyajayabharathi@gmail.com

S.Yazhini
Dept of computer science and engineering,
Kings College of Engineering,Punalkulam,TN,
yaazz.ky@gmail.com

R.Swathi
Dept of computer science and engineering,
Kings College of Engineering,Punalkulam,TN,
swathiravichandran2504@gmail.com

J.Sandhiya
Dept of computer science and engineering,
Kings College of Engineering,Punalkulam,TN,
jsandhiya0209@gmail.com

Applied Design Thinking – Publication Assignment

***Abstract*—**Artificial intelligence (AI) is increasingly used in healthcare for risk prediction, triage, imaging and remote monitoring, yet many promising tools perform well in the laboratory but fail to become part of routine care. A common reason is that they are built around available data and algorithms rather than around the real needs, workflows and constraints of patients and clinicians. This paper proposes the use of human-centred design thinking (Empathize, Define, Ideate, Prototype and Test), consistent with the principles of ISO 9241-210, as a framework for developing an AI-based healthcare application. It describes how each stage can be applied to a clinical setting, with example research methods, personas and "How might we" questions; presents a modular system architecture with candidate modelling techniques, explainability and safety features; outlines an illustrative scenario, an implementation roadmap and an evaluation plan; and discusses ethical concerns such as patient safety, privacy, bias, transparency and accountability. The paper is conceptual: it offers a design approach and does not report experimental results. It concludes that combining design thinking with AI helps create healthcare applications that are safe, useful, trusted and adopted by the people who use them.

Keywords*—*human-centred design, design thinking, artificial intelligence, healthcare, clinical decision support, explainable AI, usability, patient safety, digital health

1. INTRODUCTION

AI has attracted great interest in medicine. Deep learning systems can detect diabetic retinopathy from retinal photographs with high accuracy (Gulshan et al., 2016), and reviews describe a growing range of applications in imaging, prediction and clinical decision support (Jiang et al., 2017; Yu, Beam and Kohane, 2018; Topol, 2019). However, a high score on a test dataset is not the same as benefit to patients. Kelly et al. (2019) note that many challenges lie between a promising algorithm and real clinical impact, including workflow integration, evaluation and trust.

Real-world studies show why people and context matter. Beede et al. (2020) observed a deep learning eye-screening system in clinics in Thailand and found that variable image quality, slow network connections and the way nurses worked all affected how useful the system was in practice. Sendak et al. (2020) describe the integration of a sepsis prediction tool into routine care and stress the importance of working closely with clinicians and fitting the tool into their workflow. Similar lessons come from the study of human–AI interaction, where AI systems are described as being especially difficult to design because their behaviour is uncertain and their outputs are complex (Yang et al., 2020).

Healthcare adds further demands. Errors can harm patients, users are diverse (patients, caregivers, doctors, nurses and community health workers), clinicians work under time pressure, and data are highly sensitive. Design thinking addresses these demands by placing people at the centre of development (Brown, 2008). Instead of asking "What can our data and algorithm do?", a design-thinking team first asks "What do patients and clinicians actually need, and what would make a tool safe and worth using?"

The objectives of this paper are: (i) to explain how human-centred design thinking can guide the development of an AI-based healthcare application; (ii) to propose a system architecture that comes out of this process; (iii) to suggest an implementation roadmap and a plan for evaluation; and (iv) to discuss ethical and practical challenges.

2. BACKGROUND

2.1 AI in Healthcare

Machine learning is applied in medicine to medical imaging, risk prediction from health records, triage, remote monitoring through wearable devices, and the analysis of clinical text (Rajkomar, Dean and Kohane, 2019; Esteva et al., 2019). Clinical decision support systems have a longer history, and their benefits depend on careful design and implementation; poorly designed support can cause problems such as alert fatigue (Shortliffe and Sepúlveda, 2018; Sutton et al., 2020). Kelly et al. (2019) also point out that many AI models are evaluated only on retrospective data, so evidence from real clinical use remains limited.

2.2 Human-Centred Design Thinking

Design thinking is an iterative, human-centred approach to solving complex problems that was developed in design practice and taught at institutions such as Stanford's d.school (Plattner, Meinel and Leifer, 2011). Liedtka (2018) explains that it works by helping teams understand users deeply, consider many options and test ideas early. It alternates between *divergent* thinking (exploring possibilities) and *convergent* thinking (selecting and refining) and is not strictly linear. In health care it has been used to improve services and products, and frameworks for its use have been described by Roberts et al. (2016) and Altman, Huang and Breland (2018). The international standard ISO 9241-210 (2019) expresses similar ideas as principles of human-centred design: understanding users, tasks and environments; involving users throughout; driving design by user-centred evaluation; working iteratively; considering the whole user experience; and using multidisciplinary skills.

2.3 Human–AI Interaction and Co-design in Clinical Settings

Amershi et al. (2019) propose guidelines for human–AI interaction, such as making clear what the system can do and how well it can do it. Cai et al. (2019) found that medical practitioners need to understand an AI tool's strengths and limitations before they can use it in decisions, and Tonekaboni et al. (2019) studied what clinicians want from explanations of machine learning outputs. These studies support the idea that clinicians and patients should take part in design, not merely receive the finished product.

2.4 Foundations for Acceptance and Engagement

Three ideas guide the design proposed here. First, acceptance of health technology depends strongly on perceived usefulness and ease of use (Holden and Karsh, 2010), so the application must fit the user's task. Second, interfaces should avoid adding unnecessary mental effort for busy users, in line with cognitive load theory (Sweller, 1988). Third, patients are more motivated to manage their health when they feel autonomy, competence and relatedness (Deci and Ryan, 2000), which is why patient control and supportive, non-judgemental feedback appear throughout the design.

3. APPLYING DESIGN THINKING TO AN AI-BASED HEALTHCARE APPLICATION

To make the discussion concrete, this paper uses a running example: an AI-based application for chronic-care support, such as type 2 diabetes, that helps patients record readings and receive personalised guidance, and helps clinicians prioritise follow-up. The same approach can be adapted to other conditions. Table 1 summarises how each stage can be used and the output expected from it.

Table 1. Design thinking stages applied to the AI-based healthcare application

Stage	Activities for the AI Healthcare App	Output
1. Empathize	Interview and observe patients, caregivers, doctors, nurses and community health workers; shadow clinic workflows; study why current health apps are abandoned; map the care journey and pain points.	Empathy maps, journey maps, list of needs
2. Define	Combine findings into clear problem statements, e.g. "How might we help a patient with diabetes understand daily readings without feeling anxious or judged?"	Personas, problem statement, success and safety criteria
3. Ideate	Brainstorm features: personalised guidance, reminders, risk alerts, clinician summaries, local-language voice support. Decide which tasks AI should do and which must stay with clinicians.	Prioritised feature list, AI–human role split
4. Prototype	Paper sketches and clickable wireframes; a "Wizard of Oz" prototype where a person simulates the AI; later a small rule-based or simple model tested on retrospective or synthetic data.	Low- and high-fidelity prototypes
5. Test	Usability sessions with patients and clinicians; clinical scenario walk-throughs; compare alerts with clinician judgement; check clarity of explanations; safety and fairness checks; iterate.	Findings, design changes, next iteration plan

3.1 Empathize

The team first learns about the people and the setting. Patients may value simple instructions, reassurance, privacy and use of their own language, while clinicians may need quick summaries, few false alarms and the ability to override the system. Needs such as anxiety, fear of being judged, trust and time pressure are discovered at this stage and are often missed in a purely technical design. Useful methods include:

- **Semi-structured interviews** with patients of different ages, literacy levels and disease severity, and with caregivers, doctors, nurses and community health workers.
- **Contextual observation and shadowing** of consultations and routine monitoring to see how information is actually recorded and used.
- **Surveys** on device access, language preference, digital literacy and internet connectivity.
- **Review of existing tools** to see where current health apps are abandoned or ignored.
- **Empathy maps and journey maps** that record what users say, think, do and feel at each step of care.

Research involving patients and clinicians requires informed consent and approval from an institutional ethics committee, and personal data collected during research must be protected.

3.2 Define

Insights are turned into personas and "How might we" (HMW) questions that guide later work (Dorst, 2011). For example, a persona may describe a middle-aged patient with diabetes who has limited time, prefers a local language and is worried about test results. Sample HMW questions include:

- How might we give patients useful daily guidance without causing anxiety?
- How might we show clinicians who needs attention without adding to alert overload?
- How might we explain a risk alert so that clinicians can trust it but still make their own decision?

Success and safety criteria are also defined here, for example improved clinical indicators, safe and timely escalation of risk, regular voluntary use, and clinician-reported usefulness. Stating them early keeps later evaluation honest.

3.3 Ideate

During ideation, the team decides what AI should do, such as estimating risk, personalising advice and summarising data, and what must remain human, such as diagnosis, treatment decisions and emotional support. Ideas are generated through brainstorming and sketching and then prioritised, for instance on an impact–effort grid and a clinical-risk check. Table 2 shows an example split of responsibilities.

Table 2. Example split of tasks between AI and humans

Task	AI role	Human role
Data entry	Assists with reminders, voice or photo input	Patient confirms and corrects entries
Risk estimation	Estimates risk of poor control or complication	Clinician reviews unusual or high-risk cases
Patient guidance	Gives personalised, plain-language tips	Clinician approves the content and limits
Alerts and triage	Flags patterns early, with reasons	Nurse or doctor decides on action
Diagnosis and treatment	Provides supporting information only	Clinician makes and owns the decision

3.4 Prototype

Early prototypes can be built without any AI. Prototyping proceeds in three steps. (1) Paper sketches and clickable wireframes of patient and clinician screens test layout and flow. (2) In a "Wizard of Oz" test (Kelley, 1984), a person plays the role of the AI behind the screen, which lets the team learn whether the experience is valuable before investing in model development. (3) A small working tool, starting with clinical guideline rules or a simple model, is built using retrospective, public or synthetic data and compared with the human wizard's choices and with clinician judgement. This order reduces the risk of building complex models for features nobody needs, and no prototype is used to make real clinical decisions until it has been reviewed for safety.

3.5 Test and Iterate

Testing checks usability, clinical safety and the quality of the AI output. Typical methods are think-aloud sessions, simulated clinical scenarios, interviews, comparison of alert strategies, and agreement checks between the system and clinician judgement. Before live use, a system can be run silently alongside routine care ("shadow mode") so that its outputs are examined without affecting patients. Early-stage clinical evaluations of AI decision support can be reported following the DECIDE-AI guideline (Vasey et al., 2022). Findings may send the team back to redefine the problem, which reflects the iterative nature of the process.

4. PROPOSED SYSTEM ARCHITECTURE

The design process leads to a modular architecture, shown in Figure 1. With consent, patient and clinical data are collected and used to update a patient model. A decision-support engine uses the model to generate personalised guidance and risk alerts. A clinician dashboard presents alerts with reasons and allows the clinician to override recommendations. A safety and governance layer records actions and monitors performance.

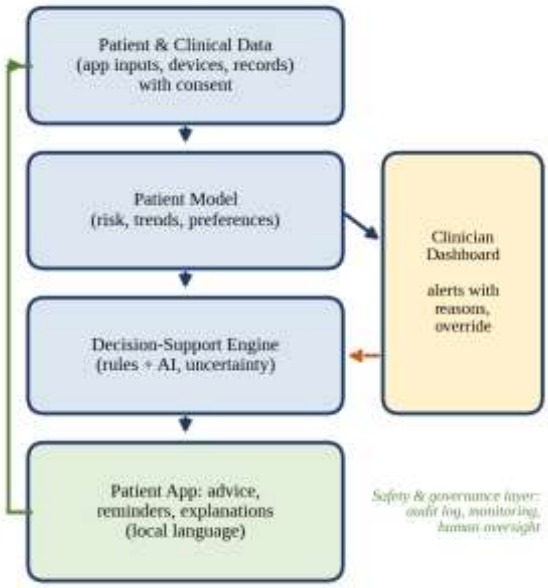


Figure 1. Conceptual architecture of the proposed AI-based healthcare application

Table 3. Main modules and their functions

Module	Function
Data collection	Records patient-entered data, device readings and relevant clinical data, using only what is needed and with informed consent.
Patient model	Estimates current status, trends, risk and preferences for each patient; handles missing or poor-quality data.
Decision-support engine	Combines guideline rules and AI models to generate guidance and alerts, with explanations and uncertainty.
Clinician dashboard	Shows prioritised patients, reasons for alerts and data quality, and allows override and feedback.
Patient app	Simple, accessible screens in local languages, with goals, progress and control over information shared.
Safety and governance	Audit logs, access control, performance and drift monitoring, incident reporting and human oversight.

4.1 Modelling Techniques

Techniques of increasing complexity can be considered. The choice should follow the needs found during empathy work and the clinical risk, not the other way round.

- **Guideline-based rules:** thresholds and pathways taken from clinical guidelines. They are simple, transparent and a good starting point.
- **Statistical risk models:** logistic regression and similar models that give calibrated risk estimates and are easier to explain.
- **Machine learning ensembles:** tree-based methods for structured health data, often accurate on tabular records.

- **Deep learning:** useful for images, signals and text (Esteva et al., 2019) but harder to explain and to validate across settings.

For high-stakes decisions, Rudin (2019) argues for using inherently interpretable models where they perform adequately rather than explaining black-box models afterwards.

4.2 Decision-Support Engine and Human Oversight

A practical design combines transparent clinical rules with a statistical model. Alert thresholds are tuned with clinicians to limit unnecessary alerts, and urgent findings follow a defined escalation path to a human. When the data are insufficient or of poor quality, the system should say so and fall back to safe guideline defaults instead of producing a confident but unreliable output, a problem noted in real clinic deployments (Beede et al., 2020). Clinicians can always override a recommendation, and overrides are recorded as feedback.

4.3 Explainability, Uncertainty and the Cold-Start Problem

Each alert should come with a short reason in plain language, such as "Raised because fasting readings were above target on most days this week". Where complex models are used, post-hoc explanation methods such as LIME (Ribeiro, Singh and Guestrin, 2016) can help, though simple models are preferred when they perform well enough. The system should also show how confident it is. A new patient has no history (the cold-start problem), so the system begins with clinical information entered by the care team and with the patient's own goals, and then refines its estimates as data accumulate.

5. ILLUSTRATIVE SCENARIO

The following hypothetical example shows how the design might work. Ravi, a 52-year-old with type 2 diabetes, lives in a semi-urban area and has a busy work schedule. During empathy interviews, patients like him said that numbers on a screen made them anxious, that they forgot to record readings, and that they preferred voice prompts in their own language. Clinicians said they had little time and feared too many alerts. The resulting application lets Ravi record readings with a few taps, gives encouragement and simple advice in his language, and shows his progress against his own earlier results.

After two weeks the model notices that his fasting readings are rising and that he has reported missed doses. It first sends a gentle reminder and tip. When the pattern continues, it adds Ravi to the nurse's dashboard with the reason in plain language and an indication of data quality. The nurse calls him, and the doctor decides whether treatment needs to change; the AI does not make that decision. The nurse marks the alert as useful, and this feedback is stored. Testing sessions with patients like Ravi would ask whether the guidance felt supportive and clear, and the results would feed the next iteration.

6. IMPLEMENTATION ROADMAP

Table 4 outlines a phased plan that follows the design thinking stages and keeps the investment in AI development small until the idea has been validated. If the software qualifies as a medical device, applicable regulation, for example the Medical Devices Rules, 2017 in India, must be considered before clinical use.

Table 4. Suggested implementation phases

Phase	Main work	Output
1	User research with ethics approval, personas, problem statement, success and safety criteria	Research report
2	Wireframes and Wizard-of-Oz test with a small group of patients and clinicians	Validated screens and flow
3	Rule-based tool, basic patient model and clinician dashboard; shadow-mode run	Working pilot
4	Supervised pilot with usability, safety and fairness checks; iteration	Pilot findings
5	Introduce statistical models, regulatory review, scale to more sites	Extended system

7. EVALUATION PLAN

Because this paper is conceptual, the following plan is proposed for later studies. A mixed-methods approach is recommended, combining numbers with the voices of users.

- **Clinical outcomes and safety:** relevant health indicators (for the example, glycaemic control), adverse events, and missed or delayed escalations, compared with usual care where ethically possible.
- **Model performance:** sensitivity, specificity, predictive value and calibration, reported overall and for subgroups, preferably in prospective use.
- **Usability:** standard questionnaires such as the System Usability Scale (Brooke, 1996) and think-aloud sessions.
- **Clinician value and workload:** interviews, time spent, and the proportion of alerts accepted, overridden or ignored.
- **Patient engagement and acceptance:** regular use, retention, and perceived usefulness and ease of use (Holden and Karsh, 2010).
- **Fairness:** comparison of performance and alerts across groups (for example by age, gender, language or income).
- **Monitoring after deployment:** checks for performance drift as patients, devices and practices change.

Evaluation should be repeated at each iteration, and negative findings should be reported and used to improve the design.

8. ETHICAL CONSIDERATIONS

Healthcare AI works with sensitive data and affects health, so ethics must be part of the design from the start. Floridi et al. (2018) propose principles of beneficence, non-maleficence, autonomy, justice and explicability, and the World Health Organization (2021) has issued guidance on the ethics and governance of AI for health. In this context they translate into the following practices.

- **Patient safety and accountability:** define who is responsible for decisions, keep audit trails, report incidents and ensure that clinicians retain final authority.
- **Privacy and data protection:** collect minimal data, obtain informed consent and store data securely. Designers should follow applicable law, such as India's Digital Personal Data Protection Act, 2023, the EU General Data Protection Regulation (2016/679) and, in the United States, HIPAA, and should define how long data are kept and how they can be deleted.

- **Fairness:** test for bias so that no group is disadvantaged. Obermeyer et al. (2019) showed that a widely used algorithm underestimated the needs of Black patients because it used health costs as a proxy for need, and Wiens et al. (2019) call for responsible development practices throughout the lifecycle. Practitioners often lack tools and processes for fairness (Holstein et al., 2019), so checks should be planned from the beginning.
- **Transparency:** explain recommendations and limits in simple language to patients and clinicians.
- **Human control:** patients and clinicians can question or override the system, so it supports rather than replaces them.
- **Wellbeing:** avoid alarming messages, labels and scores that cause anxiety or discourage patients.
- **Accessibility and inclusion:** support local languages, low bandwidth, low-cost devices and varied literacy so that digital tools do not widen existing gaps in care.

9. BENEFITS AND CHALLENGES

Benefits. Starting from real needs makes the application more relevant and easier to adopt. Early prototyping reduces the risk of building features no one needs, and involving clinicians builds trust and improves alert quality. Making safety, ethics and fairness part of the testing stage encourages responsible design.

Challenges. Design thinking needs time and access to patients and clinicians, who are often busy. Prototyping AI behaviour before data exist is difficult. Model quality depends on the quantity and quality of data, and results from one hospital or population may not apply to another. Ethics approval, data protection and regulation add effort, and many institutions lack the infrastructure to support data-driven tools. Finally, the effectiveness of the approach must still be shown through empirical studies.

10. CONCLUSION AND FUTURE WORK

This paper proposed a human-centred design thinking approach for developing an AI-based healthcare application. By empathising with patients and clinicians, defining the right problem, deciding the division of work between AI and humans, prototyping early and testing repeatedly, designers can build applications that are not only technically capable but also safe, understandable, fair and useful in real care settings. Future work should implement a prototype following this framework, compare different modelling techniques, and evaluate the effect on clinical outcomes, clinician workload, engagement and fairness with real users.

REFERENCES

Altman, M., Huang, T. T. K., & Breland, J. Y. (2018). Design thinking in health care. *Preventing Chronic Disease, 15*, Article 180128.

Amershi, S., Weld, D., Vorvoreanu, M., et al. (2019). Guidelines for human-AI interaction. In *Proceedings of CHI 2019*. ACM.

Beede, E., Baylor, E., Hersch, F., et al. (2020). A human-centered evaluation of a deep learning system deployed in clinics for the detection of diabetic retinopathy. In *Proceedings of CHI 2020*. ACM.

Brooke, J. (1996). SUS: A quick and dirty usability scale. In P. W. Jordan et al. (Eds.), *Usability evaluation in industry* (pp. 189–194). Taylor & Francis.

Brown, T. (2008). Design thinking. *Harvard Business Review, 86*(6), 84–92.

Cai, C. J., Winter, S., Steiner, D., Wilcox, L., & Terry, M. (2019). "Hello AI": Uncovering the onboarding needs of medical practitioners for human-AI collaborative decision-making. *Proceedings of the ACM on Human-Computer Interaction, 3*(CSCW), Article 104.

Deci, E. L., & Ryan, R. M. (2000). The "what" and "why" of goal pursuits: Human needs and the self-determination of behavior. *Psychological Inquiry, 11*(4), 227–268.

Dorst, K. (2011). The core of 'design thinking' and its application. *Design Studies, 32*(6), 521–532.

Esteva, A., Robicquet, A., Ramsundar, B., et al. (2019). A guide to deep learning in healthcare. *Nature Medicine, 25*(1), 24–29.

Floridi, L., Cowls, J., Beltrametti, M., et al. (2018). AI4People: An ethical framework for a good AI society. *Minds and Machines, 28*(4), 689–707.

Gulshan, V., Peng, L., Coram, M., et al. (2016). Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA, 316*(22), 2402–2410.

Holden, R. J., & Karsh, B.-T. (2010). The technology acceptance model: Its past and its future in health care. *Journal of Biomedical Informatics, 43*(1), 159–172.

Holstein, K., Wortman Vaughan, J., Daumé III, H., Dudík, M., & Wallach, H. (2019). Improving fairness in machine learning systems: What do industry practitioners need? In *Proceedings of CHI 2019*. ACM.

International Organization for Standardization. (2019). *ISO 9241-210:2019 Ergonomics of human-system interaction — Part 210: Human-centred design for interactive systems*. ISO.

Jiang, F., Jiang, Y., Zhi, H., et al. (2017). Artificial intelligence in healthcare: Past, present and future. *Stroke and Vascular Neurology, 2*(4), 230–243.

Kelley, J. F. (1984). An iterative design methodology for user-friendly natural language office information applications. *ACM Transactions on Office Information Systems, 2*(1), 26–41.

Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G., & King, D. (2019). Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine, 17*, Article 195.

Liedtka, J. (2018). Why design thinking works. *Harvard Business Review, 96*(5), 72–79.

Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science, 366*(6464), 447–453.

Plattner, H., Meinel, C., & Leifer, L. (Eds.). (2011). *Design thinking: Understand, improve, apply*. Springer.

Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. *New England Journal of Medicine, 380*(14), 1347–1358.

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. In *Proceedings of KDD 2016*. ACM.

Roberts, J. P., Fisher, T. R., Trowbridge, M. J., & Bent, C. (2016). A design thinking framework for healthcare management and innovation. *Healthcare, 4*(1), 11–14.

Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence, 1*, 206–215.

Sendak, M., Ratliff, W., Sarro, D., et al. (2020). Real-world integration of a sepsis deep learning technology into routine clinical care: Implementation study. *JMIR Medical Informatics, 8*(7), e15182.

Shortliffe, E. H., & Sepúlveda, M. J. (2018). Clinical decision support in the era of artificial intelligence. *JAMA, 320*(21), 2199–2200.

Sutton, R. T., Pincock, D., Baumgart, D. C., et al. (2020). An overview of clinical decision support systems: Benefits, risks, and strategies for success. *npj Digital Medicine, 3*, Article 17.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science, 12*(2), 257–285.

Tonekaboni, S., Joshi, S., McCradden, M. D., & Goldenberg, A. (2019). What clinicians want: Contextualizing explainable

machine learning for clinical end use. In *Proceedings of Machine Learning for Healthcare, PMLR 106*, 359–380.

Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine, 25*(1), 44–56.

Vasey, B., Nagendran, M., Campbell, B., et al. (2022). Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. *Nature Medicine, 28*, 924–933.

Wiens, J., Saria, S., Sendak, M., et al. (2019). Do no harm: A roadmap for responsible machine learning for health care. *Nature Medicine, 25*(9), 1337–1340.

World Health Organization. (2021). *Ethics and governance of artificial intelligence for health: WHO guidance*. WHO.

Yang, Q., Steinfeld, A., Rosé, C., & Zimmerman, J. (2020). Re-examining whether, why, and how human-AI interaction is uniquely difficult to design. In *Proceedings of CHI 2020*. ACM.

Yu, K.-H., Beam, A. L., & Kohane, I. S. (2018). Artificial intelligence in healthcare. *Nature Biomedical Engineering, 2*(10), 719–731.