



Ethical Frameworks for AI-Driven Multimodal Predictive Models in Personalized Healthcare

Sai Venkat Mandalapu¹, Rahul Nunna², Aarya Reddy Pamudurthy¹, Mohammed Sarfaraz³, Anulekha Chegoni⁴, Shruti Bikkumalla⁵

¹*School of Engineering, Brown University, Providence, Rhode Island, USA.*

²*Department of Biomedical Engineering, University of Michigan, Ann Arbor, Michigan, USA.*

³*Biomedical Engineering, College of Engineering & Computer Science, Wright State University, Dayton, Ohio, USA.*

⁴*Department of Information Science, University of North Texas, Denton, Texas, USA.*

⁵*Department of Biotechnology, Tandon School of Engineering, New York University, Brooklyn, New York, USA.*

Received: 12th January, 2026; Revised: 16th February, 2026; Accepted: 10th March, 2026; Available Online: 04th April, 2026

ABSTRACT

The integration of artificial intelligence (AI) and Internet of Medical Things (IoMT) into personalized healthcare introduces transformative opportunities alongside profound ethical challenges threatening patient autonomy, data privacy, and health equity. This review examines core ethical principles – autonomy, beneficence, non-maleficence, and justice, as they apply to AI-driven healthcare systems. We analyze informed consent in AI-supported environments, algorithmic bias and fairness, accountability and transparency, privacy preservation through federated learning, and equitable access. Drawing from international regulatory frameworks, including the WHO guidelines, the EU AI Act, and GDPR, we present evidence-based governance strategies. Robust ethical frameworks must be embedded at every development stage from conception through longitudinal surveillance. This review synthesizes 2021–2025 evidence providing actionable recommendations for researchers, clinicians, and policymakers advancing ethically responsible AI in personalized medicine.

Keywords: Artificial intelligence, Healthcare ethics, Algorithmic fairness, Health equity, Regulatory governance.

International Journal of Health Technology and Innovation (2026)

How to cite this article: Mandalapu SV, Nunna R, Pamudurthy AR, Sarfaraz M, Chegoni A, Bikkumalla S. Ethical Frameworks for AI-Driven Multimodal Predictive Models in Personalized Healthcare. International Journal of Health Technology and Innovation. 2026;5(1):3-7.

Doi: 10.60142/ijhti.v5i01.02

Source of support: Nil.

Conflict of interest: None

INTRODUCTION

The convergence of artificial intelligence, wearable devices, and Internet of Medical Things has fundamentally altered healthcare delivery, shifting from reactive episodic care toward continuous, individualized health management. Personalized healthcare systems leverage multimodal sensor data, electronic health records, genetic information, and AI algorithms to tailor treatment strategies to individual patient characteristics^{1,2}. These systems hold exceptional promise for early disease detection and prevention.

However, technological advancement has outpaced ethical and regulatory frameworks. Unlike traditional medical interventions, AI systems continuously learn from new data, evolving in ways that challenge conventional accountability structures^{3,4}. When algorithms influence

treatment decisions, fundamental questions emerge: How is this decision made? Why was this intervention recommended? Who bears responsibility for errors? Can patients truly consent to AI-supported care? These questions lie at contemporary healthcare's ethical core^{5,6}.

Multiple ethical challenges demand systematic attention. Algorithmic bias—where training datasets underrepresent diverse populations, threatens to perpetuate healthcare disparities^{7,8,9}. Deep learning model's opacity undermines clinician understanding and patient comprehensions^{10,11}. Privacy concerns intensify as continuous wearable monitoring collects intimate physiological data^{12,13}. The digital divide risks creating a two-tiered healthcare system concentrating benefits among affluent populations^{14,15,16}. These challenges require collaborative frameworks developed by researchers,

clinicians, patients, ethicists, and policymakers^{17,18}.

METHODOLOGY

Search Strategy and Data Sources

This narrative review synthesized literature on ethical frameworks for AI-driven multimodal predictive models in personalized healthcare published between January 2021 and March 2025. We conducted comprehensive searches across PubMed, PubMed Central, Google Scholar, Scopus, and IEEE Xplore databases. Search terms included combinations of: “artificial intelligence,” “machine learning,” “deep learning,” “healthcare ethics,” “medical ethics,” “algorithmic bias,” “algorithmic fairness,” “informed consent,” “explainability,” “transparency,” “federated learning,” “privacy preservation,” “health equity,” “digital divide,” “personalized medicine,” “precision medicine,” “Internet of Medical Things,” “IOMT,” “wearable devices,” “regulatory frameworks,” “GDPR,” “HIPAA,” and “EU AI Act.” Boolean operators (AND, OR) were used to combine terms systematically.

Inclusion and Exclusion Criteria

Sources were included if they: (1) addressed ethical, legal, or governance dimensions of AI applications in healthcare; (2) were published between 2021-2025 to ensure currency given rapid technological evaluation; (3) were peer-reviewed empirical studies, systematic or narrative reviews, policy documents from regulatory bodies, or official guidance from professional organizations; (4) were published in English; and (5) provided actionable insights for researchers, clinicians, institutions, or policymakers. We excluded: (1) purely technical papers without ethical analysis; (2) opinion pieces or editorials lacking empirical grounding; (3) conference abstracts without full-text availability; and (4) studies focused exclusively on non-healthcare AI applications.

Selection Process and Data Extraction

Initial database searches yielded 247 potentially relevant records. After removing duplicates (n=62), two authors independently screened titles and abstracts, identifying 108 articles for full-text review. Following full-text assessment, 37 sources met inclusion criteria and were retained for synthesis. Disagreements were resolved through discussion until consensus was reached. Data extraction focused on: (1) ethical principles addressed (autonomy, beneficence, non-maleficence, justice); (2) specific ethical challenges identified (bias, consent, privacy, transparency, equity); (3) proposed governance mechanisms; (4) regulatory framework analysis; and (5) stakeholder-specific recommendations. Given the conceptual breadth and interdisciplinary nature of ethical AI governance, we adopted a narrative synthesis approach rather than systematic review methodology (e.g., PRISMA), which is appropriate for integrating diverse theoretical perspectives and policy analyses^{17,19}.

Quality Assessment

While formal quality appraisal tools designed for empirical

studies (e.g., CASP, Newcastle-Ottawa Scale) were not applicable to all source types, we prioritized peer-reviewed publications from indexed journals, official regulatory documents, and guidance from established professional organizations (WHO, AMA, European Commission). Sources were evaluated for relevance, methodological rigor when applicable, and contribution to actionable recommendations.

Core Ethical Principles in AI Healthcare

Autonomy and Informed Consent

Autonomy: the right of individuals to make informed choices about their healthcare, represents a foundational bioethical principle^{17,19,20}. In AI-supported healthcare, respecting autonomy requires that patients understand how their data will be collected, processed, used in algorithmic systems, and how these systems influence clinical recommendations^{21,22}. Complex neural networks challenge traditional consent frameworks, patients cannot realistically understand mathematical operations of AI systems yet deserve comprehensible explanations of how recommendations were derived^{10,11,23}.

Emerging concepts of “explicability” provide clinically meaningful reasoning rather than technical transparency alone^{10,23}. Furthermore, as AI systems continuously learn from new data, risks and benefits change dynamically, necessitating ongoing consent rather than one-time enrollment^{3,4,29}. Meaningful informed consent requires comprehensible information, uncertainty acknowledgement particularly for underrepresented populations^{9,21,22} ongoing re-consent mechanisms, opt-out provisions, and transparency about conflicts of interest²¹⁻²⁴.

Beneficence, Non-maleficence, and Justice

Beneficence obligates maximizing patient benefits; non-maleficence demands avoiding harm. In AI applications, these principles require rigorous predeployment validation, continuous post-market surveillance, and mechanisms to identify and correct unintended harms^{17,19,20}. However, defining benefit and harm proves complex systems optimized for sensitivity increase false positives leading to unnecessary treatments; models trained on general populations may perform poorly for subgroups, inadvertently harming them^{7,8,20}.

Justice demands equitable distribution of healthcare benefits and fair risk burden sharing. In personalized medicine, justice requires ensuring diverse population representation in training data^{7,8}, validating algorithm performance across demographic subgroups, making AI-enabled treatments accessible regardless of socioeconomic status^{14,15}, and ensuring fair rather than biased decision making^{5,17}. Precision medicine concentrated among affluent populations risks widening health inequities^{14,23}, while underdiagnosis in underserved populations excludes them from advanced treatments^{11,25}.

Algorithmic Bias and Health Equity

Algorithmic bias emerges from multiple sources throughout the AI lifecycle^{7,8,9}. Training data bias occurs when datasets

underrepresent certain demographics, most healthcare AI uses data predominantly from developed countries with younger, healthier, lighter-skinned populations^{7,9}. Label bias arises when clinical diagnoses reflect historical inequities; cardiovascular disease may be underrecognized in women, biasing algorithms trained on such data^{7,9}. Representation bias emerges when specific populations are systematically omitted^{8,9}. Measurement bias reflects inherent sensor performance differences across populations, PPG sensors demonstrate lower accuracy in individuals with darker skin pigmentation due to melanin absorption^{8,9}.

Biased AI systems perpetuate, amplify, or create new health inequities. A landmark study revealed U.S. hospital algorithms identifying patients needing additional care exhibited significant bias against Black patients, for any given predicted risk level, Black patients were sicker, indicating the algorithm systematically deprioritized them. This resulted from using healthcare costs as a proxy for medical need, favoring wealthier populations with greater healthcare utilization^{7,9}.

Comprehensive bias mitigation requires data diversity and representation through actively recruiting diverse populations in data collection⁷⁻⁹, fairness metrics and auditing computing performance across demographic groups⁷⁻⁹, participatory design engaging affected communities in development^{9,26,27}, and transparency documenting training data characteristics and limitations.^{7,8,19}

Privacy, Security, and Regulatory Compliance

Privacy Risks and Preservation Technologies

Health data collected through IoMT devices reveal intimate details: activity patterns, location history, medication adherence, mental health status. Unlike discrete medical records, continuous monitoring often occurs semi-consciously, generating data streams users don't fully appreciate. Privacy risks include unauthorized access through breaches, secondary data use beyond original consent, inference of sensitive information from seemingly innocuous measurements, and discrimination by insurers based on wearable-derived metrics^{12,25}.

Federated learning addresses privacy concerns by training AI models on decentralized data while keeping raw information local^{12,13, 18}. Rather than uploading patient data to central servers, federated systems train models at each institution on local data, then aggregate only model parameters^{12,18}. This enables multi-institutional collaboration impossible under traditional architectures^{12,18}. Differential privacy adds protection by injecting calibrated noise into model updates before aggregation, preventing reconstruction of individual training examples^{12,13}. Studies demonstrate federated learning with differential privacy achieves 99% of centralized learning accuracy while providing rigorous privacy guarantees^{18,28,29}.

Regulatory Frameworks

The European Union's General Data Protection Regulation mandates explicit user consent for data processing, purpose limitation, data minimization, and user rights to access,

correct, and delete personal information^{12,30,31}. GDPR applies to any organization processing personal data of EU residents regardless of location, creating extraterritorial obligations^{12,31}. The U.S. Health Insurance Portability and Accountability Act establishes privacy and security standards for protected health information held by covered entities and business associates, requiring encryption, access controls, audit trails, and breach notification^{12,30}. However, HIPAA's scope is narrower than GDPR – consumer wearables for personal wellness often fall outside coverage^{12,30}.

The EU AI Act, effective 2025, imposes stringent requirements on high-risk healthcare AI systems, including rigorous risk assessment, high-quality diverse training data, comprehensive documentation, human oversight mechanisms, and transparency measures^{30,32}.

Accountability, Transparency, and Explainability

Traditional clinical accountability assigns responsibility to the physician as ultimate decision maker^{10,19}. AI complicates this structure, responsibility becomes ambiguous among developers, deploying institutions, clinicians and patients^{10,11,19}. Robust frameworks must clearly delineate responsibilities: developers must develop unbiased, validated algorithms with documented accuracy across populations^{5,17}; institutions must ensure clinical validation before implementation, provide clinician training, maintain audit trails, and implement escalation pathways¹⁷⁻¹⁹; clinicians must understand recommendations sufficiently to evaluate appropriateness and maintain authority to override^{10,11,19}; patients must receive comprehensible information enabling informed consent²¹⁻²³.

Explainability remains contentious, advocates argue physicians cannot ethically rely on opaque recommendations^{10,23}; critics counter that physicians themselves often cannot explain diagnostic intuitions. Emerging consensus distinguishes technical transparency (understanding mathematical operations) from clinical explicability (providing contextually meaningful reasoning). Methods improving explicability include SHAP values quantifying feature contributions, attention visualization highlighting relevant inputs, and AI-supported shared decision-making integrating algorithmic insights with patient values^{23,33}.

Regulatory bodies increasingly mandate explainability. The American Medical Association adopted policy requiring explainable clinical AI with independent third-party verification^{11,34}. The EU AI Act requires transparency documentation for high-risk healthcare AI^{30, 32}. However, tension exists between explainability requirements and accuracy, simpler, more interpretable models often perform worse than complex black-box systems^{10,11}. Regulatory frameworks must balance these tradeoffs^{5,10,11}.

Privacy-Preserving Governance and Stakeholder Collaboration

Robust AI governance requires multiple interconnected components. Ethical oversight through AI ethics committees

representing clinical, technical, ethical, and community perspectives provides multidisciplinary review^{5,17,18}. Transparency mandates require documentation of training data characteristics, model performance across demographic subgroups, identified limitations, and validation evidence^{19,30,32}. Fairness standards explicitly define acceptable performance disparities, recognizing perfect equality is impractical while systematic differences for specific populations is unethical^{5,7,8}. Human oversight ensures AI supports rather than replaces clinician decision making^{5,10,11}. Continuous monitoring implements real-time post-deployment surveillance with escalation pathways for concerning findings.^{5,18,30}

Ethical AI governance requires sustained stakeholder collaboration. Clinicians contribute domain expertise, understand workflow challenges, and identify unintended consequences^{5,17,18}. Data scientists and engineers must understand ethical constraints beyond accuracy optimization. Ethicists facilitate deliberation about values and identify emerging concerns. Patients and communities provide perspectives on which outcomes matter, acceptable risks, and real-world effects^{5,17,18,27}. Policymakers establish enabling frameworks^{5,17,30}. Biostatisticians conduct fairness audits and design validation studies.^{5,7,8}

Equitable Access and the Digital Divide

The digital divide-disparities in access to reliable internet, digital devices, and digital literacy, threatens creating a two-tiered healthcare system^{14-16,35}. Approximately 21 million Americans lack broadband internet access, with disparities concentrated in rural areas and low-income communities. Low-income individuals are significantly less likely to own smartphones; elderly and non-English speaking populations struggle with digital literacy^{14-16,35}. The digital divide intersects with existing social determinants of health, creating cascading disadvantage¹⁴⁻¹⁶.

Advancing health equity requires proactive strategies. Infrastructure investment expanding broadband in rural/underserved areas and subsidizing services for low-income households proves essential^{30,32}. Device accessibility programs and manufacturer partnerships ensuring compatibility with older devices expand access^{30,32}. Community-based digital literacy training on telehealth platforms, health apps, and wearables addresses literacy barriers^{30,32}. Culturally responsive AI ensures linguistic and cultural appropriateness for diverse populations^{30,32}. Participatory research engaging underserved communities as co-researchers ensures their needs prioritize algorithm design^{30,32}.

Emerging Challenges and Future Directions

Foundation models, large language models trained on diverse data: present novel ethical challenges beyond task-specific AI systems. Their generality creates uncertainty about appropriate use cases; performance on new tasks is difficult to predict; training data often includes unconsented health information.^{17,20} Governing foundation models requires approaches

fundamentally different from traditional validation^{17,20}. Future governance must establish processes for identifying appropriate healthcare applications, validating performance in specific clinical contexts before deployment, monitoring post-deployment performance, and discontinuing use when performance degrades^{17,20}.

Future AI governance must prioritize health equity as a central objective rather than treating it as an afterthought^{9,14,15,19}. This requires equitable data representation through actively recruiting diverse populations^{7,9}, community-engaged research ensuring priorities shape algorithm design^{9,27,35}, accessible technology functioning with limited bandwidth and lower digital literacy^{14,16,35}, and distributive justice ensuring equitable access to AI-enabled treatments^{14,25}.

Ethical frameworks must adapt continuously as AI systems evolve^{5,17,18}. Ongoing fairness auditing assesses algorithm performance across demographic groups post-deployment.^{7-9,19} Surveillance for unintended consequences monitors real-world harms^{5,17,18}. Stakeholder feedback mechanisms create pathways for concerns¹⁷⁻¹⁹. Adaptive governance allows frameworks to evolve as evidence accumulates¹⁷⁻¹⁹.

CONCLUSIONS AND RECOMMENDATIONS

The integration of AI and IoMT into personalized healthcare offers unprecedented opportunities for early disease detection, precision treatment, and improved outcomes^{1,2,20}. Realizing these benefits ethically requires sustained commitment to robust governance frameworks addressing autonomy, justice, privacy, and transparency^{5,17}.

Core findings establish ethical principles provide enduring foundations requiring careful specification for technical systems^{17,19,20}. Regulatory frameworks establish necessary minimums; organizations must implement more stringent governance than legally required^{18,19,30,32}. Algorithmic fairness demands active effort-bias will not self-correct⁷⁻⁹. Privacy is technically achievable through federated learning and differential privacy^{12,13,18,28,29}. Transparency requires reimagining explainability as clinical understanding, not technical transparency^{10,11,23,33}. Equity demands proactive intervention against widening disparities^{14,15,16,35}. Multidisciplinary collaboration is essential^{5,17,18}.

For researchers and developers: prioritize dataset diversity, conduct fairness audits across demographic groups, engage affected communities in design, document limitations transparently, and implement privacy-preserving techniques by default^{5,7-9}. For healthcare institutions: establish AI ethics committees, implement governance frameworks aligned with emerging regulations, provide clinician training, maintain real-time performance monitoring, and create stakeholder feedback mechanisms^{5,17,18,30}. For clinicians: maintain skepticism toward algorithmic recommendations, understand system limitation for underrepresented populations, engage in shared decision-making incorporating patient values, and report concerning patterns^{5,10,11,23}. For policymakers: establish accountability frameworks, mandate fairness auditing and continuous monitoring, ensure accessible challenge pathways

for algorithmic decisions, invest in digital infrastructure addressing equity, and foster international governance harmonization.^{5,18,30,32} For patients and communities: engage in participatory research, provide feedback on system effects, advocate for equity and accessibility, and demand transparency and accountability.^{14,17,18,27}

REFERENCES

1. PMC11047988 (2024). The role of AI in hospitals and clinics: Transforming healthcare. PubMed Central.
2. PMC11520245 (2022). Revolutionizing healthcare: Impact of modern AI technologies. PubMed Central.
3. PMC10930608 (2024). Global regulatory frameworks for artificial intelligence in medical devices. PubMed Central.
4. ScienceDirect (2024). Privacy preservation for federated learning in healthcare. *Healthcare Technology Review*.
5. *Frontiers in Digital Health* (2025). Ethics of AI in healthcare: A scoping review. 11(1662642).
6. CDC Preventing Chronic Disease (2024). Health equity and ethical considerations in AI. 21(24):0245.
7. *Nature Scientific Reports* (2025). Bias recognition and mitigation strategies in AI. 15:1103-1505.
8. *Health Policy and Technology* (2023). Addressing algorithmic bias and health inequities. 12(1):1-15.
9. PMC12325396 (2025). Algorithmic bias in public health AI: Threat to equity. PubMed Central.
10. *Theoretical Medicine and Bioethics* (2022). Explainability in medicine in AI-based clinical decision support. 45(5):367-388.
11. PMC11391805 (2024). Should AI models be explainable to clinicians? PubMed Central.
12. *Frontiers in Drug Safety and Regulation* (2025). Federated learning: Privacy-preserving data-centric healthcare. 10(1579922).
13. *World Journal of Advanced Research and Reviews* (2025). Privacy-preserving federated learning for healthcare. 26(2):3263-3272.
14. *JAMA Network Open* (2024). Tackling digital equity needs in health care. 7(11):e2426216.
15. Johns Hopkins Bloomberg School of Public Health (2025). Bridging the digital divide in healthcare: Equity framework.
16. PMC11705165 (2024). Digital health divide: Reducing health disparities through technology. PubMed Central.
17. PMC12075486 (2025). Scaling enterprise AI in healthcare: Governance role in risk management. PubMed Central.
18. PMC8515002 (2021). Addressing bias in big data and AI for healthcare. NIH.
19. *Frontiers in Digital Health* (2025). Ethics of AI in healthcare: A scoping review. 11(1662642).
20. PMC12076083 (2025). Ethical and legal considerations in healthcare AI. PubMed Central.
21. PMC11977975 (2025). Ethical challenges in AI integration into healthcare. PubMed Central.
22. *Nature npj Digital Medicine* (2025). A practical framework for AI implementation and review. 8(2025):1900-y.
23. PMC12331219 (2025). AI-supported shared decision-making framework. PubMed Central.
24. *Journal of Medical Ethics* (2024). Informed consent in continuously learning AI systems. 50(6):434-442.
25. *Health Equity Trends* (2025). Health equity: Provider burnout and digital divides. LexisNexis Risk.
26. *Implementation Science* (2025). Participatory design in healthcare AI development. 20(1):1-14.
27. *Community-Based Participatory Research Methods* (2024). Engaging marginalized populations in AI research. 29(2):145-168.
28. *Nature Communications* (2025). Fairness-aware and privacy-preserving federated learning. 16(1):58055.
29. *Advanced Methods in Healthcare IT* (2024). Homomorphic encryption and federated learning. 18(3):234-251
30. *Superblocks* (2025). AI Governance in Healthcare: Frameworks & Best Practices. Healthcare Technology Guide.
31. European Commission (2024). GDPR compliance in healthcare AI systems. Digital Single Market Report.
32. *Nature* (2024). Global regulatory frameworks for AI in healthcare: EU AI Act. 42(8):1-12.
33. PMC12331219 (2025). AI reasoning versus explainability in shared decision-making. PubMed Central.
34. American Medical Association (2025). Transparency policy on clinical AI tools. AMA Resolution.
35. *Health Equity Framework* (2025). Bridging the digital divide in healthcare: Community strategies. Public Health Reports.