JOURNAL OF LAW, INFORMATION AND SCIENCE
ISSN 0729-1485• Vol. 4, Issue 2 (2026) • https://jlisjournal.com/

REVIEW ARTICLE Article ID: JLIS/2026/4/2/009

AI and Healthcare Transformation: Economic, Legal, and Equity Perspectives

Amit Kumar [1] , Rituraj Soni [2]

[1] *Assistant Professor, Department of CCSIT, Teerthanker Mahaveer University, Moradabad, Uttar Pradesh – 244001, India.*
[2] *Assistant Professor, Department of CSE, Engineering College, Bikaner, Rajasthan, India.*

Received: 22 June 2026 **Revised:** 29 July 2026 **Accepted:** 25 August 2026 **Published:** 30 Sep 2026

ABSTRACT

Machine learning (ML) is increasingly presented as the enabling technology of predictive, personalized medicine. This review examines its promise and its limits from a perspective informed by economic history: health improvements in Europe have historically depended less on single inventions than on the institutions, complementary investments and distributional choices that surround them. Drawing on peer-reviewed clinical, computational and policy literature, the article maps the present scenario across four dimensions: clinical applications, data infrastructures, evidence and validation, and equity and regulation. It shows that ML now performs at or near expert level in narrow diagnostic and prognostic tasks, yet that real-world evidence remains uneven, that training data under-represent many populations, and that widely used algorithms have encoded structural inequality. The article then sets out a forward-looking agenda covering privacy-preserving collaboration, multimodal and foundation models, interpretability, economic evaluation and the productivity J-curve, and equity-by-design governance within the European regulatory environment. It concludes that predictive ML will contribute to personalized and equitable health systems only if data, validation, incentives and accountability are treated as public infrastructure rather than as by-products of algorithm development.

Keywords: *machine learning; predictive healthcare; personalized medicine; health equity; algorithmic fairness.*

© 2026. The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.

1. Introduction

The long-run improvement of European health is one of the central achievements of modern economic history. Fogel (2004) linked the decline of premature death after 1700 to gains in nutrition and the growth of human capital, while Deaton (2013) portrayed the twentieth-century rise in life expectancy as an escape whose benefits were unevenly shared across and within nations. The distributional character of that escape matters for the present argument. Cutler and McClellan (2001) showed that, in the late twentieth century, technological change in medicine delivered gains that were large relative to cost in some conditions but not in others, and Marmot (2005) demonstrated that the social gradient in health persisted even where care was universally financed. European welfare states, whose solidaristic logic Baldwin (1990) traced to the interaction of class interests and social insurance, have therefore always faced the dual task of raising average health and narrowing inequalities in it.

Health systems now confront this task under demographic ageing, chronic disease and what Baumol (2012) called the cost disease: labour-intensive services whose productivity grows

slowly become relatively more expensive over time. Machine learning, the family of methods by which algorithms learn patterns from data rather than following explicitly programmed rules, is widely regarded as a means of bending this cost curve while making care more tailored to the individual (Beam & Kohane, 2018; Rajkomar et al., 2019; Topol, 2019). Precision medicine, in the sense articulated by Collins and Varmus (2015), aspires to match prevention and treatment to individual genetic, environmental and lifestyle characteristics, and ML supplies much of the analytical machinery required to do so at scale.

Economic historians will recognise the pattern of expectation. David (1990) explained why the electric dynamo took decades to raise measured productivity: the gains awaited the reorganisation of factories around the new power source. Brynjolfsson et al. (2021) generalised the point in the productivity J-curve, in which general-purpose technologies first impose costs of complementary investment and only later yield measurable returns. Applied to health, the analogy cautions against equating algorithmic accuracy in a laboratory setting with system-level benefit. It also raises the question that motivates this review: under what conditions will predictive ML produce personalized care that is also equitable, rather than reproducing or widening existing gaps?

The article is a narrative review. It synthesises influential and recent peer-reviewed studies, landmark validation and audit studies, and European regulatory instruments, selected for relevance to prediction, personalization and equity. It does not attempt a systematic meta-analysis. Section 2 describes the present scenario, Section 3 outlines future perspectives, and Section 4 concludes.

2. Present Scenario

2.1 Technological base and clinical applications

Current predictive healthcare rests on three technical developments: deep learning applied to images and signals, sequence models applied to longitudinal electronic health records (EHRs), and, more recently, large language models. In imaging, Gulshan et al. (2016) reported that a deep neural network detected referable diabetic retinopathy in retinal photographs with high sensitivity and specificity, and Esteva et al. (2017) showed that a convolutional network classified skin lesions at a level comparable to a panel of 21 board-certified dermatologists. Such studies helped establish the credibility of ML in narrowly defined perceptual tasks (Esteva et al., 2019; Yu et al., 2018).

Prediction from structured records is the second pillar. Rajkomar et al. (2018b) trained deep learning models on EHRs from two U.S. hospitals to predict in-hospital mortality, unplanned readmission, prolonged length of stay and discharge diagnoses, outperforming traditional clinical benchmarks. This line of work moves the unit of analysis from a single image to the patient trajectory, which is precisely what personalization requires. Multimodal systems that combine imaging, genomics, laboratory data and clinical notes are the current frontier (Acosta et al., 2022), while large language models have shown that general-purpose systems can encode substantial clinical knowledge (Singhal et al., 2023). Rajkomar et al. (2019) and Topol (2019) provide accessible syntheses of these developments; Table 1 summarises major application domains, illustrative evidence and the equity considerations that accompany each.

Table 1. Principal application domains of machine learning in predictive healthcare

Domain	Typical data and methods	Illustrative evidence	Principal equity consideration
Medical imaging	Retinal, dermatological and radiological images;	Expert-level detection of diabetic retinopathy and skin cancer in retrospective validation	Underdiagnosis in under-served groups documented in chest radiography (Seyyed-

	convolutional networks	(Esteva et al., 2017; Gulshan et al., 2016)	Kalantari et al., 2021)
Risk stratification	EHRs and claims data; gradient boosting, recurrent and deep networks	Accurate prediction of mortality, readmission and length of stay (Rajkomar et al., 2018b)	Label choice can embed structural bias, as with cost used as a proxy for need (Obermeyer et al., 2019)
Early warning	Streaming vital signs and laboratory data; proprietary scoring models	Widely deployed sepsis model showed poor discrimination and missed many cases on external validation (Wong et al., 2021)	Performance may vary across hospitals and populations without local calibration (Kelly et al., 2019)
Precision medicine	Genomic and multi-omic data; multimodal models	Integration of imaging, omics and records for individualised prediction (Acosta et al., 2022; Collins & Varmus, 2015)	Genomic resources dominated by European-ancestry participants (Sirugo et al., 2019)
Clinical language systems	Clinical notes and dialogue; large language models	Encoded clinical knowledge and question-answering capability (Singhal et al., 2023)	Risk of propagating biases and errors at scale; need for governance (WHO, 2021)

Source: **Compiled by the author from the studies cited in the table.**

2.2 Data infrastructures

Prediction depends on data, and data infrastructures are unevenly developed. The UK Biobank, which follows roughly half a million participants with linked health records, and its deep phenotyping and genomic data (Sudlow et al., 2015; Bycroft et al., 2018) illustrate the value of large, well-curated population resources; these in turn support the personalization agenda by enabling risk models that integrate genetic and clinical information. Yet even such resources reflect selective participation, and Sirugo et al. (2019) reported that the great majority of participants in genome-wide association studies were of European ancestry, limiting the transferability of genetic predictors to other populations.

At European level, the European Health Data Space Regulation (Regulation (EU) 2025/327) creates a framework for the primary use of health data by patients and clinicians and for secondary use in research, innovation and policy-making. Its implementation is phased, and its consequences for ML development will depend on interoperability, data quality and the capacity of national authorities. The historical parallel is instructive: the nineteenth-century statistical movement and the later construction of national health insurance registers show that reliable population data are an institutional achievement, not a natural resource.

2.3 Evidence, validation and implementation

The evidence base is strongest for retrospective accuracy and weakest for prospective, real-world impact. Kelly et al. (2019) identified key obstacles to clinical impact, including limited external validation, poor integration into workflows and unclear accountability. A prominent illustration is the external validation of a proprietary sepsis prediction model implemented in numerous hospitals: Wong et al. (2021) found that its discrimination was substantially poorer than vendor claims and that it failed to identify most patients who developed sepsis. Reporting and development standards have accordingly received attention (Wiens et al., 2019), and

Collins and Moons (2019) called for more transparent reporting of prediction-model studies based on ML.

The regulatory record reflects the same maturity gap. Muehlematter et al. (2021) compared approvals of AI and ML-based medical devices in the United States and Europe between 2015 and 2020 and noted that many devices reached the market through pathways with limited premarket clinical evidence, and that transparency about supporting evidence in Europe was limited. For economic historians, the episode recalls earlier phases of technological diffusion in which adoption preceded rigorous evaluation of net benefit (Cutler & McClellan, 2001).

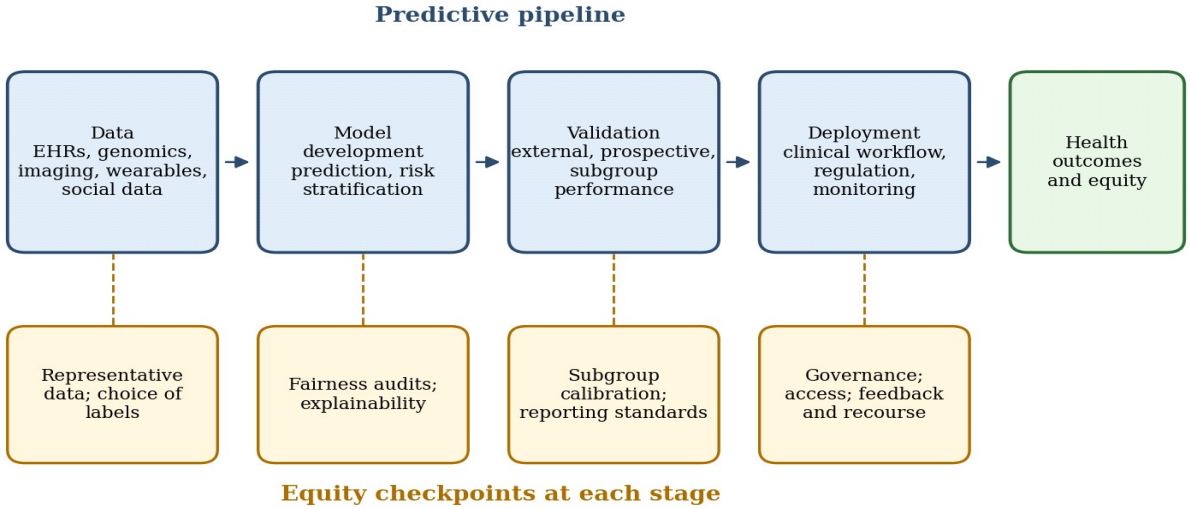


Figure 1. **Conceptual framework: a predictive machine learning pipeline with equity checkpoints at each stage (author's synthesis based on Rajkomar et al., 2018a; Wiens et al., 2019; Chen et al., 2020).**

Organisational conditions also shape present-day outcomes. Tools that perform well in development often struggle where clinical workflows, data definitions and staffing differ from those of the development site (Kelly et al., 2019). Topol (2019) argued that the real value of ML lies in giving clinicians time and information to care for patients, not in replacing them; whether this occurs depends on how alerts are designed, who is accountable for acting on them and whether staff are trained to interpret probabilistic outputs. Systems with weaker digital infrastructure, typically those serving poorer regions, are the least able to make such complementary investments, which suggests that the first wave of ML benefits may accrue disproportionately to already advantaged providers (Baumol, 2012; Brynjolfsson et al., 2021).

2.4 Equity and algorithmic bias

The equity dimension is not peripheral; it is where the promise of ML is most severely tested. Obermeyer et al. (2019) analysed an algorithm widely used to allocate care-management resources and found that, at a given risk score, Black patients were considerably sicker than White patients, because the algorithm used health expenditure as a proxy for need and expenditure was systematically lower for Black patients with equal need. They estimated that correcting the bias would raise the share of Black patients identified for extra help from 17.7% to 46.5%. The case is a textbook example of how an apparently neutral design choice reproduces an historical pattern of unequal access.

Bias arises at multiple points. Seyyed-Kalantari et al. (2021) found that chest radiograph classifiers underdiagnosed women, younger patients, Black and Hispanic patients and those with lower socioeconomic status, meaning that disadvantaged patients were more likely to be incorrectly labelled as healthy. Rajkomar et al. (2018a) described how differences in data, model design and deployment can generate inequitable outcomes, and Chen et al. (2020) argued that ML can just as well be directed at detecting and reducing disparities. Mehrabi et al.

(2021) catalogued sources of bias and fairness definitions, which are often mutually incompatible, so that choosing among them is a normative rather than a purely technical decision. Vokinger et al. (2021) proposed strategies for mitigating bias throughout the development lifecycle. These findings resonate with Marmot's (2005) insistence that health inequalities are socially produced: an algorithm trained on unequal systems will learn unequal patterns unless designed otherwise.

2.5 Governance and regulation in Europe

Ethical and legal scrutiny has intensified (Char et al., 2018; Vayena et al., 2018; WHO, 2021). The Artificial Intelligence Act (Regulation (EU) 2024/1689) classifies many AI systems used in health, including those that are regulated medical devices, as high-risk, subjecting them to requirements on risk management, data governance, documentation, transparency, human oversight and accuracy, with obligations applying in stages. The interplay between the AI Act, the medical devices framework and the European Health Data Space is still being worked out in practice. Compared with the largely market-led pattern of ML adoption elsewhere, the European approach reflects a long tradition of treating health as a collectively governed good, but it also carries risks of compliance burdens that favour large firms over public and academic developers.

3. Future Perspective

3.1 Privacy-preserving and federated learning

A structural problem of ML in health is that data are fragmented across institutions and jurisdictions while models require diversity and scale. Federated learning, in which models are trained across sites without centralising patient data, offers a technical response with obvious relevance to European data-protection law (Kaissis et al., 2020; Rieke et al., 2020). Its promise is twofold: it may allow smaller and lower-resourced hospitals to contribute to and benefit from shared models, thereby widening representation, and it may reduce the concentration of data power in a few institutions. Its limits, including heterogeneity across sites and the need for governance of model updates, must be addressed empirically.

3.2 Multimodal and foundation models

Multimodal models that combine genomic, imaging, clinical and behavioural data are likely to deliver the most personalized predictions (Acosta et al., 2022), and large language models may become interfaces to the record and to clinical guidance (Singhal et al., 2023). The same properties that make them powerful, including scale, opacity and general-purpose deployment, magnify the equity and safety challenges described above. Evaluation of such systems will need to move from single-task benchmarks to prospective, subgroup-specific assessment in the settings where they will operate (Kelly et al., 2019; Wiens et al., 2019).

3.3 Interpretability, accountability and trust

Whether explanations of complex models should be relied upon is contested. Rudin (2019) argued that for high-stakes decisions inherently interpretable models should be preferred to post hoc explanations of black boxes, while Ghassemi et al. (2021) cautioned that current explainability techniques provide less assurance than often assumed and that rigorous validation should take precedence. For equitable health systems, the implication is that accountability must rest on validation, monitoring and clear allocation of responsibility, rather than on the mere availability of an explanation.

3.4 Economic evaluation and the productivity J-curve

From an economic-history perspective, the most important future task is to establish whether and when ML raises health system productivity. The J-curve logic of Brynjolfsson et al. (2021) predicts an initial phase in which expenditure on data infrastructure, workforce training and workflow redesign outweighs measurable benefit, followed by gains once complementary intangible capital accumulates, as illustrated conceptually in Figure 2. David (1990) reached a

similar conclusion for electrification. In health there is a further complication that Baumol (2012) helps to frame: if ML augments but does not substitute for clinical labour, its main effect may be to improve quality rather than to reduce costs, and evaluation must then capture health outcomes and not only expenditure (Cutler & McClellan, 2001).

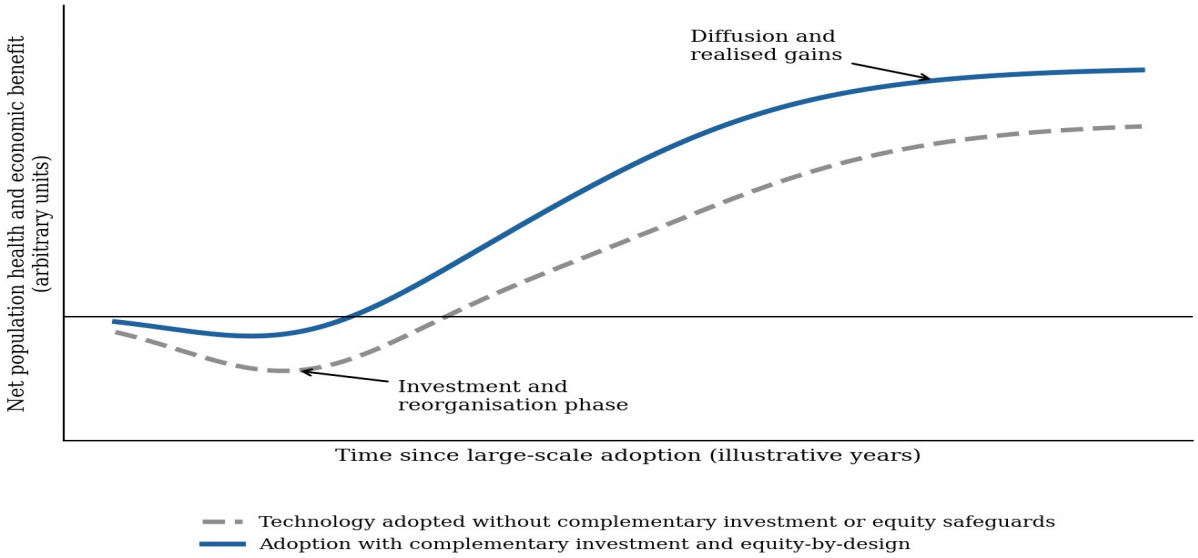


Figure 2. **Illustrative (non-empirical) J-curve of net benefit following large-scale adoption of predictive machine learning, contrasting adoption with and without complementary investment and equity safeguards (conceptual adaptation of David, 1990; Brynjolfsson et al., 2021).**

3.5 Equity by design and a research agenda

Equity must be treated as a design requirement and a public objective. This implies representative data collection, choice of clinically meaningful labels, routine subgroup performance reporting, post-deployment monitoring and mechanisms of redress (Chen et al., 2020; Rajkomar et al., 2018a; Vokinger et al., 2021). It also implies attention to the social determinants that predictive tools cannot alter alone (Marmot, 2005). Table 2 summarises priority actions and the evidence needed to assess them.

Table 2. Priority actions for equitable predictive healthcare

Priority	Action	Evidence needed
Data	Broaden participation and ancestry diversity in biobanks and registries; enable interoperable access through the European Health Data Space	Coverage of population subgroups; transferability of predictors (Sirugo et al., 2019)
Validation	Require external, prospective and subgroup-specific validation before and after deployment	Calibration and error rates by group and by site (Wong et al., 2021)
Learning infrastructure	Support federated and privacy-preserving collaboration among hospitals	Comparative performance and governance studies (Rieke et al., 2020)
Economic evaluation	Assess costs, quality-adjusted outcomes and complementary investment over the J-curve period	Long-run productivity and distributional analysis (Brynjolfsson et al., 2021)
Governance	Align AI Act, medical device and	Implementation studies

	health data rules; fund public and academic developers	of regulatory effects (Muehlematter et al., 2021)

Source: **Author's synthesis.**

Finally, the historical record offers a research agenda of its own. Comparative studies of how earlier medical technologies diffused across European health systems, and how institutions shaped who benefited, can help forecast the distributional consequences of ML and inform policy design.

3.6 Workforce, public trust and the limits of prediction

Two further conditions will influence whether the technological possibilities are realised. The first is the health workforce: training in data literacy and in the critical appraisal of algorithmic outputs is a complementary investment of the kind emphasised by Brynjolfsson et al. (2021), and without it, tools risk being ignored or over-trusted. The second is public trust. Vayena et al. (2018) noted that the legitimacy of ML in medicine depends on transparency about data use and on meaningful accountability, and the WHO (2021) guidance places human autonomy, inclusiveness and equity among its core principles. Trust, once lost through a visible failure or a discriminatory outcome, is expensive to rebuild, which reinforces the case for careful staged deployment.

It is also necessary to recognise the limits of prediction. Accurate forecasts of who will become ill do not in themselves change outcomes; they matter only if effective interventions exist and are equitably accessible. The lesson of the historical literature is that health improvements have come from nutrition, sanitation, education and institutions of social insurance as much as from clinical technology (Deaton, 2013; Fogel, 2004; Marmot, 2005). Predictive ML is best regarded as a complement to these foundations, and as a means of directing scarce resources more effectively toward those in greatest need, rather than as a substitute for them.

4. Conclusion

Machine learning has moved from proof of concept to early deployment in predictive healthcare. Evidence shows expert-level performance in narrow tasks, credible advances in multimodal prediction, and a rapidly evolving European regulatory framework. Yet accuracy is not the same as benefit, and benefit is not the same as equity. Evaluation remains weaker than development, data reflect historical inequalities, and widely used algorithms have been shown to disadvantage precisely those groups that health systems are meant to protect (Obermeyer et al., 2019; Seyyed-Kalantari et al., 2021; Wong et al., 2021).

The economic-historical lesson is that transformative technologies deliver broad gains when complementary investments, institutions and distributional commitments mature alongside them (Brynjolfsson et al., 2021; David, 1990; Fogel, 2004). For personalized and equitable health systems, this means investing in representative data, independent validation, privacy-preserving collaboration, economic evaluation and accountable governance. If these conditions are met, predictive ML can become an instrument of both individual precision and collective fairness; if not, it risks becoming another technology whose benefits were captured unequally. The choice is institutional and political as much as technical.

References

Acosta, J. N., Falcone, G. J., Rajpurkar, P., & Topol, E. J. (2022). Multimodal biomedical AI. *Nature Medicine, 28*(9), 1773–1784.

Baldwin, P. (1990). *The politics of social solidarity: Class bases of the European welfare state, 1875–1975*. Cambridge University Press.

Baumol, W. J. (2012). *The cost disease: Why computers get cheaper and health care doesn't*. Yale University Press.

Beam, A. L., & Kohane, I. S. (2018). Big data and machine learning in health care. *JAMA, 319*(13), 1317–1318.
Batar, S. CYBERCRIME AND LAW ENFORCEMENT: CHALLENGES AND SOLUTIONS. Law & Order Administration, 61.
Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics, 13*(1), 333–372.
Batar, S., Jha, R. S., Tyagi, N., Jain, V., & Chaudhary, A. (2020). Protection of Software-Patent or Copyright?. Our Heritage, 68(1), 134-143.
Batar, S. PATENTING BIOTECHNOLOGICAL INVENTIONS: A CRITICAL ANALYSIS. INTELLECTUAL PROPERTY RIGHTS, 46.
Chen, I. Y., Joshi, S., & Ghassemi, M. (2020). Treating health disparities with artificial intelligence. *Nature Medicine, 26*(1), 16–17.
Dubey, B., & Kant, N. (2024). Revitalizing Indigenous Justice: Mediation in India's Tribal Justice System. International Journal of Research in Management & Social Science, 12(1), 24-32.
Collins, F. S., & Varmus, H. (2015). A new initiative on precision medicine. *New England Journal of Medicine, 372*(9), 793–795.
Cutler, D. M., & McClellan, M. (2001). Is technological change in medicine worth it? *Health Affairs, 20*(5), 11–29.
David, P. A. (1990). The dynamo and the computer: An historical perspective on the modern productivity paradox. *American Economic Review, 80*(2), 355–361.
Deaton, A. (2013). *The great escape: Health, wealth, and the origins of inequality*. Princeton University Press.
Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature, 542*(7639), 115–118.
Fogel, R. W. (2004). *The escape from hunger and premature death, 1700–2100: Europe, America, and the Third World*. Cambridge University Press.
Ghassemi, M., Oakden-Rayner, L., & Beam, A. L. (2021). The false hope of current approaches to explainable artificial intelligence in health care. *The Lancet Digital Health, 3*(11), e745–e750.
Gulshan, V., Peng, L., Coram, M., Stumpe, M. C., Wu, D., Narayanaswamy, A., … Webster, D. R. (2016). Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA, 316*(22), 2402–2410.
Kaissis, G. A., Makowski, M. R., Rückert, D., & Braren, R. F. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence, 2*(6), 305–311.
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.
Marmot, M. (2005). Social determinants of health inequalities. *The Lancet, 365*(9464), 1099–1104.
Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys, 54*(6), Article 115.
Muehlematter, U. J., Daniore, P., & Vokinger, K. N. (2021). Approval of artificial intelligence and machine learning-based medical devices in the USA and Europe (2015–20): A comparative analysis. *The Lancet Digital Health, 3*(3), e195–e203.
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.
Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. *New England Journal of Medicine, 380*(14), 1347–1358.

Rajkomar, A., Hardt, M., Howell, M. D., Corrado, G., & Chin, M. H. (2018). Ensuring fairness in machine learning to advance health equity. *Annals of Internal Medicine, 169*(12), 866–872.
Rajkomar, A., Oren, E., Chen, K., Dai, A. M., Hajaj, N., Hardt, M., … Dean, J. (2018). Scalable and accurate deep learning with electronic health records. *npj Digital Medicine, 1*, Article 18.
Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L, 2024/1689.
Regulation (EU) 2025/327 of the European Parliament and of the Council of 11 February 2025 on the European Health Data Space. *Official Journal of the European Union*, L, 2025/327.
Rieke, N., Hancox, J., Li, W., Milletarì, F., Roth, H. R., Albarqouni, S., … Cardoso, M. J. (2020). The future of digital health with federated learning. *npj Digital Medicine, 3*, Article 119.
Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence, 1*(5), 206–215.
Seyyed-Kalantari, L., Zhang, H., McDermott, M. B. A., Chen, I. Y., & Ghassemi, M. (2021). Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. *Nature Medicine, 27*(12), 2176–2182.
Singhal, K., Azizi, S., Tu, T., Mahdavi, S. S., Wei, J., Chung, H. W., … Natarajan, V. (2023). Large language models encode clinical knowledge. *Nature, 620*(7972), 172–180.
Sirugo, G., Williams, S. M., & Tishkoff, S. A. (2019). The missing diversity in human genetic studies. *Cell, 177*(1), 26–31.
Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., … Collins, R. (2015). UK Biobank: An open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS Medicine, 12*(3), e1001779.
Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine, 25*(1), 44–56.
Vayena, E., Blasimme, A., & Cohen, I. G. (2018). Machine learning in medicine: Addressing ethical challenges. *PLoS Medicine, 15*(11), e1002689.
Vokinger, K. N., Feuerriegel, S., & Kesselheim, A. S. (2021). Mitigating bias in machine learning for medicine. *Communications Medicine, 1*, Article 25.
Wiens, J., Saria, S., Sendak, M., Ghassemi, M., Liu, V. X., Doshi-Velez, F., … Goldenberg, A. (2019). Do no harm: A roadmap for responsible machine learning for health care. *Nature Medicine, 25*(9), 1337–1340.
Wong, A., Otles, E., Donnelly, J. P., Krumm, A., McCullough, J., DeTroyer-Cooper, O., … Singh, K. (2021). External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. *JAMA Internal Medicine, 181*(8), 1065–1070.
World Health Organization. (2021). *Ethics and governance of artificial intelligence for health: WHO guidance*. WHO.

Declarations

Conflict of interest: The authors declare that they have no conflict of interest.

Funding: This review received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

Ethical approval: Not applicable. This article is a review of published literature and did not involve human participants or animal subjects.

Data availability: No new data were generated or analyzed in support of this review. All sources relied upon are cited in the References.

Use of AI tools: No generative artificial intelligence tool was used in the conception, drafting or analysis of this article. The authors take full responsibility for the content.

Author contributions: Author contributed to the conception of the review, the survey of the literature, and the drafting and revision of the manuscript. Author read and approved the final version.