This open-access article is distributed under Creative Commons licence CC-BY-NC 4.0.

Data science and AI in medicine and global health: The need for inter-philosophies dialogue, cross-cultural ethics and ecocentricity

N S Munung,[1] MSc, PhD; **G B Tangwa**,[2,3] MA, PhD

[1] *Division of Human Genetics, Faculty of Health Sciences, University of Cape Town, South Africa*
[2] *Department of Philosophy, University of Yaoundé 1, Yaoundé, Cameroon*
[3] *Cameroon Bioethics Initiative (CAMBIN), Yaoundé, Cameroon*

***Corresponding author**: N S Munung (nchangwisyntia@yahoo.com)*

Advances in data science and medical artificial intelligence (AI) raise complex philosophical and ethical quandaries about what it means to know a person or a community through data and what kinds of people and societies we are becoming in this era of predictive data science. Drawing on four lightly fictional but reality-informed case studies in mental health, radiology, genomics and environmental public health, we reflect on how AI technologies, largely built on Western biomedical traditions, may conflict with relational, spiritual and Indigenous understandings of health and wellbeing. This may manifest as epistemic friction, algorithmic fatalism and diminished trust in patient–clinician relationships. Besides familiar concerns regarding bias and transparency, this paper advances the discourse on the ethics of medical AI and data science in healthcare by shifting the analysis from epistemology (how AI systems know, classify and predict) to ontology (the study of the nature of being, as reconfigured by data and AI). We argue that AI systems may inflict significant ontological harm by reconfiguring identity, moral agency and imagined futures and advocate for a renewed medical humanism driven by inter-philosophies dialogue, cross-cultural ethics and ecocentric approaches to care. From Bamenda, Cameroon, to Mthatha, South Africa and Toronto, Canada, the future of medical AI must be defined by the moral and philosophical traditions that people already live by. African, Indigenous, Islamic, Buddhist, Confucian and marginalised Western worldviews should not be treated as peripheral critiques, but as constitutive resources for building inclusive, human-centred health technologies. We conclude that bioethics should be recognised as a core infrastructure in global health, on equal footing with data science, medicine, biomedical research and health innovation.

Keywords: Inter-philosophies dialogue, cross-cultural ethics, ecocentrism, artificial intelligence, bioethics, data science, medicine, medical humanism

S Afr Med J 2026;15(6b):e3662. https://doi.org/10.7196/SAMJ.2026.v115i6b.3662

As data science and medical artificial intelligence (AI) tools are deployed from technology hubs into clinics, schools, universities, refugee camps and community health settings, they carry not only code and technical capabilities but also embedded ontological assumptions about health, wellbeing and personhood. These tools enter healthcare and social spaces as non-neutral artefacts. Most also embody preconceptions shaped by Western traditions of risk quantification, biomedical abstraction and individualism, which influence how health and wellbeing are conceived, defined, interpreted and governed.

The rendering of granular personal, biological and environmental data into machine-readable form, known as 'datafication' and 'datafied embodiment',[1,2] raises philosophical and ethical questions about which forms of knowledge are recognised, privileged or marginalised within algorithmic systems. In settings where structural determinants of health intertwine with communal relations, lands, spirituality, cultural traditions, environmental conditions and human–non-human interactions, reducing complex experiences to discrete data points may result in epistemically partial representations and tensions with local conceptions of being, health and wellbeing.[3,4]

Across Africa, medical AI adoption is relatively nascent, however, investment, innovation and implementation efforts are accelerating rapidly.[5] Some pilots include AI-assisted diagnosis of diabetic retinopathy in Zambia,[6] rapid automated tuberculosis screening in hard-to-reach regions in Uganda, and mental health chatbots for students and adolescents in South Africa (SA) and Cameroon.[7] Although operating at small-to-moderate scale, their uptake signals broader shifts in diagnostics, clinical practice, biomedical research, public health and the social–moral–relational dimensions of health and healthcare. While often positioned in terms of efficiency, early detection and expanded access to care, these developments also raise questions about how persons and communities are known, categorised and acted upon through data-driven systems.

When AI is embedded within clinical workflows, public health surveillance and disease prediction, it inevitably interacts with local moral frameworks, cultures and everyday practices. However, the design of many such tools is informed by Western biomedical models that prioritise individualised risk assessment, statistical inference and predictive optimisation, even when they are deployed in settings where health, illness and wellbeing are understood relationally, spiritually or ecologically. This creates the risk of epistemic disjuncture between algorithmic outputs and local knowledges of being, and raises a fundamental philosophical question: whose knowledge counts in medical AI design, deployment and governance?

What does it mean to 'know' a person or a community through predictive scores, digital biomarkers and probabilistic inference? How

can agency, trust and informed consent be preserved in healthcare when AI reasoning remains partially opaque, even to clinicians? If AI and data science predictions now define clinical and life decisions, what are the implications for moral agency and the freedom to become otherwise?[8-9] These are no longer abstract questions, but ones arising daily across clinical, public health and research settings, as well as the intimate spaces of homes, communities and spiritual traditions.

Rather than treating ethical and epistemic concerns as secondary to the technical performance of medical AI, we must foreground its ontological and normative assumptions and examine how they align – or fail to align – with diverse moral worldviews of health, wellness, personhood and the meaning of life. Such analysis demands philosophical and bioethical engagement that is dialogical, pluralistic and cross-cultural.[10-12]

As a methodological device for ethical analysis, we used four fictional but reality-informed case studies in mental health, radiology, genomics and environmental public health to reflect on how data science and medical AI may generate ethical tensions around epistemic injustice, trust, agency and responsibility. Rather than representing specific individuals or institutions, they synthesised patterns in AI ethics discourse. Fictionalisation allowed us to illuminate moral assumptions without having to address issues of confidentiality or overclaim empirical generalisability. This approach is well established in bioethics and moral philosophy.[13,14] Moreover, myth, metaphor and proverb are central to the traditional African epistemological systems that strongly inform our work. The cases converged on one concern: what kinds of people and societies are we becoming in the age of predictive data science? In response, we call for a renewed medical humanism that prioritises inter-philosophies dialogue, cross-cultural ethics and ecocentric perspectives.[10-12,15] This approach makes room for African, Indigenous, Islamic, Buddhist, Confucian and other moral traditions to be treated as primary contributors to data science ethics, rather than as peripheral critiques or post hoc considerations.

While contemporary AI ethics discourse is mainly preoccupied with epistemological concerns such as accuracy, bias, transparency, explainability and algorithmic reliability, this paper advances the discourse by foregrounding the ontological stakes of AI in healthcare. Our normative contribution is threefold. First, we show that the ethical challenges of AI in medicine are fundamentally philosophical, spanning ontology, epistemology and moral agency. Second, we argue that medical humanism, enriched by inter-philosophies dialogue and cross-cultural ethics, is necessary for AI governance in global health. Third, we advance that ecocentric ethics is necessary for addressing the relational and environmental dimensions of health that AI systems routinely overlook.

We ask not how AI systems know, classify and predict, but how such systems reconfigure what it means to be a person, a patient or a community. By shifting our attention from epistemology to ontology, we offer a complementary normative lens for AI ethics. We conclude by arguing that bioethics should be treated as a key infrastructure for emerging technologies, on par with data science, medicine, biomedical research and innovation.

Case studies

Case 1: AI-powered mental health assessments, epistemic friction and ontological harm

AI-driven chatbots are increasingly promoted as low-cost tools for mental health screening and support in schools, hospitals, refugee camps and settings with limited access to mental health professionals.[16,17] Marketed as non-judgemental and scalable, these systems employ natural language processing and sentiment analysis to classify emotional distress and recommend interventions or care pathways.[18]

In a pilot digital mental health programme in a secondary school in Bamenda, Cameroon, a UK-based start-up deployed *SweetClassMH,* an AI-powered screening chatbot, to support early identification of emotional distress among students. Students were encouraged to use *SweetClassMH* as part of routine mental health monitoring. During one interaction, the system classified Sirri, a 15-year-old student grieving the recent death of her twin sister, as 'high risk' for a mental health problem. It then recommended counselling and psychotherapy. Sirri interpreted this to mean she was 'going mad'. She became anxious, withdrew socially and later told the school counsellor that her family considered her present condition to be twin-soul bereavement, and that all she needed was to keep with her traditional symbol of her twin for spiritual stability. In her culture, twins are considered to be one soul in two bodies.

In this case, the algorithmic classification of mental health clashes with the moral and spiritual frameworks held by Sirri's family. It illustrates how a well-intentioned AI mental health application can generate epistemic friction when deployed in contexts where mental health or psychological suffering is not primarily understood through biomedical categories. The ethical concern lies in algorithmic authority displacing spiritual, relational or cultural understandings of emotional stability. Can emotional distress be meaningfully reduced to data patterns?

More fundamentally, *SweetClassMH* inflicted ontological harm, understood here as damage to a person's sense of self via imposed classificatory regimes from another worldview. For Sirri, this involved reconfiguring her identity from a grieving twin sister who has lost a part of her 'soul' to a person perceived as mentally 'abnormal'. The normative issue, therefore, is not whether the diagnosis from *SweetClassMH* was accurate, but whose conception of mental health informed its design and deployment. A growing body of literature suggests that conversational AI in mental health can misinterpret culturally situated expressions of distress, reinforce stigma or pathologise ordinary emotional responses when local worldviews are ignored.[19]

This case raises three interrelated ethical questions. First, whose definitions of wellbeing should guide AI mental health tools? Second, how might such systems accommodate plural and sometimes conflicting understandings of mental health and its causes without reducing one to another? Third, what forms of accountability are owed, and by whom (in this case, the school or the developers), for ontological harm, even in the absence of clinical error? Addressing these questions requires inter-philosophies dialogue and cross-cultural deliberation on what it means to be 'emotionally well'. These are not questions only for data scientists or psychologists, but require engagement with philosophers of mind, anthropologists, educators, spiritual leaders, bioethicists and ordinary non-experts.

Case 2: The black box of AI in radiology and the philosophical limits of knowing

AI systems are increasingly integrated into radiology to support the detection and classification of disease using medical imaging such as mammograms, magnetic resonance imaging (MRI) and computed tomography (CT) scans.[20] Promoted as tools that enhance accuracy and efficiency, these systems often rely on inscrutable machine-learning models that are inaccessible to clinicians and patients.

In a SA public hospital in Mthatha, Eastern Cape Province, a multinational health technology firm has deployed an AI-powered mammography system for breast cancer screening. The system was

trained on imaging data sets from women in Kenya, Nigeria, Tunisia and SA. Nomfundo, a 42-year-old teacher, underwent a mammogram that was initially interpreted by a radiologist as normal. The AI system, however, classified the image as 'high risk' and recommended a biopsy. When Nomfundo asked for an explanation, the clinician responded: 'It detected a pattern we do not fully understand, but it is based on data from women across Africa.' Nomfundo consented to the biopsy out of fear of delaying care rather than confidence in the decision. Although the biopsy confirmed early-stage breast cancer, the experience left her anxious and distrustful of the diagnostic process.

This encounter exemplifies the philosophical limits of knowing in AI-assisted care, and how opacity in AI-assisted diagnosis can undermine patient agency even when outcomes are clinically beneficial. The 'black box' logic displaces traditional trust in the patient–clinician relationship, namely empathy, dialogue and explanation of diagnosis and care pathways. What does it mean for clinical judgement when disease is identified through data patterns that are epistemically inaccessible to both patient and clinician? What impact does this have on patient agency and trust in clinicians when patients are asked to consent to interventions based on a system that neither party fully understands? The algorithm may have been correct, but it disrupted the healthcare experience.

This case also raises questions of accountability and professional responsibility. Had the biopsy revealed no cancer, it would have been difficult to assign responsibility for the physical risk and emotional stress to which Nomfundo was exposed. Conversely, had AI not been used and the cancer been missed, clinicians might have been faulted for not using available tools. AI thus introduces a terrain of diffused responsibility and blurs established norms of clinical accountability.

The normative issue here is that AI can disrupt the moral foundations of clinical care, which are grounded not only in acceptable outcomes but also in explanation, trust, dialogue and informed consent. Although advances in explainable AI (XAI) may partially address algorithmic opacity,[21,22] technical transparency alone cannot resolve the ethical deficit revealed in this case. Similarly, while data representation in AI training is ethically and scientifically important, representation in data sets does not equate to representation in epistemic values or philosophies of healthcare. AI systems may learn from African bodies while remaining silent on African conceptions of trust, relationality and moral responsibility in medicine. Epistemic justice in medical AI demands interpretability plus participatory engagement with local practices, expectations and ontologies of care. Without this, technically sound clinical decisions may estrange patients ontologically.

Case 3: Genomic prediction, AI and narratives of fate

AI-driven genomic technologies are increasingly offered through direct-to-consumer genetic testing platforms that use machine-learning models to forecast individual risk for conditions such as Alzheimer's disease, cancer, diabetes and selected behavioural traits.[23] Advertised as predictive, personalised and empowering, these tools promise to enable individuals to anticipate and manage their future health trajectories.

Danilo, a 28-year-old man of Afro-Brazilian heritage living in Toronto, Canada, ordered one such genomic testing kit out of curiosity rather than medical necessity. The AI-generated report assigned him a 70% lifetime risk of developing Alzheimer's disease. Although framed as probabilistic, the result was experienced as quasi-decisive. His parents interpreted the result through different epistemic frameworks. Danilo and his mother practise a diasporic spiritual tradition and understood the result as ancestral fate or karma, while his father, a clinician and scientist, urged immediate preventive intervention. Confronted with these competing interpretations, Danilo began to reorganise his life. He reconsidered fatherhood, revised his career plans and adopted new dietary and exercise regimens, not as expressions of empowerment but as strategies for controlling his future. What was presented as statistical risk became a normative directive.

This case illustrates how AI-enabled genomic prediction does more than merely inform future healthcare decisions. Here, three epistemologies intersect: technoscientific, spiritual and clinical, each offering distinct accounts of what it means to know, interpret and act upon the future. In moral traditions where the future is unknowable and embedded in spiritual narratives, algorithmic prediction risks imposing a deterministic script. The question, therefore, is: does knowing our statistical or probabilistic futures free or bind us? In this case, the ethical concern is algorithmic fatalism: the belief that probabilistic genomic information determines destiny and narrows the space for meaningful choice. When individuals begin to live as if a predicted outcome were inevitable, genomic risk prediction becomes a technology of identity formation.

This case raises urgent normative questions. Who should determine how individuals live in light of probabilistic genomic information? What responsibilities do developers, clinicians and regulators bear in defining how risk is framed, interpreted and communicated in AI-enabled genomics without foreclosing the freedom to become otherwise? To resist algorithmic fatalism, AI genomic tools must be embedded within frameworks that explicitly acknowledge the contingency of futures across different moral worlds. Inter-philosophies and cross-cultural dialogue are therefore required if predictive technologies are to support healthcare while preserving the freedom to become otherwise.

Case 4: Environmental public health surveillance, AI and social injustice

Environmental public health surveillance is seeing growing demand for AI in outbreak detection, early warning, trend prediction and response modelling.[24,25]

The Mpumalanga Highveld, a major industrial hub for coal mining and energy production in SA, is among the most polluted regions in the Southern Hemisphere. Decades of mining expansion have displaced families, destroyed sacred sites and desecrated ancestral graves. In addition, chronic exposure to toxic air ('black dust') emanating from the mines is associated with high rates of asthma, bronchitis and cardiovascular disease, and disproportionately affects Black and low-income communities.

In 2023, a collaboration between a SA university and a Silicon Valley AI lab launched *UmoyaScan*, an environmental public health surveillance system that integrates satellite imagery, meteorological data, emissions modelling and self-reported health symptoms to generate population-level risk profiles for pollution-related respiratory disease. The Department of Health heralded *UmoyaScan* as a milestone in 'precision public health' in Africa. Yet for residents of the Highveld, this promise has not translated into meaningful health gains. Doctors continue to report a high incidence of asthma, pneumonia and lung cancer, and local clinics are chronically under-resourced. Traditional healers, whose Indigenous knowledge systems have long tracked ecological disruption through land, wind patterns and animal behaviour, were not consulted in the design or implementation of *UmoyaScan*. Some community members interpret the high number of lung cancer cases as a spiritual effect following the desecration of ancestral graves. Allegations have also emerged that community-level health and environmental data are being monetised in international climate and health finance markets without community consent or benefit sharing. Although the SA

government has acknowledged concerns regarding the health impact of 'black dust' in Mpumalanga, data governance and environmental injustice, its political entanglements with the fossil fuel industry and national dependence on coal for energy have slowed substantive remedial action.

UmoyaScan exemplifies how AI-enabled public health surveillance can reinforce structural, epistemic, environmental and ontological injustices rather than redress them. The system privileges algorithmic expertise and the generation of data for decision-making while neglecting local knowledge systems that guide the community responses to ecological harm. Epistemically, *UmoyaScan* 'sees' pollution, particulate matter and pulmonary risk, but not moral injury and its association with community perceptions of prevalent respiratory conditions in the region.

Can satellite data capture the slow violence of industrial displacement? Who decides what counts as legitimate environmental health knowledge? Do communities have the moral right to refuse being mapped, predicted or monetised in the name of public health or climate action? What *UmoyaScan* ultimately reveals is not a deficit of data, but a deficit of distributive, epistemic, ecological and spiritual justice. The people of the Mpumalanga Highveld are not just breathing polluted air; they are breathing the cumulative harm of land dispossession, fractured kinship and ancestral disconnection. Without attention to these dimensions, AI-driven environmental public health surveillance risks becoming extractive rather than protective, and repeating colonial patterns of observation without accountability.

A relational and ecocentric approach to AI governance in public health surveillance is needed, one that centres community agency and interrogates who defines the goals, ownership and benefits of environmental AI systems. Surveillance technologies should not treat communities as data mines or land as a neutral substrate. Their ethical legitimacy demands reciprocity and respect for Indigenous and local knowledges,[15] alongside material commitments to health system strengthening and environmental remediation.

Toward a medical humanism of inter-philosophies dialogue, cross-cultural ethics and ecocentricity

As AI gains influence in defining our medical futures, we must ask which philosophies of health, moral agency and knowing are being encoded, intentionally or unintentionally, into these systems. Across the four cases presented, a consistent ethical pattern emerges: AI systems in medicine and public health do not merely assist clinical decision-making; they co-constitute how health, illness, personhood and responsibility are understood. In Bamenda, algorithmic mental health classification transformed grief into pathology, generating fear and social withdrawal. In Mthatha, opaque algorithmic authority disrupted trust and agency despite diagnostic success. In Toronto, probabilistic genomic prediction narrowed imagined futures and reoriented life choices toward anticipatory self-surveillance. In Mpumalanga, environmental surveillance rendered communities legible as data sources while neglecting spiritual harm to land and culture.

Together, these cases demonstrate that the ethical challenges of AI extend beyond epistemic failure to ontological harm. They concern philosophical questions about what it means to know, care for and live well as human beings situated within plural social, cultural and ecological worlds. Recognising this clarifies why technical solutions such as bias mitigation or XAI, while necessary, are insufficient on their own. Normative engagement with plural conceptions of health, personhood, responsibility and healthcare is required, confirming that medical humanism, inter-philosophies dialogue and ecocentric ethics are indispensable.

Box 1: Future Research Directions

- Empirical studies should examine how differing moral ontologies shape AI problem framing, risk interpretation and acceptable trade-offs in system design.
- Empirical and conceptual work should investigate the ontological effects of AI on identity, agency and future-making.
- Future work should develop environmental justice and intergenerational indicators for medical AI, drawing on Indigenous and eco-bio-communitarian frameworks.
- Comparative studies could examine how AI systems are interpreted, resisted or re-purposed in different sociocultural contexts.

What is needed is a medical humanism that situates AI within a pluralistic understanding of human flourishing, rather than reducing care to optimisation and prediction. This is not a rejection of technology, but a call to maintain fundamental moral practices in medicine concerned with meaning, trustworthy patient–clinician relationships, dignity and relational responsibility. A medical humanism adequate to the age of AI must therefore attend not only to outcomes, but also to how technologies and data science reconfigure relationships, identities and agency. Central to this humanistic reorientation are cross-cultural ethics, inter-philosophies dialogue and ecocentric perspectives.[10,11,26] This does not mean collapsing diverse traditions into a single framework; rather, it involves creating bioethics spaces where tensions, complementarities and disagreements in medical AI can be openly discussed and negotiated. Such spaces require epistemic humility from all stakeholders, as no single culture or discipline holds all the answers. They should also allow room for moral imagination – the ability to envision futures that are just, caring and inclusive.

In the age of AI, merely improving algorithms is not enough. We need to listen and act across disciplines, traditions and cultures. Encouragingly, emerging work in community-responsive AI and participatory digital health governance demonstrates that AI systems can, under specific conditions, support local public health initiatives rather than displace them.[27]

Conclusion

Philosophy and bioethics as infrastructures in global health

This paper shifts contemporary discourse on AI ethics from predominantly epistemological concerns to the ontological effects on persons and communities. AI can generate significant ontological harm even when technically accurate. In this sense, the ethical challenge of AI in medicine is not only about how machines know, but how societies choose to define, value and care for human and non-human life in an age of predictive and surveillance technologies. Our proposal is to position AI within relational, plural and ecocentric visions of health, where humans are understood not merely as data sets or risk profiles but as meaning-making beings embedded in social, spiritual and ecological worlds. In this sense, philosophy and bioethics are not external critiques of medical AI; they are constitutive elements of its design, governance and moral imagination.

As AI and data science become central and public to global health initiatives, a renewed social contract for data-driven medicine is required,[28] one that asks not only what AI can do, but what it ought to do, for whom, at what moral cost and to whose benefit.

Going forward, what is needed is not only accurate or representative AI, but a medical humanism that recognises philosophy and bioethics

as core infrastructures in global health, on equal footing with data science centres, clinical studies and biotechnology. As Benatar *et al.*[10-12] have long emphasised, building technical and digital infrastructures for biotechnology without parallel investment in the capacity for cross-cultural ethics and inter-philosophies dialogue risks reproducing global inequities rather than alleviating them. Ethical AI and data science in medicine therefore require investment in bioethics laboratories to support conversations, research (see Box 1) and collaborations between AI developers, clinicians, patients, communities, policymakers, scientists and ordinary non-experts. Without such bioethics infrastructures, AI risks advancing as an extractive and epistemically narrow enterprise, one that draws on personal and environmental data while ignoring local meanings, values, injustice and histories.

Data availability statement.

Declaration. The content of this article is solely the responsibility of the authors and does not necessarily represent the official views of the US National Institutes of Health.

Acknowledgements. This article is written in honour of Prof. Solly Benatar and seeks to build upon his pioneering work in inter-philosophies dialogue and cross-cultural ethics in medicine and biotechnology. We are grateful to the organisers of the festschrift for the opportunity to contribute to this collection.

Author contributions. All authors contributed equally to the manuscript.

Funding. NSM receives research funding from the National Institute of Health (NIH) and the National Institute of Mental Health (U01MH127692) for the project Public Understanding of Big Data in Genomics Medicine in Africa (PUBGEM-Africa) under the Harnessing Data Science for Health Discovery and Innovation in Africa (DSI-Africa) programme, as well as The National Heart, Lung and Blood Institute (U24HL135600) as part of the Sickle Africa Data Coordinating Center (SADaCC).

Conflicts of interest. None.

1. Southerton C. Datafication. In: Schintler LA, McNeely CL, eds. Encyclopedia of Big Data. Cham: Cham, Switzerland: Springer International Publishing, 2022:358-361.
2. Lupton D, Clark M, Southerton C. Digitized and datafied embodiment: A more-than-human approach. In: Herbrechter S, Callus I, Rossini M, Grech M, de Bruin-Molé M, John Müller C, eds. Palgrave Handbook of Critical Posthumanism. Cham, Switzerland: Springer International Publishing, 2022:361-383.
3. Zhang J, Zhang Z-M. Ethics and governance of trustworthy medical artificial intelligence. BMC Med Inform Decis Mak 2023;23(1):7. https://doi.org/10.1186/s12911-023-02103-9
4. Canales C, Lee C, Cannesson M. Science without conscience is but the ruin of the soul: The ethics of big data and artificial intelligence in perioperative medicine. Anesth Analg 2020;130(5):1234-1243. https://doi.org/10.1213/ANE.0000000000004728
5. Owoyemi A, Owoyemi J, Osiyemi A, Boyd A. Artificial intelligence for healthcare in Africa. Front Digit Health 2020;2:6. https://doi.org/10.3389/fdgth.2020.00006
6. Bellemo V, Lim ZW, Lim G, et al. Artificial intelligence using deep learning to screen for referable and vision-threatening diabetic retinopathy in Africa: A clinical validation study. Lancet Digit Health 2019;1(1):35-44. https://doi.org/10.1016/s2589-7500(19)30004-4
7. Madiba M. Transforming mental health in South African universities: The promise of AI chatbots. Times Higher Education, 12 June 2024. https://www.timeshighereducation.com/campus/transforming-mental-health-south-african-universities-promise-ai-chatbots (accessed 15 May 2025).
8. Sætra HS. Science as a vocation in the era of big data: The philosophy of science behind big data and humanity's continued part in science. Integr Psychol Behav Sci 2018;52(4):508-522. https://doi.org/10.1007/s12124-018-9447-5
9. Ostherr K. Artificial intelligence and medical humanities. J Med Humanit 2022;43(2):211-232. https://doi.org/10.1007/s10912-020-09636-4
10. Benatar S, Daibes I, Tomsons S. Inter-philosophies dialogue: Creating a paradigm for global health ethics. Kennedy Inst Ethics J 2016;26(3):323-346. https://doi.org/10.1353/ken.2016.0027
11. Benatar SR, Singer PA, Viens AM, eds. Global health ethics and cross-cultural considerations in bioethics. Cambridge, UK: Cambridge University Press, 2009:341-349. https://doi.org/10.1017/CBO9780511545566.051
12. Benatar SR, Daar AS, Singer PA. Global health ethics: The rationale for mutual caring. Int Aff 2003;79(1):107-138. https://doi.org/10.1111/1468-2346.00298
13. Chan S. More than cautionary tales: The role of fiction in bioethics. J Med Ethics 2009;35(7):398-389. https://doi.org/10.1136/jme.2009.031252
14. Moen OM. Why good work in philosophical bioethics often looks strange. Theor Med Bioeth 2022;44(2):153-164. https://doi.org/10.1007/s11017-022-09601-3
15. Munung NS, Tangwa GB. Eco-bio-communitarianism: Reimagining bioethics for sustainable health. BMC Glob Public Health 2025;3:13. https://doi.org/10.1186/s44263-025-00129-x
16. Habicht J, Viswanathan S, Carrington B, Hauser TU, Harper R, Rollwage M. Closing the accessibility gap to mental health treatment with a personalized self-referral chatbot. Nat Med 2024;30(2):595-602. https://doi.org/10.1038/s41591-023-02766-x
17. Olawade DB, Wada OZ, Odetayo A, David-Olawade AC, Asaolu F, Eberhardt J. Enhancing mental health with artificial intelligence: Current trends and future prospects. J Med Surg Public Health 2024;3:100099. https://doi.org/10.1016/j.glmedi.2024.100099
18. Abd-Alrazaq AA, Alajlani M, Alalwan AA, Bewick BM, Gardner P, Househ M. An overview of the features of chatbots in mental health: A scoping review. Int J Med Inform 2019;132:103978. https://doi.org/10.1016/j.ijmedinf.2019.103978
19. Rahsepar Meadi M, Sillekens T, Metselaar S, van Balkom A, Bernstein J, Batelaan N. Exploring the ethical challenges of conversational AI in mental health care: Scoping review. JMIR Ment Health 2025;12:e60432. https://doi.org/10.2196/60432
20. Khalifa M, Albadawy M. AI in diagnostic imaging: Revolutionising accuracy and efficiency. Comput Methods Programs Biomed Update 2024;5:100146. https://doi.org/10.1016/j.cmpbup.2024.100146
21. Alkhanbouli R, Almadhaani HMA, Alhosani F, Simsekler MCE. The role of explainable artificial intelligence in disease prediction: A systematic literature review and future research directions. BMC Med Inform Decis Mak 2025;25(1):110. https://doi.org/10.1186/s12911-025-02944-6
22. Sadeghi Z, Alizadehsani R, Cifci MA, et al. A review of explainable artificial intelligence in healthcare. Comput Electr Eng 2024;118:109370. https://doi.org/10.1016/j.compeleceng.2024.109370
23. O'Daniel JM, Kobelka C, Foss K, Foreman AKM, Milko LV. Direct-to-consumer genetic testing. NEJM Evid 2025;4(11):EVIDra2400455. https://doi.org/10.1056/evidra2400455
24. Zeng D, Cao Z, Neill DB. Artificial intelligence-enabled public health surveillance – from local detection to global epidemic monitoring and control. In: Xing L, Giger ML, Min JK, eds. Artificial Intelligence in Medicine. Cambridge, USA: Academic Press, 2020:437-453. https://doi.org/10.1016/B978-0-12-821259-2.00022-3
25. Iyer HS, Karasaki S, Yi L, Hswen Y, James P, VoPham T. Harnessing geospatial artificial intelligence (GeoAI) for environmental epidemiology: A narrative review. Curr Environ Health Rep 2025;12(1):34. https://doi.org/10.1007/s40572-025-00497-4
26. Benatar SR. Global health and justice: Re-examining our values. Bioethics 2013;27(6):297-304. https://doi.org/10.1111/bioe.12033
27. Hsu Y-C, Huang T-HK, Verma H, Mauri A, Nourbakhsh I, Bozzon A. Empowering local communities using artificial intelligence. Patterns 2022;3(3):100449. https://doi.org/10.1016/j.patter.2022.100449
28. Munung NS, Ewuoso C, Agamah FE, Chimusa ER. Can a bold 'social contract' make data sharing more palatable? Nature 2025;645(8080):309-312. https://doi.org/10.1038/d41586-025-02817-y

Submitted 20 May 2025. Accepted 17 December 2025.