

The ethics and law of medical AI in South Africa: Balancing innovation with responsibility

M Ngcobo, LLB, LLM, LLD Cand

Louis D Brandeis School of Law, University of Louisville, Kentucky, USA

Corresponding author: M Ngcobo (mnotho.ngcobo@louisville.edu; mnothongcobo@gmail.com)

The rapid integration of artificial intelligence (AI) into medical practice presents both transformative opportunities and profound ethical and legal challenges. In South Africa, a country with a dual healthcare system and significant disparities in access to medical services, AI holds the promise of revolutionising healthcare delivery by enhancing diagnostic accuracy, improving patient outcomes, and mitigating resource constraints. However, the deployment of medical AI also raises critical ethical concerns regarding patient autonomy, informed consent, data protection, and accountability. From a legal standpoint, South Africa must navigate a complex regulatory terrain to ensure that AI aligns with constitutional rights and statutory obligations while fostering innovation. This article explores the legal and ethical dimensions of medical AI in South Africa, arguing for a balanced approach that encourages technological advancement without compromising fundamental principles of medical ethics and patient rights.

Keywords: medical AI, health law, informed consent, AI regulation, access to health

S Afr Med J 2025;115(5b):e3667. <https://doi.org/10.7196/SAMJ.2025.v115i5b.3667>

Medical artificial intelligence (AI) is already reshaping the landscape of healthcare in South Africa, particularly in diagnostic imaging, predictive analytics, robotic surgery and telemedicine. AI-driven tools have demonstrated remarkable efficacy in detecting diseases such as tuberculosis, cervical cancer and diabetic retinopathy—conditions that disproportionately affect South Africa's population.^[1] For example, AI-powered diagnostic platforms have been integrated into tuberculosis screening programmes, enabling quicker and more accurate detection, which is crucial in rural clinics where radiologists are scarce.^[2] Additionally, AI-driven pathology tools are being piloted in oncology to improve early detection rates of cancers, reducing diagnostic delays and improving survival rates.^[3]

Moreover, AI-powered chatbots and virtual assistants are enhancing healthcare accessibility by providing preliminary diagnoses and treatment recommendations, particularly in rural areas where medical practitioners are scarce.^[4] Companies such as HealthConnect and AI4Medical have developed chatbot-driven triage systems that direct patients to appropriate levels of care, reducing strain on emergency services. The South African government has also expressed interest in leveraging AI for epidemiological surveillance and pandemic preparedness, as seen in AI-driven modeling during the COVID-19 crisis, which helped predict outbreak trends and optimise resource allocation.^[5]

From an economic perspective, medical AI has the potential to alleviate the burden on South Africa's overstretched public health sector. AI applications can assist overburdened doctors by automating administrative tasks, optimising hospital workflows, reducing diagnostic errors, and enhancing decision-making efficiency.^[6] For instance, AI-driven scheduling systems in state hospitals have helped streamline patient appointments, minimising wait times and maximising the efficiency of available healthcare personnel.^[7] Furthermore, AI-driven solutions may democratise access to specialised medical expertise, enabling patients in under-resourced areas to benefit from high-quality care that would otherwise be unavailable. Telemedicine platforms integrated with

AI diagnostics have expanded access to specialists, allowing patients in remote areas to receive expert consultations without travelling long distances.^[8] These advancements demonstrate AI's potential not only to enhance healthcare delivery but also to create a more equitable system in South Africa, bridging the gap between urban and rural healthcare services.

Despite significant advances in diagnostic accuracy and healthcare accessibility afforded by artificial intelligence, real-world deployments frequently expose critical blind spots rooted in biased training data and insufficient performance monitoring.^[9] Recent reviews of dermatology AI programmes reveal that fewer than one-third publish performance metrics for darker skin types, and the overwhelming majority of algorithms remain trained on lighter skin tones.^[9] Such under-representation engenders algorithmic bias, leaving patients with Fitzpatrick skin types V and VI particularly vulnerable to misdiagnosis or delayed intervention.^[9,10] In the South African context, where a substantial proportion of the population presents with darker skin, these disparities are not merely theoretical; they risk exacerbating existing inequities in health outcomes and undermining AI's promise to democratise quality care across urban and rural settings.^[11]

A striking illustration of this risk emerges from studies of melanoma classification algorithms, which report a decline in discriminatory power when applied to darker skin.^[12] Whereas the same model may achieve robust accuracy among lighter-skinned cohorts, its performance can drop by nearly fifteen percentage points for Fitzpatrick types V - VI. Transposed to a Cape Town tertiary dermatology clinic, where over forty percent of patients exhibit these skin types, this performance gap suggests that up to one in three melanomas could be erroneously flagged as benign. Although formal audits within South African hospitals remain scarce, anecdotal accounts from clinicians point to near-misses identified during internal reviews. These experiences underscore the urgent need for mandatory post-market surveillance, transparent error-reporting mechanisms, and robust bias-mitigation strategies before broad clinical deployment.

To ensure proportional oversight, it is useful to distinguish between two categories of medical AI applications. Low-risk systems—such as administrative chatbots, appointment-scheduling assistants, and simple triage calculators—primarily warrant foundational transparency, data-protection safeguards, and routine performance audits. By contrast, high-risk systems—including autonomous diagnostic, prognostic or treatment-recommendation algorithms—carry direct patient-safety stakes and therefore require rigorous pre-market approval, ongoing bias and fairness audits, mandatory explainability reports, and robust post-market surveillance.^[13] This calibrated taxonomy aligns regulatory intensity with potential harm, laying the groundwork for a legal framework that can both protect patients and foster innovation.

Current legal landscape

Constitutional foundations

To appreciate the legal framework governing medical AI in South African healthcare, one must begin with the supreme law of the land, the Constitution of the Republic of South Africa, 1996.^[14] Section 27 guarantees that everyone has the right to have access to healthcare services, including reproductive healthcare, and mandates that the state take reasonable legislative and other measures, within its available resources, to achieve the progressive realisation of each of these rights.^[15] Far from being aspirational rhetoric, this provision imposes a justiciable duty on government, requiring that legislative and policy initiatives deliver tangible improvements in health services rather than mere policy statements.

The Constitutional Court's jurisprudence on section 27 has crystallised two interrelated principles: the obligation to progressively realise socio-economic rights within available resources, and the requirement to ensure a minimum core of basic health services for those in desperate need. In *Government of the Republic of South Africa v. Grootboom*, the Court emphasised that progressive realisation does not imply unconstrained discretion but rather entails a programme that is reasonable in both conception and implementation.^[16] The state's measures must be capable of adapting to changing circumstances and of extending benefits to the most vulnerable without undue delay. In that judgment, the Court explained that a mere policy commitment to housing would not suffice if emergency relief remained inaccessible to those who could not afford alternative accommodation.^[16] Likewise, in *Minister of Health v. Treatment Action Campaign*, the Court held that the state's duty under section 27 includes providing lifesaving antiretroviral therapy to prevent mother-to-child transmission of HIV.^[17] The Court rejected administrative delays and logistical excuses, insisting that the national rollout of treatment must commence without unreasonable delay, and that the state must allocate sufficient resources to ensure continuous access for all eligible persons.^[17]

Together, *Grootboom* and *TAC* establish that section 27 imposes a dual requirement: first, the state must adopt a coherent, properly funded programme to expand access to critical health services; second, it must deliver a core minimum of essential services immediately to those in dire need. These twin obligations form the constitutional foundation against which any emergent health technology must be evaluated. When AI demonstrates the capacity to enhance diagnostic accuracy, expedite treatment decisions, or optimise scarce resources, the state cannot treat such innovations as optional or as premium add-ons available only to those who can pay. Rather, in fulfilling its constitutional mandate, the government must assess whether and how AI tools can strengthen public health programmes in both urban and rural settings and must take steps to integrate effective systems into the public sector within its resource constraints.

Applying these principles to medical AI, one begins by recognising that AI-driven diagnostics and predictive analytics hold the potential to address long-standing inequalities in South Africa's healthcare system. For example, some public clinics often lack specialist radiologists, leading to diagnostic delays that exacerbate morbidity and mortality.^[18] If an AI-powered imaging tool can reliably identify tuberculosis or early-stage cancer in underserved communities, the state's failure to pilot, evaluate and ultimately deploy such a system in the public sector may amount to an unreasonable delay or an incomplete rollout of core health services. Under the minimum-core doctrine, eligibility for immediate relief extends to those whose conditions can be mitigated or prevented through proven interventions.^[19] Therefore, once robust evidence emerges that a particular AI application materially improves health outcomes for vulnerable populations, the state's obligation to ensure equitable access crystallises.

Moreover, the principle of progressive realisation requires that the state's AI adoption strategy be neither *ad hoc* nor sporadic. Rather, it must form part of an overarching health innovation plan that contemplates phased integration, ongoing training of personnel, infrastructure upgrades, and monitoring of outcomes. Just as the *TAC* litigation compelled the Department of Health to expand antiretroviral programmes nationwide, strategic litigation or parliamentary oversight could compel the government to adopt AI road maps that meet constitutional standards.^[17] Failure to articulate such a plan, or to fund and implement it where cost-effective, risks judicial intervention on the grounds that the state's measures are neither reasonable in conception nor in implementation.

Finally, the constitutional right to healthcare encompasses more than access; it also embodies principles of dignity and non-discrimination.^[20] AI tools must therefore be deployed and regulated in ways that respect the dignity of patients. Systems that operate opaquely, that deprive individuals of meaningful choices, or that perpetuate bias against historically marginalised groups, contravene the constitutional commitment to human dignity. If an AI application systematically under-detects illnesses in certain communities, or if it makes treatment recommendations without regard to cultural or linguistic differences, the state's endorsement or procurement of such technology could be challenged as inconsistent with section 27's guarantee of equitable service delivery. Accordingly, constitutional foundations demand not only that AI enhances clinical outcomes, but also that it does so in a manner aligned with South Africa's commitments to dignity, equality and non-discrimination.

Statutory regimes

Beyond constitutional mandates, South Africa's statutory landscape establishes the specific legal contours within which medical AI must operate. Three principal regimes are particularly salient: the National Health Act, 61 of 2003;^[21] the Protection of Personal Information Act, 4 of 2013;^[22] and the Medical Device Regulations administered by the South African Health Products Regulatory Authority.^[23] Each regime addresses distinct but interlocking aspects of healthcare provision, data governance and medical device oversight, and together they form the scaffolding for responsible AI integration.

The National Health Act serves as the legislative backbone for healthcare delivery, codifying patient rights, provider obligations, and the organisational structure of health services. Chapter Two of the Act enshrines the requirement of informed consent, stipulating that no healthcare service may be provided without a patient's full and informed consent, that consent must be voluntary, and that the patient must receive sufficient information regarding the benefits, risks and available alternatives.^[21] Although drafted before the advent

of AI-driven diagnostics, the Act's broad definition of 'health service' encompasses any procedure, investigation or treatment performed for diagnostic or therapeutic purposes. Consequently, AI-based clinical decision-support systems fall within the ambit of services requiring informed consent.^[21]

However, the Act's consent provisions do not grapple with the unique challenges presented by algorithmic decision-making. Traditional consent encounters involve explaining the anticipated risks and benefits of a surgical intervention or pharmaceutical regimen.^[24] By contrast, AI tools operate through complex, often non-intuitive, processes that may consider hundreds of variables, from genomic markers to patterns in social determinants of health.^[25] Ensuring that patients truly comprehend an AI recommendation and its underlying uncertainties requires methodologies beyond standard disclosure scripts. Yet the National Health Act does not prescribe a statutory duty to explain algorithmic mechanics or to provide model-specific transparency. It leaves open the question of how clinicians should translate opaque machine-learning outputs into information that patients can meaningfully understand when consenting to care.

Turning to data governance, the Protection of Personal Information Act regulates the collection, processing, storage and dissemination of personal information, including sensitive health data. POPIA's foundational principles require purpose specification, data minimisation and the prohibition of excessive processing.^[22] Data subjects must provide informed consent for processing, and any secondary use or cross-border transfer of personal information demands additional safeguards.^[22] These requirements are particularly pertinent for AI systems trained on large, aggregated datasets. Machine-learning algorithms thrive on continuous feeds of patient information, and iterative retraining may occur long after the initial data collection.^[26] POPIA does not explicitly distinguish between static databases and dynamic, evolving models, nor does it provide clear guidelines on how to secure ongoing consent for secondary uses.

Furthermore, while POPIA advocates for de-identification to protect privacy, it does not fully address the risk of re-identification through sophisticated linkage attacks or inference algorithms. Advanced AI tools may reconstitute personal identities from anonymised datasets by correlating disparate data points.^[27,28] In the absence of explicit statutory rules for differential privacy or federated learning protocols, POPIA's broad mandates on security safeguards and consent may prove insufficient to guard against emerging technical threats. Practitioners and regulators must therefore interpret POPIA's principles in light of evolving AI capabilities and consider sectoral guidance that mandates privacy-preserving technologies tailored to healthcare contexts.

Finally, software intended for diagnostic or therapeutic purposes is regulated as a medical device under SAHPRA's Medical Device Regulations, published in terms of the Medicines and Related Substances Act, 101 of 1965.^[23] These regulations adopt an internationally harmonised approach, classifying devices according to risk classification rules and requiring pre-market registration, performance evaluation, and post-market surveillance for higher-risk categories.^[23] AI-powered diagnostic algorithms and decision-support tools, particularly those that directly influence patient management, typically fall into Class C or Class D, which demand the most rigorous review. Manufacturers or importers must demonstrate compliance with technical standards, including quality management systems and clinical evidence of safety and performance.

Despite this comprehensive framework, ambiguities remain regarding AI's adaptive nature. Traditional medical devices are static in design; once cleared, their specifications remain constant

until a formal update triggers re-evaluation.^[29] By contrast, modern AI models often incorporate continuous learning mechanisms that refine performance as new data are ingested.^[30] The regulations do not explicitly address how frequently adaptive algorithms must undergo re-certification, nor do they clarify whether minor model updates or parameter adjustments require fresh submissions. Moreover, the process for reporting adverse events associated with algorithmic errors is borrowed from hardware and pharmaceutical analogues, potentially overlooking algorithm-specific failure modes such as concept drift or training data bias.

Identified legal gaps

Despite the robust constitutional and statutory frameworks that govern South African healthcare, the unique characteristics of AI reveal significant lacunae in our existing laws. As AI systems become increasingly integrated into clinical workflows—performing tasks ranging from image interpretation to risk prediction and treatment recommendation—regulators, legislators and judges confront novel legal questions that cannot be fully resolved within the parameters of statutes drafted prior to the AI revolution. This part identifies four core legal gaps that must be addressed to ensure responsible deployment of medical AI: the opacity of algorithmic decision-making and its implications for informed consent; data privacy in the context of continuously learning systems; the attribution of liability among clinicians, institutions and developers; and the risk of exacerbating inequities through two-tier deployment. Each subsection analyses the shortcomings of existing provisions and lays the groundwork for targeted reforms.

Opacity and informed consent

At the heart of the legal challenges posed by medical AI lies the 'black box' problem—the opacity inherent in many machine-learning algorithms that derive conclusions through complex, non-linear transformations of high-dimensional data.^[31] Traditional medical interventions allow clinicians to explain, in plain language, the causal mechanisms by which a drug, a surgical procedure or a diagnostic test produces its effects. By contrast, an AI model may weigh hundreds or thousands of variables—genetic markers, imaging features, demographic data, social determinants—in ways that even its creators cannot fully disentangle.^[31] This opacity undermines the fundamental precondition of informed consent under the National Health Act, which requires that patients receive sufficient information regarding the nature, risks and benefits of a proposed health service to make a voluntary and knowledgeable decision.

The 'black box' problem in clinical context

Machine-learning models such as deep neural networks and ensemble methods routinely achieve high predictive accuracy by discovering latent patterns in training datasets.^[32] Yet these patterns lack intuitive explanations: a neural network's weights, for instance, do not correspond to discrete clinical variables in a manner that a physician can easily convey to a patient.^[33] In practice, when an AI tool flags a patient as 'high risk' for a cardiovascular event based on algorithmic analysis of imaging data, the physician may know that the risk threshold was set at, say, the 85th percentile of predicted probabilities, but cannot meaningfully translate to the patient which specific variables tipped the balance or how small changes in input data might alter the verdict. The patient, in turn, is told that 'the algorithm indicates increased risk,' but remains unable to probe the underlying causal inferences.

This situation is legally untenable under the Act's informed-consent regime. Section 7 of the National Health Act stipulates that every healthcare provider must ensure that the patient is made aware of the nature of the healthcare service, its risks and benefits, and any alternative forms of health services that may be appropriate.^[21] Where the 'nature' of the service is an algorithmic recommendation furnished by an inscrutable model, the requirement collapses into a rote recitation of algorithmic outputs rather than a substantive explanation of how the service operates. A patient consenting to AI-assisted diagnosis under these conditions cannot be said to have given fully informed consent, because the decision rests on premises that remain obscure, unchallengeable and unreviewable at the point of care.

Liability in AI-driven healthcare: A legal grey area

AI also complicates liability in medical malpractice claims, posing unresolved legal challenges under South African law. As previously discussed, traditional negligence principles require that a plaintiff demonstrate a duty of care, a breach of that duty, and causation between the breach and harm suffered. However, AI blurs the lines of responsibility: if a physician relies on an AI-generated diagnosis that later proves incorrect, who bears liability? Is it the doctor who is acting on the AI's recommendation? The hospital, for integrating AI-driven diagnostic tools? Or the developer for designing an AI system that failed to account for critical medical variables?

Internationally, courts have begun exploring alternative liability models, including:

- strict liability for AI developers, holding manufacturers and software creators accountable for defective AI systems^[34,35]
- hybrid liability models, where responsibility is distributed across healthcare providers, AI developers, and institutions
- duty-to-explain standards require physicians to exercise independent clinical judgment rather than defer uncritically to AI-generated recommendations.^[36]

South African courts have yet to rule on AI liability in medical malpractice, but legal precedents from European and US jurisdictions may serve as persuasive authority.^[37] Given the likelihood that South African courts will need to balance professional accountability with the need for innovation, statutory guidance on AI liability is urgently needed. Without clear legal standards, there is a risk that AI-driven errors may fall into a liability vacuum, leaving patients without recourse while absolving developers and institutions of responsibility.

AI and the right to healthcare: Public-private disparities

Beyond individual liability, AI raises broader constitutional questions about equitable access to healthcare services. As previously noted, South Africa's dual healthcare system, a well-funded private sector alongside an under-resourced public sector, raises concerns that AI-driven medical advancements may primarily benefit private hospitals, further entrenching healthcare inequalities.

Under Section 27 of the South African Constitution, the state has a duty to progressively realise access to healthcare services.^[20] If AI-powered diagnostics become widely available in private facilities but remain inaccessible in public hospitals, a legal argument could be made that this disparity violates the state's constitutional obligation to ensure equitable healthcare access. A particularly pressing concern is whether failing to provide AI-assisted diagnostics in public hospitals could constitute an unjustifiable limitation on the right to healthcare, especially if AI tools significantly improve treatment outcomes.

To mitigate these disparities, policymakers should implement technology transfer initiatives, ensuring that AI-driven tools are deployed across private and public hospitals. Additionally, public-sector AI funding programmes and legal protections against AI-induced healthcare inequities should be established to ensure that access to AI diagnostics is not dictated solely by socioeconomic status.

Strengthening AI regulation: A policy roadmap

South Africa's dual regulatory architecture—anchored in the Health Professions Council of South Africa and the South African Health Products Regulatory Authority—is uniquely suited to operationalise international AI-ethics norms at the national level. Both bodies possess explicit statutory mandates to safeguard patient safety, uphold professional standards and ensure the quality of health products and services. The HPCSA already enforces ethical conduct for practitioners and has instituted continuing professional development requirements, making it well-positioned to require algorithmic bias audits and explainability training under its existing licensing framework. Likewise, SAHPRA's pre-market approval processes for medical devices and *in vitro* diagnostics can be readily extended to cover high-risk AI algorithms, leveraging its expertise in clinical evaluation, post-market surveillance, and adverse-event reporting. By embedding the WHO's *Ethics and Governance of AI for Health* guidelines^[38] and the OECD's AI Principles^[39] into the HPCSA's professional codes and SAHPRA's conformity-assessment protocols, South Africa can achieve a seamless translation of global best practices into enforceable local rules. This approach avoids the creation of a redundant agency, capitalises on established institutional capacity, and ensures that both clinicians and manufacturers remain accountable to a consistent set of ethical and safety standards.

Possible regulatory measures include:

1. Mandatory bias audits – AI models should undergo periodic testing to detect and mitigate discriminatory outcomes in medical decision-making.
2. Licensing standards for AI-driven medical devices – AI systems should be subject to approval requirements that assess accuracy, fairness and transparency before clinical deployment.
3. Professional guidelines for AI in clinical practice – the HPCSA should expand its ethical codes to define physicians' responsibilities when using AI in patient care.
4. Explainability and transparency requirements – AI-driven medical tools should be designed with interpretability mechanisms, ensuring that both healthcare providers and patients can understand and challenge algorithmic decisions.

By embedding these safeguards into South Africa's legal framework, AI can serve as a tool for expanding healthcare access and quality rather than introducing new risks of inequity, discrimination and liability uncertainty.

Conclusion: AI as an ethical and legal imperative

The ethical and legal challenges surrounding AI in healthcare are not merely theoretical—they have direct implications for patient rights, clinical accountability, and equitable healthcare access. While AI offers unparalleled opportunities to address chronic healthcare challenges, from improving diagnostic accuracy to optimising resource allocation in hospitals, these benefits will remain elusive if they are not grounded in a robust legal and regulatory framework.

As South Africa moves toward wider AI adoption, it must strike a delicate balance: harness AI's transformative potential while upholding the core legal and ethical principles that underpin trust in healthcare systems. If this balance is achieved, South Africa could position itself as a global leader in responsible AI deployment, demonstrating that cutting-edge medical technology can harmonise with patient rights, fairness and social justice. However, this outcome hinges on collective action from policymakers, healthcare professionals, AI developers and civil society to shape AI in ways that amplify its promise while mitigating its risks. In this sense, the ethical concerns surrounding medical AI should not be seen as barriers but as guiding principles for building a transparent, equitable and human-centered future in healthcare.

Declaration. No AI-generated content is present in the article. AI tools were utilised solely for reference formatting.

Acknowledgements. None.

Author contributions. This article was entirely researched and independently written by the author.

Funding. None.

Conflicts of interest. None.

1. Kumar Y, Singh JP, Shukla A, et al. Artificial intelligence in disease diagnosis: A systematic literature review, synthesizing framework and future research agenda. *J Ambient Intell Humaniz Comput* 2023;14(7):8459-8486. <https://doi.org/10.1007/s12652-021-03612-z>.
2. Xiong Y, Ba X, Hou A, Zhang K, Chen L, Li T. Automatic detection of mycobacterium tuberculosis using artificial intelligence. *J Thorac Dis* 2018;10(3):1936-1940. <https://doi.org/10.21037/jtd.2018.01.91>
3. Lotter W, Hassett MJ, Schultz N, Kehl KL, Van Allen EM, Cerami E. Artificial intelligence in oncology: Current landscape, challenges, and future directions. *Cancer Discov* 2024;14(5):711-726. <https://doi.org/10.1158/2159-8290.CD-23-1199>
4. Aggarwal A, Mishra P, Basu S, Raj R, Shukla A. Artificial intelligence-based chatbots for promoting health behavioral changes: Systematic review. *J Med Internet Res* 2023;25:e40789. <https://doi.org/10.2196/40789>
5. Mashishi A. South Africa's artificial intelligence (AI) planning: adoption of AI by government. Pretoria: Department of Communications and Digital Technologies (DCDT); 2023. 1-47.
6. Krishnan G, Talby D, Lansky D, et al. Artificial intelligence in clinical medicine: Catalyzing a sustainable global healthcare paradigm. *Front Artif Intell* 2023;6:1227091. <https://doi.org/10.3389/frai.2023.1227091>.
7. Maleki Varnosfaderani S, Forouzanfar M. The role of AI in hospitals and clinics: Transforming healthcare in the 21st century. *Bioengineering (Basel)* 2024;11(4):337. <https://doi.org/10.3390/bioengineering11040337>.
8. Perez K, Johnson A, Smith B, et al. Investigation into application of AI and telemedicine in rural communities: A systematic literature review. *Healthcare (Basel)* 2025;13(3):324. <https://doi.org/10.3390/healthcare13030324>
9. Florent R, Fardman B, Podwojniak A, et al. Artificial intelligence in dermatology: Advancements and challenges in skin of color. *Int J Dermatol* 2024;63(5):455-461.
10. Adamson AS, Smith A. Machine learning and health care disparities in dermatology. *JAMA Dermatol* 2018;154(11):1247-1248. <https://doi.org/10.1001/jamadermatol.2018.2348>

11. Statistics South Africa. Mballo Brief: Monthly economic and social statistics—October 2023. Pretoria: Stats SA; 2023. <https://www.statssa.gov.za/>
12. Daneshjou R, Vodrahalli K, Novoa RA, et al. Disparities in dermatology AI performance on a diverse, curated clinical image set. *Sci Adv* 2022;8(32): eabq6147. <https://doi.org/10.1126/sciadv.abq6147>.
13. European Parliament and Council of the European Union. Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence and amending certain Union legislative acts (Artificial Intelligence Act). *Off J Eur Union* 2024 Jul 12; L 2024/1689. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
14. Republic of South Africa. Constitution of the Republic of South Africa, 1996. Pretoria: Government Gazette; 1996. <https://www.justice.gov.za/legislation/constitution/>
15. Republic of South Africa. Constitution of the Republic of South Africa, 1996. Section 27. Pretoria: Government Gazette; 1996. <https://www.justice.gov.za/legislation/constitution/>
16. Government of the Republic of South Africa and Others v Grootboom and Others (CCT11/00) [2000] ZACC 19; 2001 (1) SA 46 (CC); 2000 (11) BCLR 1169.
17. Minister of Health and Others v Treatment Action Campaign and Others (No 2) 2002 (5) SA 721 (CC); 2002 (10) BCLR 1033.
18. Schoeman R, Haines M. Radiologists' experiences and perceptions regarding the use of teleradiology in South Africa. *SA J Radiol* 2023;27(1):2647. <https://doi.org/10.4102/sajr.v27i1.2647>
19. Tasioulas J. Minimum core obligations: human rights in the here and now. Washington (DC): World Bank; 2017. <https://openknowledge.worldbank.org/handle/10986/28778>
20. Republic of South Africa. Constitution of the Republic of South Africa, 1996. Sections 10, 12, and 27. Pretoria: Government Gazette; 1996.
21. Republic of South Africa. National Health Act 61 of 2003. Pretoria: Government Gazette; 2004.
22. Republic of South Africa. Protection of Personal Information Act 4 of 2013. Pretoria: Government Gazette; 2013.
23. South African Health Products Regulatory Authority. Regulations relating to medical devices and in vitro diagnostic medical devices. Pretoria: SAHPRA; 2017. <https://www.sahpra.org.za>
24. Raab EL. The parameters of informed consent. *Trans Am Ophthalmol Soc* 2004;102:225-230; discussion 230-232.
25. Alum EU. AI-driven biomarker discovery: Enhancing precision in cancer diagnosis and prognosis. *Discov Oncol* 2025;16(1):313. <https://doi.org/10.1007/s12672-025-02064-7>
26. Parikh RB, Hasler JS, Zhang Y, et al. Development of machine learning algorithms incorporating electronic health record data, patient-reported outcomes, or both to predict mortality for outpatients with cancer. *JCO Clin Cancer Inform* 2022;6:e2200073. <https://doi.org/10.1200/CCI.22.00073>.
27. Lubarsky B. Re-identification of 'anonymised' data. *Georgetown Law Technol Rev* 2017;1(2). <https://georgetownlawtechreview.org>
28. Rocher L, Hendrickx J, de Montjoye YA. Estimating the success of re-identifications in incomplete datasets using generative models. *Nature Commun* 2019;10(1):3069. <https://doi.org/10.1038/s41467-019-10933-3>
29. Geisler J. Software for medical systems. In: Fowler K, editor. Mission-critical and Safety-critical Systems Handbook. Oxford: Newnes; 2010. 147-268. <https://doi.org/10.1016/B978-0-7506-8567-2.00004-4>.
30. Soori M, Ghaleh Jough FK, Dastres R, Arezoo B. AI-based decision support systems in Industry 4.0: A review. *J Econ Technol* 2024. <https://doi.org/10.1016/j.ject.2024.08.005>
31. Gordijn B, ten Have H. What's wrong with medical black box AI? *Med Health Care Philos* 2023;26:283-284. <https://doi.org/10.1007/s11019-023-10168-6>
32. Choudhary K, DeCost B, Chen C, et al. Recent advances and applications of deep learning methods in materials science. *NPJ Comput Mater* 2022;8:59. <https://doi.org/10.1038/s41524-022-00734-6>
33. Reid F, Pravinkumar SJ, Maguire R, et al. Using machine learning to identify frequent attendance at accident and emergency services in Lanarkshire. *Digit Health* 2025;11. <https://doi.org/10.1177/20552076251315293>
34. Kowert W. The foreseeability of human-artificial intelligence interactions. *Tex Law Rev* 2017;96:181-182.
35. Escola v. Coca Cola Bottling Co. of Fresno, 150 P.2d 436 (Cal. 1944).
36. Blackman J, Veerapen R. On the practical, ethical, and legal necessity of clinical artificial intelligence explainability: An examination of key arguments. *BMC Med Inform Decis Mak* 2025;25(1):111. <https://doi.org/10.1186/s12911-025-02891-2>
37. Garcia v. Character Techs., Inc., No. 6:24-CV-01903 (M.D. Fla. filed Oct. 22, 2024).
38. World Health Organization. Ethics and governance of artificial intelligence for health: WHO guidance. Geneva: WHO; 2021. <https://www.who.int/publications/i/item/9789240029200>
39. Organisation for Economic Co-operation and Development (OECD). Recommendation of the Council on Artificial Intelligence. Paris: OECD; 2019 [updated 2024]. <https://www.oecd.org/en/topics/ai-principles.html>

Received 26 February 2025. Accepted 14 May 2025.