Ethical governance of artificial intelligence in healthcare: Critical reflections of the AI Task Team of the South African Medical Association on the Health Professions Council of South Africa guidelines on the ethical use of AI in healthcare (Booklet 20)

S Mahomed,[1] LLM, PhD; **M V Ncube,**[2] MMed Sci, PhD; **A Dhai,**[3] MB ChB, FCOG, LLM, PhD , PGDip IntResEthics; **W Janneker,**[4] Dip Comp Sci; **M Nodikida,**[5] MPH, Dip Obst&Gyn (SA), MBA, Cert AI in Healthcare (Harvard), MD

[1] *Department of Jurisprudence, School of Law, University of South Africa, Pretoria, South Africa*
[2] *Head of Unit: Health Policy and Research Unit, South African Medical Association NPC*
[3] *Editor,* South African Medical Journal*; Editor,* South African Journal of Bioethics and Law*; Chair, HSRC Research Ethics Committee; Chair, SANBS Research Ethics Committee; Honorary Professor, Steve Biko Centre for Bioethics, University of the Witwatersrand; Honorary Specialist Clinical Ethics Consultant, Nelson Mandela Children's Hospital, Johannesburg, South Africa; Certified Mediator – Conflict Dynamics and Centre for Effective Dispute Resolution, UK*
[4] *Chief Executive Officer, AfriTech AI, Johannesburg, South Africa*
[5] *Chief Executive Officer, South African Medical Association NPC, Pretoria, South Africa*

Corresponding author: *S Mahomed (mahoms1@unisa.ac.za)*

Artificial intelligence (AI) is rapidly reshaping clinical practice, and has prompted the Health Professions Council of South Africa to publish Booklet 20 – its first ethical framework for AI use in healthcare. This review critically evaluates Booklet 20 through the lens of the South African Medical Association AI Task Team, examining its alignment with national legislation, emerging regulatory mechanisms, broader policy commitments and the ethico-social context. Drawing on the Protection of Personal Information Act, the South African Health Products Regulatory Authority's 2025 guidance for AI-enabled medical devices, the National AI Policy Framework and the 2024 National Health Research Ethics Council ethics guidelines, the analysis identifies key operational gaps relating to human oversight, disclosure, data sovereignty, equity, accountability and risk categorisation. The article argues that while Booklet 20 establishes an important foundation, its principles require concrete implementation tools, including risk-tiered safeguards, structured consent templates, meaningful governance co-ordination and context-appropriate standards for transparency, explainability and bias mitigation. Grounding these enhancements in South Africa's communitarian ethic of *ubuntu* highlights the need for relational accountability, fairness and community participation to ensure safe and equitable AI integration. The article concludes with a set of practical recommendations aimed at strengthening ethical governance and supporting patient trust and professional integrity as AI becomes embedded in clinical workflows.

Keywords: artificial intelligence, healthcare, HPCSA Booklet 20, ethico-regulatory governance, SAMA AI Task Team

S Afr Med J 2026;116(5):e5072. https://doi.org/10.7196/SAMJ.2026.v116i5.5072

The Health Professions Act 56 of 1974 established the Health Professions Council of South Africa (HPCSA) as the statutory authority responsible for regulating health professionals. Under this Act, the HPCSA oversees registration, education, training and ethical conduct across various health disciplines. Its ethical guidelines, derived directly as a mandate from the Act, form the foundation for professional practice, and function as normative standards against which professional practice is evaluated. In early 2025, the HPCSA released a draft of its ethical guidelines for the Responsible Use of Artificial Intelligence (AI) in Clinical Practice (Booklet 20), for public comment – marking a significant step in aligning emerging technologies with ethical healthcare standards, and signalling a way forward towards embedding ethical guardrails for AI across clinical practice. In December 2025, the HPCSA confirmed the publication of Booklet 20.[1]

Booklet 20[2] marks the HPCSA's first comprehensive stance on AI, and includes, among others, sections on patient welfare, transparency and practitioner accountability, and outlines the ethical expectations for the use of AI in healthcare. The AI Task Team of the SA Medical Association (SAMA), comprising medically, scientifically and legally qualified individuals, engaged the draft guideline after its initial release in early 2025 and offered structured comments on definitions, disclosure, accountability, equity, safety, clinical decision-making, data protection and regulatory alignment, which for reasons best known to the regulator were not considered when the approved version was published. Neither was there a second version sent out for comment. This article reflects on the suggestions of the AI Task Team with respect to Booklet 20, assesses them in light of SA's regulatory environment, and proposes enhancements to the guidelines for improvement.

Regulatory and policy context in SA

POPIA: The right to privacy, automated decision-making and cross-border transfers

The right to privacy is protected as a fundamental right under section 14 of the Constitution of the Republic of South Africa, 1996. This constitutional safeguard is reinforced by sector-specific legislation and policy frameworks within the health environment. Notably, the National Health Act 61 of 2003 protects patient confidentiality, while the Protection of Personal Information Act 4 of 2013 (POPIA)[3] serves as the central statute governing the processing of personal information. In section 1, POPIA defines personal information broadly, and embeds principles of autonomy and self-determination by regulating how such information may be collected, used, stored and shared. Its eight conditions for lawful processing, ranging from accountability to data subject participation, must be met by the responsible party, in this case a health practitioner or researcher, to ensure that processing respects constitutional privacy rights. Data subjects must be informed of the nature, purpose, retention, sharing and potential cross-border transfer of any personal information being collected.[4]

Of particular relevance to AI in the context of healthcare is section 71 of POPIA,[3] which places a general prohibition on decisions that are based solely on automated processing without any human oversight. This section restricts decisions with legal or substantial effects that are based solely on automated processing, requiring appropriate measures, including an opportunity to make representations, and sufficient information about the underlying logic of the automated processing of the information (section 71(3)(b)). Section 71 of POPIA directly engages emerging uses of algorithmic and AI-driven decision-making, where systems can generate rapid yet opaque outcomes that may significantly affect individuals, and ensures that data subjects (patients/participants) retain the right to challenge, request reasons for or contest decisions made purely by automated means, recognising the risk that AI systems may reproduce or amplify existing biases.[5] This safeguard is essential in protecting individuals from unfair or prejudicial automated determinations, particularly in sensitive areas such as healthcare, research participation and eligibility assessments. Furthermore, it has direct implications for clinical AI that profiles patients or informs diagnoses, especially where decisions might be perceived as 'solely automated', unless human oversight is substantive and documented.

In the context of data transfers, section 72 of POPIA[3] becomes operative, and introduces an additional safeguard for the protection of personal information. Domestic transfers are permissible where informed consent has been obtained and ethics review has been undertaken. Although POPIA provides five grounds for the transfer of personal information to foreign jurisdictions, a single ground appears practically viable: an international transfer is permissible where the receiving party in the foreign country is bound by a legal framework, binding corporate rules, or a binding agreement that affords an adequate level of protection and upholds principles that are substantially equivalent to those contained in POPIA (section 72(1)).

Consequently, a binding contractual mechanism that incorporates safeguards consistent with POPIA's requirements for processing personal information appears to offer a realistic and legally sound basis for cross-border transfers. In the context of AI being used in clinical practice, where many AI vendors operate international servers, a health practitioner will have to ensure that POPIA-compliant operator contracts are in place.

SAHPRA: AI-enabled medical devices

The definition of a medical device is contained within the Medicines and Related Substances Act 101 of 1965; however, the Act and its 2015 amendment were drafted before the emergence of AI-specific medical technologies as regulatory categories. Thus, it historically regulated devices in broad terms without differentiating between software-driven, algorithmic, or adaptive technologies. However, this position changed in the last quarter of 2025, when the SA Health Products Regulatory Authority (SAHPRA) communicated supplementary guidance to stakeholders regarding this gap. The Regulatory Requirements for AI and Machine Learning (ML)-Enabled Medical Devices (2025)[6] issued by SAHPRA in September 2025 describes an AI/ML-enabled medical device as 'a product that conforms to the definition of a medical device and utilises one or more AI or machine-learning algorithms to perform, in part or in whole, its intended medical purpose'. This includes but is not limited to AI/ML applications in:

- medical imaging analysis (e.g. software using AI to detect tumours or fractures in radiological images);
- predictive algorithms (e.g. an ML model that forecasts risk of patient deterioration);
- clinical decision support systems (e.g. AI-driven diagnostic aids or treatment recommendation systems for healthcare professionals); and
- wearable health monitoring technologies (e.g. wearables that analyse bio-signals and alert users to abnormalities).

An AI/ML-enabled medical device can either be software as a medical device (SaMD) or software embedded in a medical device (SiMD). The former is distinguished by the fact that it performs its intended functions without being part of a hardware medical device. Mobile applications that meet the definition above are considered SaMD.[6] Therefore, AI/ML devices must meet existing requirements for medical devices, including the correct risk classification, International Organization for Standardization (ISO)-aligned quality management, risk assessment, and vigorous technical and clinical evidence indicating safety, effectiveness and bias mitigation.[7] The core regulatory principles emphasise patient protection, the need for meaningful human oversight, and transparency and explainability. They also stress robust cybersecurity measures, preservation of data quality and adherence to POPIA. Manufacturers are expected to conduct ongoing post-market monitoring to identify any decline in system performance, and to report adverse events. SAHPRA further notes the regulatory complexities introduced by adaptive algorithms and generative AI, indicating that additional guidance, particularly on change-control mechanisms, will be issued to promote responsible technological development while maintaining patient safety and public confidence.[6] Therefore, alignment between the HPCSA's ethical duties and SAHPRA's technical assurance is essential, particularly as SAHPRA updates device classification and regulatory requirements relevant to AI in the healthcare sector.

Africa contains the highest level of human genetic diversity globally, yet <2% of human genomes analysed are those of African people.[8] This under-representation has material consequences for clinical safety, pharmacovigilance and the applicability of evidence used in routine care in SA.[9] Most therapeutic guidelines, predictive models and medical-enabled device algorithms are developed using data from non-African populations, which may not adequately reflect local biological, environmental and phenotypic variation.

Policy framework in SA

The SA National AI Policy Framework (2024)[10] positions healthcare as one of the critical sectors where AI can deliver major societal benefits. Although the document does not contain a stand-alone healthcare chapter, it incorporates multiple strategic provisions directly relevant to clinical practice, health system governance, medical AI safety, and ethical oversight. It highlights the increasing social demand for AI-enabled solutions in essential public services, including healthcare, noting that such technologies hold the potential to enhance service delivery and improve overall quality of life. Central to the framework's orientation towards medical and clinical applications is a commitment to human-centred AI, which prioritises the augmentation, rather than replacement, of professional clinical judgement, similar to the provisions contained in Booklet 20. The framework stresses that ethically grounded AI must remain transparent, accountable and fair to support responsible clinical decision-making and maintain public trust. It also places significant emphasis on explainability, arguing that interpretable systems are vital for high-stakes environments such as medicine, as they enable clinicians to justify decisions, detect biases and uphold accountability standards.[10]

In addition, the framework underscores that responsible deployment of AI in healthcare requires robust privacy and data-protection measures, aligned with national data-governance principles, to ensure the ethical handling of sensitive patient information. It calls for strengthened data-protection regulations, transparent data-usage practices and rigorous cybersecurity protocols to safeguard both personal data and AI-driven clinical systems. The document furthermore identifies the infrastructural and capacity-building foundations necessary for a resilient AI-enabled health sector, including investment in digital connectivity, supercomputing capacity and workforce training to ensure that clinicians and health-system actors can safely adopt and oversee AI tools. Finally, the framework signals that sector-specific regulatory strategies, including those governing medical AI, will flow from this national policy foundation, creating a coherent governance environment designed to promote innovation while upholding patient safety, professional responsibility and public confidence in AI-assisted clinical practice. SAHPRA's 2025 Regulatory Requirements for AI/ML Enabled Medical Devices, as described above, operationalise the general principles of the National AI Policy Framework, and translate the vision of the Policy into enforceable regulatory obligations for AI/ML enabled medical devices.

The *South African Ethics in Health Research Guidelines: Principles, Structures and Processes* (2024), issued by the National Health Research Ethics Council (NHREC) (the NHREC Guidelines),[11] introduces an ethics framework that responds to the intensified privacy and data-protection risks emerging from new digital technologies in healthcare and, increasingly, in health-related AI research. Now in the third edition, the guidelines set the national minimum standard for responsible research practice in SA, and explicitly recognise that data sharing raises unique ethical concerns. Decisions about sharing or transferring data necessarily require balancing the protection of individual privacy with the benefits of facilitating scientific progress. In this regard, the guidelines provide detailed direction for both researchers and research ethics committees (RECs), reiterating the data-processing obligations established under POPIA and advocating for the use of dedicated material transfer agreements and data transfer agreements when research data are exchanged (sections 4.2.2. and 4.2.2.1).

When situated within the AI and health landscape, these guidelines serve a critical governance role by addressing the ethical complexities introduced by AI-driven tools in clinical practice and health research. They align closely with the broader national emphasis on human-centred, safe and equitable AI, requiring RECs to evaluate AI research using criteria such as transparency, explainability, responsibility and accountability, fairness, benefit sharing, and risk-mitigation strategies. These considerations are particularly salient in healthcare, where algorithmic outputs influence high-stakes clinical decisions, and where risks of bias, discrimination and harm are magnified. The guidelines also underscore that ethical oversight of AI applications should be continuous and adaptive, reflecting the iterative nature of ML systems that may evolve over time. To support rigorous review, they provide RECs with a set of 11 guiding questions, including assessments of whether an AI system is appropriate for the SA context, how the privacy interests of vulnerable groups will be protected and what safeguards apply to the handling of participant data after death. Recognising the technical complexities inherent in AI and big-data research, the guidelines recommend that RECs draw on specialised data-science expertise, and engage both researchers and participants in evaluating large-scale data-driven projects (section 3.4.4.2). In doing so, the NHREC guidelines play a foundational role in anchoring AI-enabled health research within a robust ethical, participatory and context-sensitive governance framework, one that complements emerging regulatory instruments such as SAHPRA's requirements for AI/ML-enabled medical devices and broader national policy commitments to safe, responsible AI integration across the health sector.

SAMA AI Task Team's reflections on HPCSA Booklet 20

As described above, Booklet 20 provides ethico-regulatory guidance for the use of AI in clinical practice, and emphasises that AI may support, but never replace, the clinical judgement of healthcare practitioners. The guidelines apply to all practitioners registered under the HPCSA, and outline expectations around ethical use, disclosure, accountability, equity, safety, continuing professional development (CPD), clinical decision-making, data privacy and regulation, and include in section 12 principles on the 'pillars of AI'. The SAMA AI Task Team welcomes this direction, and identifies critical operational gaps that could hinder safe, fair and practicable deployment.

From the Task Team's submission to the HPCSA, priority themes included:

- a systems approach to promote workflow transparency and distributed accountability;
- clarity on governance alignment across the HPCSA, SAHPRA, NHREC and the information regulator;
- embedding equity into data and system design (bias mitigation) and disclosure;
- aligning definitions, e.g. automated decision-making, with POPIA section 71;
- including definitions, e.g. sovereign data as a critical concept;
- providing risk categorisation tools and an informed consent template;
- specifying complaints pathways and dispute resolution; and
- expansion on CPD competencies and explainability standards.

A discussion of certain reflections by the Task Team follows in the section below.

Human oversight and clinical accountability

Booklet 20 emphasises that while AI can support clinical practice, the practitioner remains fully accountable for all decisions, reflecting national and global standards such as those endorsed by the 2021 World Health Organization (WHO) guidance on *Ethics and Governance of Artificial Intelligence for Health*,[12] which states that

'In the context of health care … humans should remain in control of health-care systems and medical decisions.'[12] However, a gap arises in relation to section 71 of POPIA, where the prohibition on 'solely automated decision-making' may unintentionally be breached if human involvement is superficial rather than substantive. To prevent this, Booklet 20 should clearly articulate what constitutes *meaningful human oversight* in clinical environments, encompassing interpretable AI outputs, practitioner capability to assess and question algorithmic reasoning, and transparent documentation whenever AI influences a clinical decision. In this context, explainability becomes both a legal safeguard and an ethical requirement, demanding practical methods to ensure that AI-supported decisions remain clinician-led rather than algorithm-driven. Additionally, appropriate mechanisms should be put in place to facilitate equitable dispute resolution practices, which are currently lacking in Booklet 20. Jurisdictional uncertainty is a significant gap in Booklet 20. AI in healthcare routinely involves cross-border data transfers, cloud-based processing and foreign-developed algorithms, all of which create ambiguity regarding which legal system governs disputes or harms. The HPCSA's mandate extends only to regulating practitioners registered in SA, yet many AI tools originate from companies domiciled outside the country, making accountability for malfunction, bias, or data breaches difficult to operationalise. To ensure fairness and legal certainty, Booklet 20 should expressly require that AI vendors operating in the SA health sector agree to dispute-resolution within SA. Without such mechanisms, practitioners may be unfairly burdened with liability for harms that originate from opaque or foreign-controlled AI, undermining patient protection and professional accountability.

Transparency, disclosure and consent

A more rigorous disclosure framework is required to ensure that patients fully understand when and how AI contributes to their care. This should include clear explanations of what the tool does, its limitations, the nature and representativeness of its training data, and any foreseeable risks, supported by an HPCSA-approved consent template appended to Booklet 20 and adjusted according to the risk level of the AI system. Given WHO guidance and the European Union (EU) AI Act's[13] emphasis on transparency and integrity, SA's multilingual and unequal digital environment makes standardised, practical templates essential to avoid superficial disclosure. Ultimately, informed consent for algorithmic tools must be re-envisioned to reflect contextual realities. Substantive sections of a proposed informed consent template could include:

(*i*) Purpose of the AI system: a clear explanation of why the AI tool is being used and the specific role it plays in the patient's care.
(*ii*) Role of AI in clinical decision-making: a description of how the AI contributes to clinical assessments or recommendations, and how the clinician will interpret and apply its outputs.
(*iii*) Benefits, risks, and limitations: an outline of the anticipated advantages, potential risks, uncertainties and known limitations associated with the AI tool.
(*iv*) Data sources and localisation: information about the types of data the AI relies on, where data are stored or processed, and whether the system has been trained or validated on populations relevant to SA.
(*v*) POPIA and data-protection rights: a summary of the patient's rights under POPIA, including access, correction, objection to processing, and how their information is protected.
(*vi*) Right to decline or opt out (where clinically appropriate): a statement indicating when patients may refuse or opt out of AI-assisted processes without compromising essential care, if clinically safe to do so.

Additionally, in section 4.1 of Booklet 20 it is indicated that a health practitioner should only use technology that is not a secret or claimed to be secret. However, instead of leaving this decision up to an individual health practitioner, a benchmark or set of minimum standards should first be met, to ensure that the AI tool is suitable for its intended purpose. SAHPRA could be central to setting these benchmarks and developing a list of SAHPRA-approved technologies that meet the requirements of all necessary standards regulators.

Data privacy and sovereignty

Health information is classified as special personal information under section 26 of POPIA, meaning that cross-border data flows, cloud-based processing and the use of third-party digital platforms raise significant concerns regarding data sovereignty and cybersecurity. To address these risks, Booklet 20 should explicitly reference section 72 of POPIA, require formal data processing or data transfer agreements for any external handling of health data and mandate robust technical safeguards such as encryption, strict access controls and comprehensive audit logs. Sovereign data, which refers to data that are generated, stored and governed within a specific nation's legal and regulatory framework,[14] ensuring that the country retains full control over their use, access and security, must be included in the definitions within Booklet 20. This clarification is increasingly urgent as SAHPRA continues to update its governance of AI-enabled and digital medical devices, making it essential for the HPCSA's expectations on data protection to align coherently with POPIA and SAHPRA's evolving approval processes, and the information regulator's guidance. Furthermore, POPIA requires that personal information be kept only for as long as is necessary to achieve the specific purpose for which it was collected, after which it must be destroyed, deleted, or de-identified (sections 13 and 14). As many AI systems continually store, replicate or learn from patient data across distributed servers (often outside SA), clear guidance is needed to ensure that practitioners comply with lawful retention periods, and that operators maintain verifiable deletion protocols. Incorporating post-processing safeguards into Booklet 20 would align it with POPIA's sections on purpose-specific retention and international data transfers, and would protect patients from having their personal information stored indefinitely.

The SAMA Task Team further referenced the Human Sciences Research Council's risks and mitigation guideline[15] for using AI or cloud-based data analysis applications in its submission to the HPCSA (Table 1).

Ubuntu, equity and context-specific guidance

SA's constitutional commitments and communitarian ethos, most notably the value of *ubuntu*, call for an approach to clinical AI grounded in relational accountability, collective welfare and substantive fairness. In this context, the SAMA Task Team's recommendation to embed *ubuntu* within accountability frameworks must be understood as a practical directive rather than a symbolic gesture: it requires genuine co-creation with communities, transparent channels for reporting and responding to harm, and solidarity-oriented deployment of AI tools so that innovations do not privilege urban, advantaged patients while marginalising rural or underserved populations. By centering *ubuntu*, the ethical terrain becomes clearer: equity audits, plain-language disclosure, community participation and scrutiny of structural biases within training datasets are not optional enhancements but essential components for ensuring legitimacy, trust and relational autonomy in AI-supported healthcare.[16] Imported AI systems carry a significant risk of being poorly calibrated for SA's diverse populations, which can reinforce or deepen existing health

Table 1. HSRC's risks and mitigation for using AI or cloud-based data analysis applications

Risk	Definition	Mitigation
1. Data protection	By uploading raw data to a third-party AI platform, healthcare practitioners and researchers may face difficulty ensuring full compliance with these regulations, as they may have limited control over data storage, transfer and processing by the provider.	Clear agreements with service providers, including guarantees on data processing practices and certifications, can help ensure a level of data protection. Additionally, pseudonymisation or encryption of raw data before uploading can mitigate some risks.
2. Privacy	Raw data often include PII or sensitive information that could compromise individual privacy if accessed by unauthorised parties or misused by the AI platform. Data uploaded to cloud-based platforms may also be subject to foreign government access requests, potentially infringing on the privacy rights of individuals in jurisdictions with stricter privacy laws.	Privacy can be protected by minimising PII in datasets before upload, restricting data access to only those parties explicitly involved in the analysis, and choosing AI service providers that meet international privacy standards and adhere to strict privacy policies.
3. Confidentiality	Confidentiality entails that only authorised individuals or systems have access to data. Uploading raw data to AI models often involves sharing data with third-party platforms or providers, which increases the risk of data leaks, breaches, or misuse. If the platform's confidentiality measures are inadequate, sensitive data could be exposed, especially if contractual or technical safeguards are weak.	To preserve confidentiality, healthcare practitioners and researchers should use AI models or cloud-based platforms that offer end-to-end encryption, controlled access and detailed audit trails. Confidentiality agreements and third-party audits can further ensure that service providers uphold high standards of data security.
4. Anonymity	Many AI models require large datasets to function effectively, which can make full anonymisation challenging, especially if the data contain unique identifiers or patterns that could reveal individual identities indirectly. If the AI model is not properly secured, data may be subject to re-identification attacks, especially with sophisticated cross-referencing.	Healthcare practitioners and researchers should consider de-identifying or anonymising data before uploading them. Techniques such as differential privacy or aggregation can help reduce re-identification risks. Additionally, using synthetic datasets or limiting the granularity of uploaded information can help ensure individual anonymity while still enabling valuable insights from the analysis.

HSRC = Human Sciences Research Council; AI = artificial intelligence; PII = personally identifiable information.

inequities. To mitigate this, any AI tool deployed clinically should be required to demonstrate rigorous local validation, including bias testing across relevant demographic and clinical subgroups, and routine performance monitoring. Developers and practitioners should also have a continuing duty to report and address any material bias that becomes apparent after deployment. This approach aligns with global best practice, echoing the WHO's emphasis on inclusiveness and equity, the EU AI Act's obligations for robust data governance and continuous monitoring in high-risk systems, and SA's own National AI Policy Framework, which identifies fairness as a core principle for trustworthy AI.

Risk categorisation and safeguards

A three-tiered clinical AI risk classification (low, moderate and high) is recommended for Booklet 20 to guide practitioner use, aligned with SAHPRA's existing medical-device categories and consistent with the WHO's risk-based approach. For each tier, the guidelines should specify the minimum required safeguards, including the appropriate level of disclosure, the degree of clinical oversight and the expected frequency of performance audits and monitoring activities. For example:

- Low-risk AI (administrative functions without clinical inference): use should be recorded, with minimal disclosure provided through standard privacy notices.
- Moderate-risk AI (decision-support tools subject to clinician review): requires structured patient disclosure, documented human oversight and routine checks for accuracy and bias.
- High-risk AI (systems informing diagnosis, triage, or therapeutic recommendations): necessitates a dedicated informed-consent addendum, evidence of local validation, maintained system logs, a robust audit schedule and a clear incident-reporting pathway.

Recommendations

From the commentary to Booklet 20 provided by the SAMA Task Team, the following recommendations provide practical approaches to theoretical discourse.

(*i*) A strengthened regulatory approach for clinical AI should begin by defining what constitutes *meaningful human oversight*, and establishing standards to prevent superficial approval of algorithmic outputs. This must be complemented by a set of risk-tiered safeguards that align the HPCSA's ethical duties with SAHPRA's device-risk classifications, together with mandatory structured disclosure supported by an HPCSA-endorsed consent template for AI use.

(*ii*) Robust local validation and subgroup-specific bias testing should be required, accompanied by periodic audits and a clear obligation to report newly detected biases.

(*iii*) Effective governance also depends on clarifying the respective roles of the HPCSA, SAHPRA, NHREC and the information regulator. In addition, operationalising sections 71 and 72 of POPIA in the clinical AI context is essential, ensuring explainability for affected patients and strong safeguards for any cross-border data transfers.

(*iv*) Finally, *ubuntu*-based equity benchmarks, including language accessibility, rural inclusion and suitability for the public-sector context, should be embedded directly into approval criteria and deployment guidance to ensure that AI systems advance fairness rather than deepen existing inequalities.

Conclusion

The HPCSA's Booklet 20 provides an important and timely foundation for guiding the ethical use of AI in SA healthcare. To ensure that its principles are effectively translated into daily practice,

the guideline should furnish clinicians and institutions with concrete implementation tools, such as risk-tiered safeguards, standardised consent templates, practical oversight checklists and clear audit requirements, while also clarifying the respective roles of multiple regulatory bodies and ensuring alignment with POPIA, SAHPRA, NHREC, and the National AI Policy Framework. Grounding these measures in the values of *ubuntu* and equity will be essential for fostering patient trust, strengthening safety and upholding professional integrity as AI becomes increasingly integrated into clinical workflows.

Declaration. None.

Acknowledgements. This article was presented as a paper at the Digital Technology Conference hosted by UNISA's College of Law. MVN and SM presented the paper on behalf of the SAMA AI Task Team.

Author contributions. The SAMA AI Task Team conceptualised and prepared a comprehensive submission to the HPCSA on Booklet 20, after which the submission was further developed and expanded into this manuscript for publication.

Funding. None.

Conflicts of interest. None.

1. Health Professions Council of South Africa. Updated ethical guidelines. HPCSA Corporate Affairs, 4 December 2025. https://www.hpcsa-blogs.co.za/updated-ethical-guidelines/?utm_source=ezmed.solutions (accessed 27 March 2026).
2. Health Professions Council of South Africa. Ethical guidelines on the use of AI. Booklet 20. Pretoria: HPCSA, September 2025. https://www.hpcsa-blogs.co.za/wp-content/uploads/2025/11/ETHICAL-GUIDELINES-USE-OF-AI.pdf (accessed 27 March 2026).
3. South Africa. Protection of Personal Information Act No. 4 of 2013. https://www.gov.za/sites/default/files/gcis_document/201409/3706726-11act4of2013popi.pdf (accessed 18 March 2026).
4. Mahomed S. Data privacy and protection in AI-driven healthcare. S Afr Med J 2025;115(5b):e3666. https://doi.org/10.7196/SAMJ.2025.v115i5b.3666
5. Klang E, Tessler I, Freeman R, Sorin V, Nadkarni GN. If machines exceed us: Health care at an inflection point. NEJM AI 2024;1(10). https://doi.org/10.1056/AIp2400559
6. South African Health Products Regulatory Authority. Communication to stakeholders: Regulatory requirements of AI and machine learning-enabled devices. Pretoria: SAHPRA, 26 September 2025. https://www.sahpra.org.za/wp-content/uploads/2025/09/MD08-20252026_-SAHPRA-Communication-to-Industry-AI-Medical-devices_Acknowledgements.pdf (accessed 27 March 2026).
7. South African Medical Association. AI in clinical practice. Pretoria: SAMA, 2025. https://samedical.org/wp-content/uploads/2025/12/AI-MEDICAL-REPORT-SPREADS.pdf (accessed 27 March 2026).
8. Wonkman A. Sequence three million genomes across Africa. Nature, 10 February 2021. https://www.nature.com/articles/d41586-021-00313-7 (accessed 27 March 2026).
9. Ojewunmi OO, Fatumo S. Driving global health equity and precision medicine through African genomic data. Human Molecular Genetics 2025:ddaf025. https://doi.org/10.1093/hmg/ddaf025
10. Department of Communications and Digital Technologies, South Africa. SA AI Policy Framework, 2024. Pretoria: DCDT, 2024. https://www.dcdt.gov.za/sa-national-ai-policy-framework/file/338-sa-national-ai-policy-framework.html (accessed 18 May 2026).
11. National Health Research Ethics Council, South Africa. South African Ethics in Health Research: Principles, Processes and Structures. 3rd ed. Version 3.1. Pretoria: National Department of Health, 2024.
12. World Health Organization. Guidance, Ethics and Governance of AI for Health 2021. Geneva: WHO, 2021. https://www.who.int/publications/i/item/9789240029200 (accessed 18 May 2026).
13. European Union. Regulation (Eu) 2024/1689 of the European Parliament and of the Council of 13 June 2024. EUR-Lex, 2024. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689 (accessed 18 May 2026).
14. Department of Public Service and Administration, South Africa. Determination and directive on the usage of cloud computing services in the public service. Pretoria: DPSA, 2022. https://www.dpsa.gov.za/dpsa2g/documents/egov/2022/egovernment_02_02_2022.pdf (accessed 18 May 2026).
15. Human Sciences Research Council. Standard operating procedure ethical data handling for AI and cloud-based platforms. Pretoria: HSRC, 2025. https://hsrc.ac.za/wp-content/uploads/2025/06/Research-Ethics-Data-Handling-for-AI-and-Cloud-SOP_2025_v1_-March2025.pdf (accessed 27 March 2026).
16. Odero B, Nderitu D, Samuel G. The ubuntu way: Ensuring ethical AI integration in health research. Wellcome Open Res 2024,9:625. https://doi.org/10.12688/wellcomeopenres.23021.1

Received 1 April 2026; accepted 13 May 2026.