TGA consultation paper - Clarifying and strengthening the regulation of AI

Submission from The Social Policy Group

The Social Policy Group (SPG) is a national, non-government, not-for-profit body with specialist expertise in social policy and program design with a focus on population diversity, social and community cohesion, gender equality, health and inclusion. The Social Policy Group is funded as a Health Peak and Advisory Body.

Introduction

The Social Policy Group (SPG) welcomes the opportunity to respond to the Therapeutic Goods Administration's (TGA) consultation on clarifying and strengthening the regulation of Artificial Intelligence (AI) in healthcare. While we recognise the TGA's vital role in safeguarding public health through effective regulation, the current approach outlined in the consultation paper does not fully address the challenges posed by AI technologies. The consultation paper, which proposes incremental amendments to the Therapeutic Goods Act (1989), reflects an attempt to fit AI within an existing regulatory structure not designed to accommodate such technologies. In our view, AI represents not merely a new class of medical device but a fundamental transformation in how healthcare is delivered, decisions are made, and systems are integrated. This requires more than adjustments to existing legislation - it requires a new regulatory paradigm.

The SPG contends that AI technologies - due to their dynamic, evolving nature and far-reaching impact across clinical, administrative, and patient-facing applications - necessitate a comprehensive, AI-specific regulatory framework. This submission will critically examine why the Therapeutic Goods Act is inadequate for regulating AI, explore how the TGA's current approach may leave regulatory gaps and accountability issues unresolved, and propose the necessary components of a robust evaluation and oversight framework for AI in healthcare.

AI as a Dynamic Learning System

The Therapeutic Goods Act and its amendments are built around static, clearly defined medical devices that perform specific functions within predictable parameters. Medical devices such as imaging machines, prosthetics, and surgical tools can be assessed for safety and efficacy at a single point in time, with their performance monitored over their lifecycle through periodic post-market surveillance. This model assumes that once a device is certified, its performance remains consistent and predictable, provided that it is used as intended and any defects are identified through routine surveillance.

AI fundamentally challenges this static model, particularly machine learning (ML) and deep learning systems. AI systems are designed to evolve by learning from new data inputs and adapting to clinical environments. While this adaptability offers significant benefits, such as improved performance over time, it also introduces risks. An AI system that initially performs well may begin to degrade in accuracy or develop biases as it processes new data or encounters patient populations underrepresented in its training set. These changes can lead to “scope creep,” where AI outputs gradually diverge from their original function.

It is understood that AI technology performs differently than conventional technologies and that the capture of emerging tools under TGA's medical device classifications is weaker than that of more established medical technologies. Acknowledging this, the Therapeutic Goods Act does not currently provide a mechanism for continuous oversight or real-time monitoring, which are critical for ensuring the safety of AI systems throughout their lifecycle. Traditional post-market surveillance is ill-suited to detect these subtle but potentially harmful shifts in performance, leaving healthcare providers and patients exposed to unforeseen risks. AI systems require ongoing validation and real-time evaluation to ensure they continue to meet safety and efficacy standards as they evolve.

Expanding the Definition of “Medical Device”

The Therapeutic Goods Act defines medical devices in a way that reflects the traditional boundaries of healthcare - tools and instruments used to diagnose, treat, or prevent disease in patients. These devices have precise, tangible functions, and their role in healthcare is relatively easy to delineate. AI, however, blurs these lines. AI systems are not necessarily discrete tools but can be embedded, like all software, within more generalised systems and networked systems, performing a wide range of functions that extend beyond the traditional scope of a medical device.

AI influence in healthcare is expansive, touching not only clinical diagnostics and treatment recommendations but also administrative tasks, resource management, patient engagement, and population health. For instance, AI systems are used to optimise hospital workflows, predict patient admission rates, and allocate staffing resources. These systems may not fit neatly within the existing definition of a medical device, yet their performance can directly affect patient outcomes. If an AI system designed to manage hospital staffing levels assigns human resources inappropriately, this could lead to delays in patient care, strained hospital capacity, and adverse clinical outcomes.

AI systems integrated into healthcare operations but not strictly "medical devices" as defined by the Act fall outside the TGA's regulatory scope. Many AI systems that impact health outcomes are not captured under the medical device classification and do not behave like a therapeutic good. Iterative software products, consumer technologies, supermassive integrated systems, and public and preventive health systems are just the tip of the iceberg. The narrow definition of "medical device" leaves significant gaps in oversight, particularly for AI systems that influence patient outcomes indirectly. The regulatory framework must evolve to reflect the broad and systemic role AI plays in healthcare, ensuring that all relevant AI applications are subject to appropriate regulation and oversight.

AI as a Service: Moving Beyond Product-Based Regulation

One of the fundamental limitations of the current TGA framework is its product-based approach to regulation. This model works well for traditional medical devices, which are evaluated as discrete products with defined functions. However, AI technologies do not fit neatly into this framework. AI systems operate as services rather than products, continuously interacting with data and making real-time decisions.

The concept of "AI-as-a-service" recognises that AI systems are not standalone tools but are embedded within the broader healthcare environment. For example, an AI system used to predict patient outcomes or recommend treatment pathways is not a one-time-use product; it is a service that continuously interacts with patient data, healthcare providers, and other systems to deliver its outputs. This service-oriented nature of AI requires a new regulatory approach that goes beyond the product-based model.

Under an AI-as-a-service framework, regulatory oversight would extend beyond the initial approval process to include ongoing monitoring and independent evaluation of AI systems throughout their lifecycle. This would involve regular evaluations of AI systems, focusing on how they are evolving, interacting with other systems, and whether they continue to meet safety and performance standards. The framework would also include clear guidelines on the role of human oversight in AI-driven decision-making. Healthcare professionals must retain ultimate responsibility for decisions made with the assistance of AI, and there should be protocols in place to ensure that AI recommendations are reviewed and validated by clinicians.

The AI-as-a-service model would also ensure that AI systems are designed with transparency and accountability in mind. AI developers should be required to provide clear documentation on how their systems operate, how decisions are made, and how the system learns and adapts over time. This transparency is critical for ensuring that healthcare providers can trust AI systems and that patients understand how AI is used in their care.

AI as an Augmentation to Workflow and Workforce Management

In addition to its role as a diagnostic integration, AI increasingly augments healthcare workflows and workforce management. AI systems are being deployed to integrate data across electronic health records (EHRs), assist in clinical decision-making, and support operational efficiency in hospitals and clinics. These systems do not operate in isolation but are embedded within the fabric of healthcare delivery, augmenting the decision-making processes of clinicians, administrators, and policymakers.

For example, AI systems can streamline hospital operations by integrating clinical and administrative data to forecast patient admissions, optimise bed occupancy, and predict staffing needs. These systems are integral to the smooth functioning of healthcare organisations, and their influence extends beyond direct patient care to the broader management of healthcare resources. If an AI system used to predict patient flow fails or produces inaccurate results, the ripple effects could impact everything from hospital bed availability to the allocation of critical care resources. This, in turn, affects the quality of care patients receive.

However, the Therapeutic Goods Act does not adequately account for these AI systems that operate at the intersection of clinical and operational functions. AI integration into workflow and workforce management necessitates a broader regulatory approach that recognises its systemic impact and the potential risks associated with its failure.

Challenges in Assigning Accountability

As AI becomes more autonomous and integrated into healthcare workflows, the question of who is responsible for AI-driven decisions becomes increasingly complex. Traditional medical devices are typically used by healthcare professionals, who retain responsibility for their use. However, as AI systems begin to make decisions or recommendations without human intervention, the line between human and machine accountability becomes blurred.

For instance, if an AI system used to recommend treatment plans makes an error that results in patient harm, who is responsible - the developer of the AI system, the healthcare provider who relied on the system, or the hospital that deployed it? The Therapeutic Goods Act must provide clear guidance on assigning accountability in such cases, leaving healthcare providers and developers uncertain about their legal and ethical responsibilities. Without a clear accountability framework, the TGA risks being unable to effectively manage the legal and regulatory implications of AI-driven errors, particularly as AI systems become more autonomous.

Incremental Amendments Are Insufficient for AI Regulation

The TGA's consultation paper suggests that amendments to the Therapeutic Goods Act will be sufficient to bring AI within its regulatory framework. However, this approach fails to recognise the fundamental differences between AI and traditional medical devices. AI systems' continuous learning, adaptive nature, and integration into broader systems cannot be managed through incremental changes to legislation designed for static technologies.

It is essential to acknowledge that the TGA's consultation seeks feedback on how well medical devices legislation accommodates or conforms to the evolving Department of Industry's guardrails model rather than focusing on how well TGA legislation serves to govern AI in health more broadly. This approach evades the bigger question of what role

the TGA will have in the total governance of AI technologies used in health or why the TGA's legislative model should be preserved in this context.

One key limitation of the proposed amendments is their focus on expanding post-market surveillance without addressing the need for real-time oversight of AI systems. AI systems require continuous monitoring to detect issues such as performance drift, bias, or degradation in accuracy. Retrospective checks cannot adequately safeguard patients from these risks, as the evolution of the AI system is constant, and new risks can emerge over time.

Furthermore, the TGA's proposed amendments do not account for network effects, where AI systems interact with other technologies, including other AI systems. As AI systems become more deeply integrated into healthcare operations - interacting with EHRs, wearable devices, and real-time monitoring systems - they create dependencies between historically separated systems. If one system performs unacceptably, it could cause cascading failures across the healthcare environment. The current regulatory framework does not provide sufficient mechanisms for managing these interdependencies, leaving healthcare providers vulnerable to system-wide risks.

Continuous Monitoring and Real-Time Validation

To effectively regulate AI systems in healthcare, the TGA must significantly expand its independent evaluation capacity for AI and develop a regulatory ecosystem to support continuous oversight. Traditional post-market surveillance mechanisms, which rely on periodic checks to ensure compliance, are inadequate for managing the constant evolution of AI systems. Given AI systems' ability to learn, adapt, and adjust their algorithms based on new data, real-time monitoring and validation are essential to ensure that AI systems remain safe and effective throughout their lifecycle.

Continuous monitoring would involve using advanced tools to track AI systems in real time, identifying performance deviations, biases, or errors as they occur. For example, an AI system used in diagnostics should be continuously monitored to ensure its predictions remain accurate as it encounters new types of patient data. If the system shows performance degradation or bias, it should be flagged for immediate review and revalidation. This real-time oversight would allow the TGA to intervene before AI systems cause harm rather than relying on post-market surveillance to detect issues after the fact.

Real-time validation is critical for ensuring AI systems remain safe as they adapt to new data and environments. Unlike traditional devices, which can be revalidated through periodic audits, AI systems require continuous checks to ensure they are not drifting from their intended performance. A comprehensive evaluation capacity for AI would need to include mechanisms for regular reviews and revalidation, with the ability to dynamically adjust oversight based on the evolving risks associated with AI systems.

Conclusion

The integration of AI technologies into healthcare holds immense potential to revolutionise care delivery, enhancing quality, accessibility, efficiency, and patient outcomes. However, to harness this potential responsibly, it is essential to establish a robust regulatory framework that prioritises public safety, accountability, and alignment with healthcare values. The recommendations outlined in this report emphasise the need for a multi-faceted approach that includes sector-specific governance, lifecycle-based evaluation, dynamic regulatory reviews, and an AI-as-a-service framework.

To achieve a regulatory system that genuinely supports the safe and effective use of AI in healthcare, synergistic alignment between sector-specific and whole-of-economy AI governance is crucial. A harmonised approach ensures that healthcare-specific safety and ethical standards are maintained while benefiting from consistent national AI policies. Recognising AI as an augmentation of human services further highlights the need for AI regulation to align with the standards expected of the healthcare workforce, ensuring patient-centred and ethical care.

The adoption of comprehensive definitions of safety and risk and the inclusion of diverse stakeholder inputs are vital for addressing both individual and societal impacts of AI systems. An emphasis on sociotechnical evaluations helps to account for the broader ethical and societal implications, ensuring that AI tools are safe, fair, and inclusive. Additionally, accounting for generative AI, consumer technologies, and massive-scale software systems within the regulatory framework ensures comprehensive oversight that leaves no significant technology unaddressed.

Ultimately, the successful integration of AI into healthcare relies on building public trust through accountability, transparency, and equity. Establishing independent evaluations and a meaningful recourse mechanism for the public will further reinforce the trustworthiness of AI technologies, ensuring that they genuinely serve the interests of patients and the broader community.

By implementing these recommendations, Australia can create an AI ecosystem in healthcare that is safe, ethical, and aligned with core healthcare values. This framework will not only enhance the effectiveness of healthcare delivery but also ensure that the integration of AI strengthens, rather than undermines, the trust that forms the foundation of Australia's healthcare system. In doing so, AI can truly be leveraged to serve the public good, supporting a healthier and more equitable society.