AUSTRALIAN MEDICAL ASSOCIATION
ABN 37 008 426 793
T I 61 2 6270 5400
F I 61 2 6270 5499
E I ama@ama.com.au
W I www.ama.com.au
39 Brisbane Ave Barton ACT 2600
PO Box 6090 Kingston ACT 2604

SUBMISSION

Friday, 4 October 2024

AMA submission to the DISR – discussion paper on mandatory guardrails for safe and responsible AI

Consultation closes: 4 October 2024

By email: aiconsultation@industry.gov.au

Introduction

The AMA welcomes ongoing consultation with the Department of Industry, Science and Resources' (DISR) regarding future regulation for the safe and responsible use of Artificial Intelligence (AI) in Australia.

The AMA notes the broader government activities occurring across national, state and territory governments towards establishing a consistent regulatory framework for the application of AI across industry settings. This is an important step in responding to the many opportunities and attendant risks presented by artificial intelligence across industries.

The AMA will provide further submissions to the Department of Health and Aged Care's Safe and Responsible Artificial Intelligence in Health Care – Legislation and Regulation Review, and the Therapeutic Goods Administration's consultation, Clarifying and strengthening the regulation of Artificial Intelligence, which complement the DISR's proposed guardrails for safe and responsible AI.

Given the complexity of assessing existing regulatory structures relating to digital technologies and AI in healthcare, and the broad scope of the government's proposed measures, the AMA is disappointed the DISR's consultation has not provided more time for stakeholder feedback. The time frame is not sufficient to facilitate thorough review of proposed changes.

As the peak professional body for doctors in Australia, the AMA's submission will focus on the government's proposed mandatory guardrails and cross-industry regulatory approach in relation to the essential regulation of AI in health and medicine as high-risk contexts. This submission will not comment on voluntary standards. Reliance on self-regulation in healthcare poses an unacceptable risk to consumers and medical professionals, as technological progress in the AI space will be largely driven by free market, profit driven entities.

The AMA advocates for a comprehensive, ethically grounded regulatory framework that enhances healthcare delivery while safeguarding patient rights. We recognise a balance in regulation must be struck to support the advancement of increasingly sophisticated and secure AI capabilities that will achieve the best health and efficiency outcomes.

The healthcare tools using AI are not perfect and will continue to break the bounds of their current scope of application. It is within the nature of technology to perpetually outpace regulation. Therefore, a once-for-all fixed determination for regulation will not suffice. The AMA recognises the DISR's proposed approach is largely in line with the principles we outlined in our submission to the Safe and responsible AI in Australia Discussion Paper in 2023. In summary, regulation for healthcare and high-risk contexts must ensure:

- Clinicians are meaningfully involved in governance of AI in healthcare at every level, including human intervention points during the decision-making process.
- AI never compromises medical practitioners' clinical independence or is used by non-medical individuals to second-guess.
- Clearly established responsibility and accountability for any errors in diagnosis and treatment caused by AI products.
- Use of AI in healthcare protects the privacy and security of patient health information. Any data collected in using AI in healthcare must be subject to stringent data governance and security measures.
- Visibility of AI generated components of care and advice, treatment or diagnostic procedure to be undertaken, with full disclosure and patient consent for the use of health data.
- Data undergirding machine learning algorithms for AI in healthcare is inclusive and representative to mitigate against bias.
- AI systems are required to substantiate the information which supports decision making at any point in time.

The AMA has voiced concern Australia is lagging behind other countries in AI policy development. The DISR's proposed coordinated all-of-government response, integrated with existing regulation regarding the use of AI in Australia, is a positive step towards keeping pace with international standards and safeguards.

The AMA considers the proposed mandatory guardrails provide a starting point — the basis for an overarching national regulatory framework for using AI in high-risk contexts — so long as it is accompanied and integrated with contingent sector-specific AI regulation. These resources will equip developers and deployers of technology to be responsible users and drivers in the ongoing evolution of AI in Australia. While the government carries key responsibility for establishing and maintaining a regulatory framework, it is not the chief agent in the development of AI. This role lies with the technology developers and other stakeholders as active innovators in technology across the various spaces of industry.

The AMA views the entire health sector as a high-risk setting requiring a dedicated regulatory framework and clinical governing body. Significant existing regulation in healthcare provides a strong framework to embed new provisions for the challenges of AI in healthcare and ensure patient safety, privacy and ethical standards.

The AMA calls upon the government to implement dedicated mandatory AI standards for healthcare, and to establish a dedicated governance body consisting of practising clinicians, medical professionals, consumers and technology developers for the active oversight of AI application in healthcare.

1. Defining high-risk AI

The AMA has consistently advocated for a risk-based approach to the implementation of AI. We are pleased to see this approach is being consulted and incorporated into a national regulatory framework that will build trust and confidence in the rapidly evolving digital health solutions using AI. The AMA is broadly supportive of the multifaceted risk approach outlined in the proposals paper, which considers levels of risk and key characteristics, and balances preventative and remedial regulatory measures to mitigate against known risks.

The AMA approve of the government's proposed approach to defining high-risk AI (where the mandatory guardrails are to be applied) according to the identified principles a – f outlined in the paper. These principles consider:

- human rights
- health and safety risks
- protection to individuals from adverse legal outcomes
- prevention against further disparities among population groups
- addressing systemic impacts of AI to maintain social cohesion, protect democratic processes and uphold the rule of law
- a systemic approach to evaluating risk to ensure regulations protect vulnerable groups while allowing innovation in those lower-risk areas of application.

We acknowledge these principles have drawn from international examples such as the EU and Canada and focus on promoting interoperability – an essential component of building data integrity supporting application of AI in Australia. This approach aims to avoid misclassifying low-risk applications and is adaptable to technological advancements.

In a submission to the DISR consultation in 2023, the AMA outlined the key considerations that must be embedded in high-risk regulation for AI in healthcare. We reiterate these key points below.

- **High risk designation:** Regulating the application of AI in healthcare requires a tailored approach that allows for the adoption and integration of safe technologies while protecting practitioner and patient safety. The application of AI in healthcare must always be considered high-risk, primarily due to potential negative consequences that could arise from systemic errors, patient privacy issues, and algorithmic biases. AI associated errors can lead to significant harm, such as misdiagnosis or inappropriate treatment, with potentially irreversible impacts on patient health.

- **Health advisory regulatory body for AI:** While the Australian Government's *temporary artificial intelligence expert group* will be of value in the early stages of regulatory implementation, AI's rapid development may require a permanent national body. The EU's approach to high-risk AI includes continuous review through the mechanism of a European Artificial Intelligence Board to oversee implementation and regulation. The AMA call for the establishment of a health sector regulatory advisory board consisting of clinical and legal experts. This body may advise Commonwealth and state/territory governments and guide evolving policy for AI's consistent application in medicine and healthcare over time. These subject-matter experts must be engaged in ongoing testing and evaluation of AI products.

- **Application, tiered approach:** Regulations should not be one-size-fits-all, but adaptable to the specific risks associated with various AI applications. Governance structures tailored to different services and programs within healthcare must be developed. The AMA recommend a tiered, application approach to regulation that categorises risk levels based upon potential consequences and establishes corresponding governance structures to manage them. This structured approach is inspired by international examples such as the EU, which distinguishes risk levels associated both with the AI product itself and according to its use setting.

- **Evidence base and consultation:** Regulatory frameworks should be founded on solid evidence, incorporating insights from clinical studies, expert opinions, and data analytics. This approach helps ensure regulations are grounded in the realities of healthcare practice and patient safety. Evidence-based regulation must be maintained through ongoing collaboration with experts in AI, healthcare, ethics, and law to adapt to the evolving landscape of AI technology.

- **Human oversight and intervention:** Perhaps the most important principle the AMA has enunciated is that final clinical decisions influenced by AI must involve human oversight. High-impact AI applications require regulated, defined human intervention points during the decision-making process. Healthcare professionals must validate and take responsibility for any decisions made with the assistance of AI tools. Final decision-making cannot be a mere formality - it must reflect meaningful human judgment, ensuring that clinical context, patient values, and ethical considerations are taken into account. This principle is reflected in Guardrail 5 and must be implemented in subsequent regulatory frameworks for clinical practice.

- **Mechanism of clinical responsibility:** A principles approach to determining a high-risk application is highly compatible with healthcare's ethical foundation - practitioners' clinical responsibility. AI applications which influence critical medical decisions (e.g., diagnostic tools, treatment recommendations), receive stricter scrutiny and support the mechanism of clinical judgement tethered to Ahpra and the National Boards' oversight of clinical practice. Legislating to complement these existing mechanisms is an efficient means of supporting medical practitioners using AI to interpret their obligations to patient safety and best clinical practice.

- **Accountability and liability:** Aligning AI regulation with existing clinical standards mechanisms also ensures we can delineate responsibility for clinical errors or adverse outcomes stemming from AI applications. This is crucial for accountability, allocating liability and enabling proper avenues for redress in cases of misdiagnosis or inappropriate treatment. It must be recognised medical practitioners are not technology experts and cannot be expected to bear full knowledge of how an AI application operates. The AMA emphasises the importance of clarifying accountability and ensuring legal responsibility is demarcated clearly and appropriately between developers and deployers.

The AMA considers the DISR's proposed Option 2 – Framework approach as the most effective regulatory approach to satisfy the key considerations outlined above. We will provide further comment on this under 7. Regulatory Options for mandating guardrails.

2. Proposed introducing mandatory guardrails for AI in high-risk settings

The proposed guardrails are adequate for general application and respond to stakeholder feedback provided in ongoing consultation. The AMA considers this a firm basis for the further development of healthcare-specific regulation and calls upon the government to implement fit-for-purpose guardrails for clinical application.

As a high-risk environment, regulatory safeguards must be mandatory in healthcare. In combination with TGA regulation of medical devices and Ahpra regulation of health practitioners, robust AI standards which keep pace with AI developments must be embedded in clinical practice. They must be tailored to the healthcare context and maintained by a health-specific governing authority. Ongoing collaboration with experts in AI, healthcare, ethics, and law is critical to developing effective regulations that adapt to the evolving landscape of AI technology.

The AMA supports a regulatory structure which imposes accountability upon developers and deployers of technology to satisfy standards requirements as the surest guarantee of continued monitoring and improvement to AI technologies. We are pleased to see the DISR's proposal targets these cohorts as key drivers of AI innovation.

In the DISR's 2023 consultation, the AMA highlighted the key challenge facing the healthcare sector was the central agency of technology deployers, who can be least informed about how technology functions and how it should be governed. Patients, while end users, are appropriately excluded from being considered deployers. However, practice managers, medical practitioners and a range of professionals in the health space have vastly differing degrees of knowledge regarding AI. Regulation must provide consistent guidance and resourcing at the organisational and system level.

Standards must ensure health outcomes analysis and risk assessment associated with AI application are integrated into clinical workflows. Clinical governance should be responsible for regularly reviewing and updating the standards to account for emerging trends and new capabilities. The AMA suggests Guardrail 10 is essential to supporting developers' and deployers' joint responsibility for the application of AI within healthcare. A new conformity assessment should be required for action by deployers on a regular basis to ensure ongoing compliance with standards - perhaps annually.

The AMA notes Guardrail 7 - compensating patients harmed by AI-related errors - works in tandem with the provisions of Guardrail 6 regarding transparency. Regulatory safeguards must require developers to be transparent regarding design and ensure AI application adheres to agreed ethical principles. Healthcare providers must be educated, supported and held accountable for assessing where AI can be safely and effectively integrated into clinical workflows. These measures are appropriate to promote patient safety and engender trust in the digital health system.

3. Guardrails ensuring testing, transparency and accountability of AI

The AMA supports testing, transparency and accountability requirements captured under Guardrails 1, 4 and 6 in particular. AI systems must be subject to testing and evaluation, and this should be performed by technology providers and developers. Ongoing monitoring is also essential to ensure conformity with safety and functionality requirements, as technology will continue to evolve.

The AMA reinforces activities under Guardrail 4 in the clinical context must include ongoing interaction between model developers and upstream / downstream stakeholders, and require

ongoing consultation with clinical experts regarding use of AI applications in healthcare. Again, we emphasise a health-specific governance body would be equipped to provide this mechanism. This measure should support transparency and open-source development of new AI applications in the health space so healthcare expertise is integrated into the design process.

The engagement between clinical experts – likely to be administrators of the healthcare systems AI is integrated into – is particularly important in light of the broad definition applied to 'deployers'. The AMA is concerned the DISR definition creates significant overlap between the end user and the deployer in the guardrail framework. While DISR's approach does not directly regulate consumers, clinician users are deemed 'deployers' by virtue of using AI software in providing care to a patient. Medical practitioners carry responsibility to manage causes of potential harms to patients from AI decisions. However, they are not the most appropriate agent to manage risks posed within the operation of the technology itself. For example, when an AI application produces inaccurate diagnostics from a poor algorithm or inequitable data base.

As the regulatory body for medical devices, the TGA provides a proven and highly effective role in promoting clinical safety with the use of technology in healthcare. Furthermore, medical practitioners trust determinations of safety and efficacy made by the TGA. The TGA already regulates some AI as medical software, but not all. We support an expanded role for the TGA in AI regulation, integrated into existing TGA regulatory structures to enact the recommended tiered, application-based approach for governing the use of AI. The TGA would carry legal liability and actively participate in monitoring to guarantee safety standards. There would need to be resourcing considerations to ensure this does not impede existing regulatory roles. The AMA is preparing a submission to the TGA's current review into AI regulation.

By distinguishing between risk levels associated with an AI product and those determined by its use setting, the guardrails approach enables existing legislative instruments to apportion responsibility for risks to the actor with appropriate agency to manage it. It also utilises the necessary partnership between developers and deployers in the act of ongoing monitoring and compliance. This reciprocal relationship between 'upstream' developers and 'downstream' deployers must be clearly captured in healthcare legislative provisions to ensure the principle is established in clinical processes. It is essential our regulatory approach facilitates the ongoing innovation of AI technology and advances in its clinical application.

4. General purpose AI (GPAI)

General purpose AI (GPAI) will be comprised of increasingly broad-ranging applications that could become threatening interlopers where regulation is concerned. By the DISR's own definition, GPAI is "capable of being adapted for use, for a variety of purposes, both for direct use as well as for integration in other systems." Therefore, GPAI cannot be restricted to administrative (and generally less risky) functions in healthcare. As innovators work towards interoperability within digital health solutions, software will breach the boundaries of clinical advice and diagnostics in its pursuit of multifunctionality.

The AMA position is that a GPAI tool, when used in the clinical context, must be deemed clinical in function and therefore subject to regulation.

AI scribes provide a case example for general purpose AI. Due largely to their time-saving capabilities for GPs and other practitioners, AI scribes are becoming increasingly prevalent in the healthcare space. The supportive role they provide to clinical decision making requires they be subject to TGA regulation. This does not need to be onerous at base levels, with providers stating that they meet basic data governance standards, have appropriate insurances etc. Some monitoring to ensure compliance will be required.

The capacity to scale up regulatory requirements will be required as AI scribe accuracy improves over time and incorporate additional AI tools to streamline consultation processes, such as diagnostic supports. Risks associated with inaccuracy or expanded functions must be managed jointly through TGA regulation and the mechanism of practitioner liability for clinical decision making.

A principles method to managing AI best facilitates TGA regulatory structures. Tiered, application-based regulation, consulted by the proposed AI health advisory body, can support the TGA in determining the terms and parameters of its use according to the AI's specific function in the clinical context. The TGA's current regulatory mechanisms can differentiate GPAI tools and provide for exemptions and exclusions where appropriate. The intention is to ensure clinicians can trust the AI tools used in their practice have met basic requirements.

Regardless of their prescribed category, clinical responsibility for decisions based on scribed consultation notes remains with the medical practitioner. Principles and standards for healthcare platforms the AMA has already outlined, such as the requirement for transparency in medical software regarding AI calculation, are essential to ensure clinicians are supported to review the outputs of AI in making clinical determinations. These would have to be captured within mandatory guardrails for healthcare.

GPAI is generative AI with the built-in capability to generate content from a database, trained on an ongoing basis by scraping, analysing and processing publicly available data from the internet. When used by medical practitioners, GPAI may also collect and integrate sensitive health data even in routine, administrative functions. Privacy, security and quality considerations apply equally to GPAI as to purpose-built clinical AI systems.

5. Data governance

For both cross-industry regulation and for the health sector alone, special attention must be given to data governance (Guardrail 3). AI capabilities are built upon the imperatives of accessible, quality data, which undergirds machine learning. Robust data governance must be integrated into healthcare regulatory structures to safeguard clinicians and patients from the harms of system errors, patient privacy breaches or systemic bias.

Systemic bias and clinical integrity

Erroneous or imbalanced data will produce poor and even dangerous health outcomes. In terms of safe and responsible application in healthcare, AI regulation must anticipate and identify every instance where data may influence algorithmic or systemic bias in AI, generate erroneous reports, or produce ill-founded advice to patients. Embedded human points of intervention within the clinical process are the surest safeguard against AI-generated errors. Additionally, error reporting mechanisms must be built into processes beyond those of the TGA.

Data quality

The AMA calls for data quality assurance measures in the regulatory response to AI. Healthcare databases will experience exponential growth as AI proliferates. The efficacy of machine learning - a vital function to AI performance and integrity - depends greatly upon huge amounts of available, high-quality data in the action of mimicking data patterns. AI - itself an 'end user' in many emerging applications - will over time generate more data in both closed databases and the public domain. Technology experts are already flagging potential 'model collapse' in the future if the quality of data should decline as AI-generated information proliferates and comes to dominate databases.

Safeguards for AI use in clinical practice must consider how data quality can be maintained, possibly by ensuring raw clinical data is consistently incorporated into databases shared across healthcare services. Policy must mitigate against any reduction in the quality and diversity of AI model behaviour associated with 'regurgitative data' in healthcare applications and tools. This may require automated labelling for AI-generated content within applications to differentiate datasets used by AI and clinical decision making in the future.

Cyber security

Health data without sufficient cyber protections poses a serious risk to patients' wellbeing and healthcare services. The sensitive and imperishable nature of health information makes it a valuable commodity for software providers and a predictable target for pernicious cyber security attacks. Health data leaks are devastating to the individual and ruinous to health services and technology providers alike. Regulation must incorporate cyber security standards for all digital health technologies, to ensure AI applications do not compromise health data.

Profiling

Patient profiling is necessary and occurs through accessing health services and receiving care. The AMA recognises AI will enable healthcare services to connect and integrate broader sources of patient information (system interoperability) to construct comprehensive patient profiles that support medical professionals to evaluate patient health. Increasingly comprehensive, centralised patient data is also a more valuable commodity for data brokers and cyber hackers alike. The proposed changes to the Privacy Act currently in review do not address profiling through data sharing. The AMA emphasises this review of AI regulation must incorporate protections against tracking, targeting and profiling by data brokers, major retailers, rental platforms and data-matching firms.

Data sharing regulation

Patients are the owners of their health data. Healthcare providers, private health insurance providers and clinical software developers/operators are the custodians of patient data, not the data owners. National legislation and regulation must ensure data custodians are only permitted to share data within approved health systems and service providers. Data custodians' use of health data within AI applications must adhere to defined usage parameters appropriate to the application and management of data security risks. Regulatory frameworks must also clarify data protocols in the

instance when software ownership changes with the sale of a business, a merger, or shareholder-related restructuring.

Patient consent disclosure

Regulations must facilitate full disclosure and patient consent for the use of health data and any AI-generated health information, advice, treatment or diagnostic procedure to be undertaken when feasible.

As part of a national regulatory response to AI, the AMA recommends the Privacy Act Review adopt national-level digital data privacy and ownership protections, based upon the General Data Protection Regulation (GDPR) models of the EU and UK. Patients must be the assured owners of their health data, and transparency limitations must be imposed to determine how, when and by whom patient data can be accessed.

6. Regulatory options for mandating guardrails

In considering a regulatory approach, the proposal paper highlights the issue of regulatory inconsistency across sectors and regulatory bodies. There is an obvious need for legislative amendments that will foster effective AI risk management and promote a unified regulatory response to emerging AI challenges.

The AMA recognises many of the obligations explicit in DISR's proposed guardrails are satisfied already within workplace regulations and TGA regulations under industry-specific frameworks. Additionally, there is already a national framework within government administrated by the Department of Finance. Each state and territory jurisdiction has implemented its own complementary framework. It is not only necessary, but far more effective, that guardrails are integrated with and operate alongside existing industry laws and regulatory instruments.

The DISR cite recent regulatory initiatives that work towards this end - the case examples of TGA guidance on software-based medical devices and the eSafety Commissioner's establishment of the Designated Internet Services (DIS) standard, aimed at preventing online harms related to generative AI. These provisions require revision to ensure comprehensive regulation of AI in healthcare.

The AMA therefore support DISR's proposed Option 2 Framework approach, 'introducing framework legislation, with associated amendments to existing legislation', as a whole-of-government approach to supporting proactive, industry-specific regulation of AI to establish pre-market safeguards and risk mitigation measures that together construct a cohesive Australian regulatory environment.

The proposed Option 2 Framework approach:

- is congruent with existing regulation in various high-risk contexts across industries in Australia,
- leverages familiarity with existing health regulatory regimes, making implementation smoother for medical practices and regulators.
- promotes a consistent approach to AI reform across the economy by establishing a centralised source for AI-related concepts, aiding clarity and reducing regulatory burdens.
- will facilitate the establishment of health-specific regulatory structures for the application of AI in clinical practice.
- is most able to adapt and respond to clinical requirements, including alignment with international standards, helping Australian AI innovators integrate into global AI supply chains.

Regulatory harmonisation

The AMA recognises healthcare requires a tailored regulatory structure to deal with the high-risk environment of clinical care and the importance of harmonisation between sector-specific legislative instruments dealing with the management of AI risks. Framework legislation would best establish an enabling environment for government decision-making, with standardised regulatory terminology and powers across jurisdictions and industries.

The government's guardrails and identified high-risk settings and instructive use cases would become enforceable only when integrated into existing laws, which could be done through amendments or guidance from health sector regulators when deemed appropriate. The AMA considers this approach will support the clinical engagement required to effectively regulate AI in healthcare. Under this method, enforcement of provisions will also fall to existing regulators in the health space, using the penalties established in their respective frameworks.

Risks to approach

While a framework approach enhances consistency, implementation may be slow due to the need for coordination and consensus between health regulatory agencies and national reforms. The AMA's recommendation to establish a dedicated governance body consisting of health expertise for oversight of AI application in healthcare would facilitate this process and ensure framework legislation is integrated with the existing powers of current regulations in healthcare.

The AMA could potentially support Option 3 with regards to the establishment of an AI Act. However, this option is highly likely to create increased complexity and overlapping obligations with existing frameworks unless it is very carefully designed to include exemptions. The effort associated with identifying or establishing a new, suitable regulator to enforce the new guardrails could be deemed an unnecessary and resource-intensive endeavour when highly effective regulatory mechanisms already exist in the health sector.

Contact

president@ama.com.au

See AMA Position Statement Artificial Intelligence in Healthcare.

See AMA Submission to Automated Decision Making and AI Regulation (2022).

See AMA Position Statement on Data Governance and Patient Privacy

See AMA Position Statement System Interoperability in Healthcare.