Avant
By doctors for doctors

Avant Mutual Group Limited
ABN 58 123 154 898

Registered office
Level 6, Darling Park 3
201 Sussex Street, Sydney NSW 2000

Postal address
PO Box 746 Queen Victoria Building
Sydney NSW 1230

DX 11583 Sydney Downtown

avant.org.au

Telephone 02 9260 9000 **Fax** 02 9261 2921
Freecall 1800 128 268 **Freefax** 1800 228 268

15 October 2024

Department of Health and Aged Care

Via email **DigitalFutures@health.gov.au**

Safe and Responsible Artificial Intelligence in Health Care – Legislation and Regulation Review

Thank you for the opportunity to provide a response to the Department of Health and Aged Care's consultation on Safe and Responsible Artificial Intelligence in Health Care – Legislation and Regulation Review.

Our submission is attached. Avant would welcome the opportunity to work with the Department as part of any stakeholder working groups should they be arranged in the future.

Please contact us on the details below if you require any further information or clarification of the matters raised in the submission.

Yours sincerely

Georgie Haysom

Georgie Haysom
General Manager, Advocacy, Education and Research
Email: georgie.haysom@avant.org.au

Tracy Pickett

Tracy Pickett
Legal and Policy Adviser
Email: tracy.pickett@avant.org.au

Avant submission to the consultation on Safe and Responsible Artificial Intelligence in Health Care – Legislation and Regulation Review

Avant is a member-owned doctors' organisation and Australia's largest medical indemnity insurer, committed to supporting a sustainable health system that provides quality care to the Australian community. Avant provides professional indemnity insurance and legal advice and assistance to over 90,000 healthcare practitioners and students around Australia (more than half of Australia's doctors). Our members are from all medical specialities and career stages and from every state and territory in Australia.

We assist members in civil litigation, professional conduct matters, coronial matters and a range of other matters. Our Medico-legal Advisory Service provides support and advice to members and insured medical practices when they encounter medico-legal issues. We aim to promote quality, safety and professionalism in medical practice through advocacy, research and medico-legal education.

Over the past 18 months we have received an increasing number of requests for advice from our members about the potential medico-legal risks associated with using artificial intelligence (AI) in health care, including general purpose AI such as ChatGPT and AI scribes for clinical notetaking. A recent survey of 600 members indicated that three in four respondents had a fair to poor knowledge of AI overall. While only a small number (around 10%) were using AI, around 40% of respondents indicated they were likely to use an AI scribe in the future.

As a medical indemnity provider, our focus is on advocating for clear frameworks and guidelines to address the complexities of AI in healthcare and manage the medico-legal risks. We want to ensure that responsibility and liability are clear and properly managed and that both doctors and patients are adequately protected in the case of patient harm.

As well as this submission we have made submissions to the consultations being conducted by the Department of Industry, Science and Resources (DISR) and the Therapeutic Goods Administration (TGA), and this submission should be read together with those.

Our key points are:

- Overall, the use of AI in health care is high-risk, based on the potential for adverse impacts on an individual's physical or mental health or safety. However, there is a scale of risk for the AI products that may be used in a health care setting.

- Existing regulatory frameworks are ill-equipped to address the unique challenges posed by the use of AI in health care. There is a range of medico-legal risks that need to be considered. Current regulations may not sufficiently cover these risks.

- The regulatory settings for AI products used in health care should be commensurate with the scale of risk associated with the product and its intended use. Not all uses need to be treated in the same way. The regulation of some products, such as AI scribes and other AI currently assumed to be low-risk, should be reconsidered.

- There should be mandatory minimum standards for AI tools used in health care that fall outside the TGA's regulatory framework. This should include any AI tools that suggest clinical findings and make recommendations that if not acted upon or if inaccurate could cause adverse patient outcomes.

- There should be appropriate insurance and indemnity arrangements for developers, deployers and users of AI in the health care sector so that there is accountability and responsibility for patient safety across the AI lifecycle.

- Legislative and regulatory obligations should be placed on the entities across the AI supply chain and throughout the AI lifecycle that can most effectively prevent harms before people interact with it. If harm does occur there should be mechanisms for appropriately determining liability and obtaining redress. It is not appropriate that health care professionals bear the sole liability, particularly where they do not control the risk.

- Ongoing consultation is essential to ensure the appropriate regulation of AI in health care and should involve all relevant stakeholders including insurers.

Responses to consultation questions

What benefits will AI have for healthcare consumers and providers?

1. How can AI benefit health and care in Australia and how can we measure and deliver these benefits?

Based on our experience and interactions with members and stakeholders, we understand that AI can benefit health and care in Australia through:

- enabling more efficient use of the health workforce through analysis of data (for example bed scheduling, patient deterioration monitoring)
- improved access to health care for those members of the community who are isolated
- medication monitoring and management
- improved patient outcomes, through data analysis, early diagnosis and treatment and reduced impact of clinician fatigue and cognitive bias
- improved management of workload for clinicians due to efficiencies and productivity gains
- information management and data analysis including electronic medical record summarising.

Some ways the benefits offered by AI can be measured include:

- productivity analysis within the health system, particularly bed usage and patient care in facilities using AI and data associated with AI triage
- clinical studies relating to the implementation of AI systems in health care
- clinical audits and registries.

2. Can AI improve access to care, and what regulations could be amended or added to enable this?

No response.

Are there specific risks to using AI in healthcare?

3. What risk does AI pose to patients/consumers or health care professionals? Are the risks high or low? What criteria could be used to characterise risk? Should consumers be informed when AI is used in these low-risk ways?

Level of risk

In general, the risk that AI poses for patients and health care professionals is high due to the potential for adverse impacts on a patient's physical or mental health or safety. We agree with the characterisation in the DISR consultation that the use of AI in health care is high risk.

However, there is a scale of risk for AI that may be used in health care. The regulatory settings for AI products used in health care should be commensurate with the scale of risk associated with the product and its intended use. The severity of the risk of harm which may eventuate should determine the level of regulation required. We explain our position further in response to question 6.

Categorising risk

Key medico-legal risks associated with using AI in health care and some examples include:

- Risks of inaccuracies and errors:
    - AI prioritising tools may contain unforeseeable errors due to hidden bias within the AI system (for example prioritising one patient group over another which is also a risk for the patient).
    - AI scribes may misunderstand terms used during a consultation and substitute incorrect terms (which the tool may have been trained on). If the error is not identified by the health care professional this could lead to misdiagnosis.

  - AI scribes will not capture non-verbal cues and examination findings unless they are verbalised by the health care professional.
  - Some AI scribes suggest which MBS item numbers to bill without ensuring that all the requirements for those item numbers are met and recorded in the medical record. This could lead to financial risk for the health care professional if they accept these recommendations and bill suggested MBS item numbers without ensuring they have met the MBS billing requirements. If an audit is later undertaken, it would not be a defence to suggest that the AI told the practitioner to bill the item numbers.
  - There may be risks of diagnostic error because the AI tool has been trained on inappropriate population data or contains hidden biases.

- Risks associated with automation bias:
  - The health care professional can rely too heavily on an AI tool, thereby delegating responsibility to it rather than remaining vigilant and in doing so overlook an aspect of patient care. For example, AI scribes can appear to be logical and comprehensive and yet may be inaccurate. Errors can be repeated in subsequent documents based on initial notes.
  - There may be risks of misdiagnosis with AI scanning tools due to over-reliance on the AI and by the health care professional not applying their own clinical reasoning.

- Risks associated with loss of clinical skills:
  - If there is over-reliance on AI diagnostic and management tools, there is a risk of reduced clinical skills for the health care professionals who use them. This may have impacts for clinical care if the AI system is inaccessible.
  - These risks are amplified in instances where AI is over-trusted or where underlying decision-making processes are not well understood.

- Risks associated with the “black box” nature of some AI tools:
  - These risks are associated with not understanding how the AI works and not being able to evaluate its outputs as a result. On example is “black box” predictive AI tools where there is no way to check whether the basis of the prediction is correct because the output is based on statistical associations rather than accessible deterministic programming.
  - AI systems can “evolve” over time in a manner that is not foreseeable or intended by the developer or deployer. There are significant risks for health care professionals who seek to rely on the output of the AI in these circumstances particularly if they are not aware that this is a possibility.

There are also significant privacy risks around how patient information is collected, used and disclosed by AI tools.

As well as these key medico-legal risks, there is legal uncertainty around accountability, responsibility and liability across the AI lifecycle. It can be difficult to determine who is

accountable when AI-driven decisions result in errors or adverse outcomes. The complexity of AI algorithms and machine learning makes it challenging to trace the decision-making process, creating medico-legal risks and complicating the implementation of risk management processes and the assignment of liability. This could ultimately result in inadequate outcomes for consumers of health care if they are unable to obtain redress in the case of harm.

Health care professionals may be deployers of AI and end users (as defined in the DISR consultation paper) and therefore may have multiple levels of responsibility, the limits of which are unclear.

Depending on the type of AI scribe used and the location of the consultation, there may be risks for health care professionals associated with breach of state and territory surveillance devices legislation. If an AI scribe "listens" to a consultation and the patient has not provided consent, in some jurisdictions, this may amount to a breach of the legislation.

Informing consumers about the use of AI

Consumers should be informed about the use of AI in all health care settings. The responsibility for doing so should not rest with health care professionals alone as they may not be aware that AI is being used or how it is being used. It should be the obligation of developers to provide this information to deployers and end-users.

4. What factors are important for rural and regional Australia when assessing the benefits, risks, and safety of AI? Are there other communities that face specific risks when implementing AI-driven health care? What considerations should be made to ensure all Australians have access to the benefits of AI?

From our communications with members and stakeholders, we understand that the important factors for rural and regional Australia when assessing the benefits, risks and safety of AI include:

- providing safe, appropriate and clinically effective health care to as many people in rural and regional Australia as possible
- using appropriate and population specific data sets to inform the AI
- providing health care which is not affected by human factors such as fatigue or cognitive bias and does not require extensive travel, by either the patient or the health care professional, to obtain or provide the care
- having an appropriately qualified and experienced human in the loop to check the outcomes of the AI tool
- ensuring that the outputs of the AI tool have not "evolved".

Marginalised and isolated communities should not be further adversely impacted by the use of AI which cannot be checked to ensure that its outputs are appropriate for that

population, for example, AI trained on overseas data which is not appropriate to the recipient population group in Australia.

5. Should health care professionals have a choice about whether they use AI as part of their work?

Whether or not health care professionals should have a choice about using AI as part of their work will depend on several factors including the type of AI, whether its use is supported by evidence, whether it has been subject to regulatory oversight, and the level of risk involved.

From our experience assisting members, for doctors to trust the use of AI in healthcare they need:

- confidence that the AI does what it purports to do, is reliable, explainable and secure, and will protect their patient's data
- AI systems that are evidence based, clinically relevant and have been subject to appropriate approval or accreditation protocols
- clarity around liability, insurance and indemnity, particularly if something goes wrong.

As the use of AI is in a transitional phase, we agree that promotion of AI literacy among health care professionals would be beneficial (consultation paper page 6).

What are the possible regulatory changes?

6. What unique considerations are specific to AI in health care, and why? Should the government address them through regulatory change?

There are several considerations specific to health care, primarily around managing risk and we have articulated these in answer to question 3.

Avant supports a risk-based approach to regulation of AI. As stated in response to question 3, we agree that generally the use of AI in health care is high risk, based on the potential for adverse impacts on an individual's physical or mental health or safety. We also broadly support the introduction of guardrails as outlined in the DISR consultation and agree that they should be mandatory in high-risk settings including health care.

Importantly however, there is a scale of risk for the AI products that may be used in a health care setting. Uses range from providing and collecting information, triage and prioritisation, document summarisation, prediction and clinical decision support, to diagnosis, treatment, prognosis and ongoing clinical management.

Existing regulatory frameworks are ill-equipped to address the unique challenges posed by the use of AI in health care. Current regulations may not sufficiently cover the medico-legal risks.

Our view is that the regulatory settings for AI products used in health care should be commensurate with the scale of risk associated with the product and its intended use. Not all uses need to be treated in the same way. The regulation of some products, such as AI scribes and other AI currently assumed to be low-risk, should be reconsidered.

Where the definition of “medical device” under the Therapeutic Goods Act 1989 is met, a product should fall under the TGA’s regulatory regime. This includes where a product incorporates prediction, clinical decision support, diagnosis, treatment, prognosis and ongoing clinical management.

We are concerned, however, about AI and AI-enabled products used in health care that are currently subject to no regulation if they do not meet the definition of “medical device” but still represent a higher level of risk for users (including consumers and clinicians). An example is AI scribes.

We believe there should be mandatory minimum standards for AI tools used in health care that fall outside the TGA’s regulatory framework. As well as AI scribes, this should include consumer health products, digital mental health tools and any AI tools that suggest clinical findings and make recommendations that if not acted upon or if inaccurate could cause adverse patient outcomes.

This would include minimum regulatory standards for:

- privacy and security
- transparency and explainability around how the tool works and how it has been trained
- record keeping, access to meta data and access to historical information for relevant time periods
- service agreements and the extent to which disclaimers and indemnity clauses can or cannot be used
- controls in place to reduce the risk of patient harm, including monitoring and error and adverse event reporting
- insurance and indemnity cover.

Standards development, accreditation and education could be within the remit of the Australian Commission on Safety and Quality Health Care, the Australian Digital Health Authority or a newly created independent AI regulatory body (see response question 8). Having an “approved” mark of assurance or certification from the TGA (or other appropriate body) would provide reassurance for medical practitioners and their patients of regulatory compliance.

7. How does the use of AI differ in healthcare settings compared to general or other sectors such as finance, education, etc.?

As stated in response to question 3, the use of AI differs in health care settings compared to other sectors because of the potential for adverse impacts on an individual's physical, or mental health or safety.

8. Should there be an Australian body specifically dedicated to overseeing AI in health care? If so, how would this body differ from a broader organisation like the National AI Centre?

We consider that that there should be an Australian body dedicated to overseeing the use of AI in health care. This role could be undertaken by an existing body, such as the Australian Commission on Safety and Quality in Health Care (ACSQHC) or the Australian Digital Health Agency (ADHA), or by a newly created, independent AI regulatory body (as outlined in the consultation paper page 7). This would need to be coordinated with the current role of other regulators such as the TGA.

9. Are there any specific changes to existing healthcare laws that would address AI-related harms or help AI to be used safely?

We broadly support DISR's proposal for mandatory guardrails for high-risk AI. The adoption of mandatory guardrails in the context of existing healthcare laws would assist to ensure the safe use of AI in health care.

We support several of the changes to existing TGA regulations in our submission to the TGA's consultation.

As noted in answer to question 6, we consider there should be mandatory minimum regulatory standards for AI tools used in health care that fall outside the TGA's regulatory framework. As well as AI scribes, this should include consumer health products, digital mental health tools and any AI tools that suggest clinical findings and make recommendations that if not acted upon or if inaccurate could cause adverse patient outcomes.

AI regulation should cover liability when there is misuse or adverse events when using AI. Legislative and regulatory obligations should be placed on the entities across the AI supply chain and throughout the AI lifecycle who can most effectively prevent harms before people interact with it. If harm does occur, there should be mechanisms for appropriately determining liability and obtaining redress.

We are concerned about the risk that responsibility for the design or function of AI could be unfairly shifted onto health care professionals. This has been exacerbated by the presence of broad indemnity clauses in some AI provider contracts, which attempt to absolve those companies of responsibility and liability and place the burden on the health

care professionals using the AI system when they are not responsible for controlling the risk. Professional indemnity insurance policies do not cover product liability.

One example of current legislative complexity relates to the application of state and territory surveillance devices legislation to AI scribes. As AI scribes "listen in" to consultations, and depending on the state or territory where the doctor is located, AI scribes may be caught by the definitions in the legislation and impose obligations on doctors to obtain consent before the AI scribe is used. Because of this we have been advising our members to obtain consent at each consultation, even though surveillance devices legislation is not intended to cover a situation like this. Although this is within the legislative power of the states and territories, it would be helpful for this legislation to be reviewed and for AI scribes in healthcare to be exempt from this legislation.

10. Which international approaches should we consider, if any, that are specific to health care?

No response.

Is AI always right?

11. Should humans be able to overrule a finding or decision made by AI?

Appropriately qualified and experienced humans should be able to overrule a finding or decision made by AI if required. That is, we consider that AI should assist rather than replace human input in health care.

12. Should there always be a person or "human in the loop" to make decisions or deliver a health care service? Are there any circumstances in which it would be acceptable to have fully automated health or care decisions made by an AI product?

We consider that generally AI should assist rather than replace humans in health care.

Ultimately, there are functions that humans perform that can't be replaced by AI. Humans are able to consider factors on various levels, including visualising the patient, and reach a conclusion based on all relevant information.

As technology develops, it may be acceptable to have fully automated health or care decisions made by an AI product where there is reliable evidence that the AI outputs are as good as or better than those of a human, and where the AI has been through a rigorous independent quality and safety assessment and assurance program. However, in that context there would need to be review and continual monitoring of outputs and reporting of unexpected and adverse events to ensure the outputs continued to meet the required standards and intended purpose.

13. Should errors made by AI be reported? If yes, how should they be reported?

Yes, errors made by AI should be reported. Failure to require this may disincentivise the safe use of AI in health care. While we acknowledge that the level of reporting needs to be carefully managed to prevent notification fatigue, errors where AI operates in a manner which was not intended (that is, the AI operates erroneously or in a way which is unexpected) should be reported, reviewed and actioned as appropriate.

Continued reporting and monitoring of the operation of AI is essential to ensure that operation it functions in the manner intended.

Errors by an AI system that is a medical device and subject to regulation by the TGA should be reported through the TGA's adverse event reporting for therapeutic goods.

Errors by an AI system that is not a medical device should be reported to the developer and the developer be required to assess and monitor their system as part of the minimum regulatory standards outlined above in answer to question 6.

How should healthcare data and transparency be managed for AI use?

14. Should there be transparency about when AI is involved in health care, and should consent be requested from the consumer or health care professional?

For AI to be effective and widely accepted in health care, particularly in this transitional stage, both doctors and patients should understand when and why it is being used.

Avant's general position is that doctors need to understand how the AI they are using works and should inform patients about the use of AI, supported by information provided by developers, as outlined in our response to question 4. We note that in some circumstances clinical devices may currently use AI (for example within medical software) but doctors and patients are not aware that this is the case.

Obtaining consumer or health care professional consent to using AI in every instance is not practical, but at a minimum, consumers and health care professionals should be informed that AI is being used. Labelling of AI-based products would be of assistance in this regard.

15. Generative AI may be developed for general use, yet used in health care. Should generative AI developed have any special treatment, regulatory or otherwise?

Generative AI that has been developed for general use can be problematic when used in healthcare, as noted in the consultation paper (page 9).

As noted in our submission to the DISR consultation, the mandatory guardrails recommended in the DISR paper should apply to all general purpose AI models because there is the potential for these models to be used for a wide range of purposes and the possible applications and risks cannot be foreseen.

16. What protections are needed for patient data used or generated by AI that are different for health care?

We agree with the concerns noted in the consultation paper (page 9) about AI produced or operated by overseas organisations, namely:

- Data on which the AI is based is from other countries and may be biased and less accurate when used in the Australian context.
- Data may be sent and stored overseas and may be used to develop new products.
- It may not be clear that AI is part of a website or software.

Health care professionals and consumers may not be aware of any of these. The same protections that apply to identifiable patient data under privacy legislation should be applied to identifiable patient data used or generated by AI. In our experience, some doctors have been using AI in healthcare without consideration of the privacy issues, while others have resisted using AI because they cannot confirm it complies with Australian privacy law.

We understand that sometimes doctors use redacted health information in general purpose AI models thinking that it is safe to do so from a privacy perspective, without understanding that redaction is not necessarily the same as de-identification. Current privacy obligations apply to identifiable and re-identifiable health information and further guidance from the OAIC would assist doctors and others to comply with their privacy obligations when using AI.

In the context of AI scribes, many developers state that their product is compliant with Australian privacy laws, but it can be difficult for doctors to obtain information to validate this. We have suggested in answer to question 6 that there should be minimum standards for privacy and security to ensure that there are protections for patient data.

17. Is it acceptable for developers of AI products to use patient data to develop their products or to sell patient data collected from use of AI?

It is not acceptable for developers of AI products to use identifiable patient data to develop products or sell identifiable patient data collected from the use of AI without specific patient consent to do so, as required under current privacy legislation.

Some analogous situations where there have been legal and ethical concerns about lack of transparency and secondary use are:

- the use of cell lines from tissue samples in commercial product development
- the use of genetic information provided for healthcare being used for other purposes, for example criminal profiling.

Usually, the way these challenges are managed is through obtaining specific patient consent for the intended uses. However, the potential uses of patient data to develop AI products and sale of patient data are outside the knowledge and contemplation of the individual medical practitioner or practice. Therefore, the practitioner or practice cannot obtain effective consent for those uses and should not be responsible for doing so.

Data that has been de-identified and aggregated is not protected from use under current privacy legislation. At a minimum, developers of AI products using data to train and refine models and for commercial purposes should be required to disclose this in advance to users of their products.

18. Should your healthcare information be kept in Australia? If yes, would your view change if this reduced ability to access advances in AI made overseas?

Where possible, healthcare information used by AI should remain in Australia to ensure that Australian privacy protections apply.

If the use of AI requires identifiable patient information to be sent overseas, there are already obligations under Australian Privacy Principle 8 that must be complied with.

If overseas advances in AI mean Australians would have reduced ability to access innovative technology, patients should be informed and agree before their identifiable information is sent overseas due to the unforeseeable risks of data loss and other privacy risks.

19. Are there any specific safety considerations that have not been raised elsewhere?

No response.

Avant Mutual
15 October 2024