

Australian Alliance for Artificial Intelligence in Healthcare

Response to Australian Government Department of Health and Aged Care
Consultation on Safe and Responsible AI in Health Care

October 2024

Introduction

The Australian Alliance for AI in Healthcare (AAAiH) is the peak organisation representing the Healthcare AI sector. It was established in 2018, and its more than 100 member organisations are drawn from academia, government, consumer, clinical, industry organisations, and peak bodies.

AAAiH wishes to see a **fully funded national plan by 2025** designed to create an AI-enabled Australian healthcare system capable of delivering personalised healthcare, safely, ethically and sustainably supported by a vibrant AI industry sector that creates jobs and exports to the world, alongside an AI-aware workforce and AI-savvy consumers.

Our clear view is that AI in healthcare has many unique elements, and these will require healthcare-specific responses, beyond generic cross-industry policy. AI cannot be treated as an isolated technology, but to work, must be embedded within healthcare workflow, policy, culture, workforce, technology and practices. Any policy response to the arrival of AI will be unable to separate these unique elements of healthcare context from AI itself. Healthcare has been regulating AI in one form or another for some time, and the Therapeutic Goods Administration (TGA) has deep experience in regulating AI via the Software as a medical Device (SaMD) Framework – making healthcare an early leader in the regulation of AI, with much to share across other sectors.

AAAiH welcomes the current set of consultations on AI being run by the Department of Health and Aged Care (DoHAC), the TGA, the Department of Industry, Science and Resources (DISR), and the Senate Select Committee on Adopting AI.

This document represents our response to the consultation by DoHAC.

To facilitate engagement between government and the healthcare sector around matters relating to AI, AAAiH has undertaken a number of activities over the last 12 months:

- In late 2023 the AAAiH launched its AI National Policy Roadmap[1] after a series of extensive consultations with its membership and the broader healthcare community, making 16 recommendations. We have been pleased to be able to discuss these recommendations with DoHAC on several occasions over the last 12 months.
- In July of this year, Professors Enrico Coiera and Farah Magrabi were invited to address the Senate Select Committee, and were able to table the AAAiH National Policy Roadmap, and answer questions from Committee members related to the recommendations, with specific focus by Senators on AI Safety and the new challenges of Generative AI.
- AAAiH was pleased to assist in facilitating the DoHAC current consultation process by co-ordinating and running a series of three webinars reaching approximately 195 individuals over the last month. AAAiH looks forward to DoHAC receiving individual submissions from its membership as a result.

AAAiH has not undertaken a new community conversation to address the detailed questions in this DoHAC consultation. Consequently, our late 2003 AI National Policy Roadmap remains the best indication of the community consensus amongst Alliance members. For this response we thus will link the AI Roadmap recommendations to the consultation questions, identify progress over the last 12 months, and draw attention to questions where we think particular attention must be paid. We additionally draw attention to critical Roadmap recommendations, which are not obviously explored in the specific questions asked by DoHAC.

In Table 1, we align consultation questions with specific recommendations in the AI National Policy Roadmap.

Critical issues in response to the consultation questions – Generative AI Safety (Questions 3, 15, 19)

We wish to draw specific attention to our response to questions on the safety and regulatory treatment of Generative AI (GenAI).

Regulatory approach to GenAI – Pre and Post Market mechanisms

Generative AI is substantially different to other recent forms of AI such as deep learning created predictive algorithms, and these differences have safety and regulatory consequences:[2]

- GenAI outputs are **not deterministic** but probabilistic. The same input can produce different outputs derived from the same underlying probability distributions captured within the model.
- Real-world GenAIs can be **dynamic**, changing their behaviour as they are trained on new data or have their architecture or inference procedures updated or recalibrated. Such changes may not be announced.
- Because GenAIs are **general purpose**, they can be used for a wide variety of often unanticipated tasks. However, performance on predefined benchmark tasks cannot predict real-world performance for other "off-label" (non-benchmark) uses. Small variations in prompts may yield large variations in performance.
- Large-scale GenAIs like large language models (LLMs) are typically kept **opaque to reverse engineering** because model owners gain commercial benefit from not sharing technical information. Without critical structural information, it is hard to predict how model behaviours may vary across different tasks.
- Even if there were commercial transparency, GenAIs are **not fully describable** in a way that allows behaviours to be inferred from first principles.

GenAI evaluations become snapshots of one unique moment in a model's life and one small part of their addressable task space. This means that pre-market assessment is not sufficient, given that we expect changes in models, software, or user behaviour to significantly alter performance after a GenAI comes to market.

The alternative is to adopt a strategy of **rigorous post-market surveillance**, where performance is monitored continuously for unexpected errors or changes in performance. Post-market surveillance seems essential if GenAIs are used in mission-critical industries like healthcare, where changes in performance can lead to patient harm or death. The form that such surveillance takes will require careful consideration but will need to be substantive enough to mitigate GenAI use in high-risk settings.

Should GenAI digital scribes be regarded as medical devices or treated separately outside the TGA's remit?

Currently AI scribes that solely support clinical documentation are excluded from medical device regulation by the TGA, and the question arises whether they require separate treatment e.g. by DoHAC.

We believe that the **TGA remains the preferred agency for dealing with AI scribes**, given that relatively minor changes in scribe function can turn them into a form of decision support capable of making explicit recommendations to clinicians around clinical matters such as diagnosis and treatment. If the TGA had oversight of scribes across the continuum of their design and function, they would be best placed to flag to manufacturers when their software moves from one regulatory treatment to another. Separate agencies would impose a bureaucratic barrier to rapid recategorization.

Next, we note the **difference between scribes that transcribe speech verbatim and those that summarise**. Transcription systems create a direct record of every spoken word without editorialization, whilst summarisation systems create compressed variants of what was said, sometimes enriched with data from the electronic record. The act of including or excluding items, or otherwise rephrasing information, into a clinical document such as a note or letter, requires computational reasoning informed by clinical knowledge. Decisions are being made about what to say and how it should be said. The purpose of documentation is to provide information that shapes human decisions. Consequently, we have expressed the view in the past that the **generation of clinical documentation is a form of decision support**.

We know that communication errors are a common cause of a patient harm and a frequent component of medicolegal cases, making communication a high-risk element of care. We thus believe that it makes sense to **treat software for the automated generation of clinical documentation as a medical device**.

Critical Issues not covered in the consultation questions

The AI National Policy Roadmap made a number of recommendations that are not directly exposed through the questions asked in this consultation.

National coordination of government agencies

The AI Roadmap's first recommendation was for the establishment of a National AI in Healthcare Council to better coordinate and harmonise the responsibilities and activities of those entities responsible for oversight of AI safety, effectiveness, and ethical and security risks. We note the clear evidence of cross-portfolio and cross-agency collaboration federally in the AI sphere, and welcome this. Without formalisation of such collaborative structures however, there is concern that in future years we will see fragmentation of response. Consequently, we continue to believe that formalisation of collaboration through a Council remains a priority.

Industry

The DoHAC consultation paper does not specifically discuss the local AI industry, no doubt reflecting the separate consultation undertaken by the Department of Industry, Science and Resources. However, there are healthcare specific issues that do need addressing, given the need for sovereign capability to implement and operate AI. We consider the AI Roadmap recommendations for industry to be important i.e.

- the need for national clinical AI procurement guidelines;
- expanding the R&D tax incentive scheme to cover TGA regulatory compliance costs;
- finding mechanisms to provide safe and privacy preserving access to data for model training;
- the development of a National AI Capability Centre in Healthcare to assist industry (and SMEs in particular) to bring products to market, and to assist healthcare services and organisations in the safe and effective adoption of AI.

Research

The AI Roadmap identified the critical role that research plays in ensuring that the development and deployment of AI in healthcare is based on the most up to date evidence, and that Australia retains world-class sovereign capability to use AI to support industry and in the national interest. At present the level of government support for research in healthcare AI is minimal when compared to that in other comparable nations and puts Australian Research at a substantial disadvantage.

The breakthrough science of the next two decades will be built upon two technology platforms – genomics and AI. Australia took the bold decision to create a national mission in genomics, ringfencing substantial funding to feed multiple research and translational activities, and the success of this investment has now seen the embedding of Australian Genomics within DoHAC.

We believe that **the time has come for a national mission in healthcare AI**. Such a bold and nation building investment would be shaped by the lessons of the Genomics mission, and would secure the future of Australia's healthcare system.

Conclusions

AAAiH is grateful for the opportunity to contribute to the national discussion surrounding the safe and effective adoption of AI in healthcare. We are happy to answer questions on any matters arising from this submission or provide additional support as needed.

References

1. Coiera E, Magrabi F, Hansen D, Verspoor K. A National Policy Roadmap for Artificial Intelligence in Healthcare. Australian Alliance for Artificial Intelligence in Healthcare, 2023. (https://aihealthalliance.org/2023/11/16/ai-can-revolutionise-healthcare-but-only-with-a-national-plan/).
2. Coiera E, Fraile-Navarro D. AI as an Ecosystem — Ensuring Generative AI Is Safe and Effective. NEJM AI 2024;1(9):AIp2400611. DOI: doi:10.1056/AIp2400611.