Consumers Health Forum of Australia

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

SUBMISSION: November 2024

Consumers Health Forum of Australia (2024)
Submission to the *Safe and Responsible Artificial Intelligence in Health Care – Legislation and Regulation Review.* Canberra, Australia

P: 02 6273 5444
E: info@chf.org.au

twitter.com/CHFofAustralia
facebook.com/CHFofAustralia

Office Address
7B/17 Napier Close,
Deakin ACT 2600

Postal Address
PO Box 73
Deakin West ACT 2600

Consumers Health Forum of Australia is funded by the Australian Government as the peak healthcare consumer organisation under the Health Peak and Advisory Bodies Programme

CONTENTS

Contents

Introduction 4

Key Recommendations 4

CHF's advocacy in automated decision making and artificial intelligence 5

Interim findings from the Australia's Health Panel AI survey 6

AI Governance 14

Safety and Quality risks of AI 15

Data, Privacy and Ethics 18

Consumer Engagement 19

References 21

Introduction

Consumers Health Forum (CHF) is the national peak body representing the interests of Australian healthcare consumers and those interested in healthcare consumer affairs. CHF works to achieve safe, quality, and timely healthcare for all Australians, supported by accessible health information and systems. At the heart of CHF's policy agenda is consumer-centred care.

In the context of CHF's commitment to quality improvement in healthcare, CHF supports and recognises the value of automated decision-making and discriminative Artificial Intelligence (AI) technologies, which have been used in healthcare for many decades.

However, the sudden rise of generative word modelling, such as ChatGPT, prompts CHF to raise concerns, as many features of generative AI software may pose significant risks to consumers when used in a healthcare setting. CHF welcomes the opportunity to provide a submission to the Department of Health and Aged Care, "Safe and Responsible Artificial Intelligence in Health Care – Legislation and Regulation Review".

While healthcare consumers, healthcare services, and healthcare financing institutions could all potentially benefit from the broader adoption of generative AI, the risks are primarily borne by consumers.

In this submission, CHF will present interim findings from Australia's Health Panel AI survey, performed in June 2024. We will then discuss the risks that the unregulated adoption of generative AI would have in healthcare. CHF will also provide recommendations on mitigating such risks while implementing legislative and regulatory reforms.

Key Recommendations

1. Healthcare consumers must be involved in all levels of governance and the implementation of legislative reforms for generative AI in Australia. So far, the formative process of reform has focused on consulting and capturing the ideas of the technical experts that design, deploy, and research AI – in other words, those that are most likely to benefit from broader applications of AI. CHF recognises that their views and input are fundamental but argues there has not been sufficient consultation with consumers.
2. Government to implement risk-based legislation in which healthcare applications of AI are classified as high-risk
3. Healthcare safety and quality regulatory institutions must work cooperatively with consumers, clinicians, and AI developers to address critical regulatory grey areas.

4. Requirements of AI healthcare software must include the assessment of the employed algorithms against minimum standards of safety and quality.
5. AI algorithms must be vetted against bias. This requirement is a fundamental step to prevent algorithms from exacerbating existing discriminatory practices towards marginalised groups.
6. The government should leverage data and consumer protection legislation to engage with consumers, clinicians, and developers in creating a solid data governance framework for AI software in healthcare.
7. Legislation must protect consumers against data being used for perverse profiteering.
8. Resources must be developed to increase AI literacy in the Australian population, particularly its applications within healthcare provision.
9. The integration of AI in healthcare must happen under the stewardship of the Australian Commission on Safety and Quality in Healthcare. The National Safety and Quality Health Service Standards must be amended to explicitly address generative AI use and its implications.

CHF's advocacy in automated decision making and artificial intelligence

CHF has been involved in advocating for consumers regarding Automated Decision Making (ADM) and artificial intelligence for several years, including:

- Advocating for including consumer perspectives in the safety and quality of Electronic Clinical Decision Making Support (ECDS) software in February 2022[1]. Our issue paper - developed through extensive consumer consultation - highlighted the importance of governance and oversight of AI at the federal government level. This ensures that guidelines related to AI (safety, quality, privacy, and security) are developed at a national level.
- Organising webinars to educate and collect feedback from consumers on ADM and AI, such as the June 2022 "Automated Decision Making and Artificial Intelligence in Health"[2] webinar, which focused on the importance for consumers to have solid data guardianship frameworks when their data is collected and utilised by AI software.

[1] "GP Data and Electronic Clinical Decision Support | Consumers Health Forum of Australia." 2023. Chf.org.au. June 15, 2023. https://chf.org.au/publications/gp-data-and-electronic-clinical-decision-support.

[2] "Webinar: Automated Decision Making and Artificial Intelligence in Health | Consumers Health Forum of Australia." 2022b. Chf.org.au. June 2, 2022. https://chf.org.au/events/webinar-automated-decision-making-and-artificial-intelligence-health

- Since 2019, CHF has been a member of the Australian Alliance for Artificial Intelligence in Healthcare (AAAiH).[3] The alliance presents several priorities for developing regulatory frameworks for AI in healthcare. In October 2023, AAAiH released a new updated roadmap, calling for healthcare consumers to be engaged in co-designing AI healthcare services and systems. The roadmap also calls for developing consumer-focused AI healthcare literacy guidelines and resources.
- Two of CHF's Special Interest Groups (SIGs), the Digital Health SIG and Safety and Quality SIG, have provided various policy input regarding AI over the years, focusing on the importance of data safety, guardianship and the need for a legislated framework.
- CHF Australia's Health Panel (AHP) is an interactive platform dedicated to collecting Australians' views about the state of the nation's healthcare system. In June 2024, CHF undertook an online survey of AI in Health Care with members of CHF's opt-in Australia's Health Panel. Topics included trust and confidence in generative AI software aiding clinical decision-making, concerns with using generative AI in healthcare, and assessing current familiarity with generative AI software.

Interim findings from the Australia's Health Panel AI survey

In June 2024, the Consumers Health Forum of Australia (CHF) undertook an online survey on AI in Health Care with members of CHF's opt-in Australia's Health Panel. Preliminary results from CHF's upcoming report *"What Australia's Health Panel have to say about AI in Health Care?"* are summarised below.

In total, 136 responses were received, of which 61.0% came from females, and most respondents were aged over 45 years. There were respondents from all Australian States and Territories. The majority of respondents (87.3%) had a long-term health condition, and 39.6% of respondents reported being in "fair or poor" health at the time of the survey.

Our survey found that:

[3] Dorricott, Pippa. n.d. "AI Can Revolutionise Healthcare but Only with a National Plan | Ai Health Alliance." https://aihealthalliance.org/2023/11/16/ai-can-revolutionise-healthcare-but-only-with-a-national-plan/.

There is a need for tools that allow consumers to better understand the benefits and risks of AI in healthcare. Respondents often stated that they currently have little knowledge of AI and the extent to which it is being used in their health care. Only about one in seven (13.6%) respondents regarded their knowledge of how AI is used in health care as excellent or very good, 59.2% stated it was good or fair, whilst one in four (24.0%) regarded their knowledge of this topic as poor. If AI is used more in the future, greater education of the potential benefits and risks of AI might be required to support consumers to make decisions about their health.

Consumers must always be informed if AI is being used to diagnose or treat their condition. Three out of four respondents (75.0%) felt it was very important to them to be told if "AI has been used in your diagnosis or treatment", a further 17.7% stated it was "somewhat important" to be told. Only around 1 in 20 (4.8%) respondents indicated that it was "not important" to be told.

Consumers see the human element of healthcare as fundamental. Consumers don't want AI to make decisions about their health care without thorough clinician supervision and output validation. Respondents appeared to appreciate the potential convenience of AI in health care, for example when scheduling appointments or improving their access to advice outside of regular business hours. However, almost a quarter of survey respondents said they do not trust AI alone to decide what their diagnosis or treatment should be. The majority of others (68.5 %) only trusted AI to be involved in making decisions about their healthcare if it was verified by a human. When asked "Do you trust AI to be involved in making decisions about your healthcare?" only 2.4% responded with "Yes", while the most common response was "Yes, only if verified by a human" (68.5%). Almost 1 in 4 respondents (23.4%) answered "No". It is also relatively common to have some reservations about how AI might erode their interactions human health care staff.

The extent by which AI use is allowed in healthcare must be transparent and communicated to consumers. AI tools are characterised by algorithm opacity, which describes the inherent inability of humans to comprehend how an algorithm produces a decision. This is due to the complexity of machine learning. Even having access to the code that generated an AI decision, the code would be so long and complex that no human would not be able to validate it. This creates a lack of transparency that is not easily addressed, leading to mistrust and discomfort in consumers. One third of respondents (32.3%) stated they have no idea of how often AI is being used to support their health care at present. Additionally, more than 9 in 10 respondents indicated their beliefs that it is important for there to be appropriate monitoring and public accountability about the impacts of AI has on health care in Australia in the future.

Figure 1: Consumer comfort towards the use of AI in Health Care (%) n=128-129

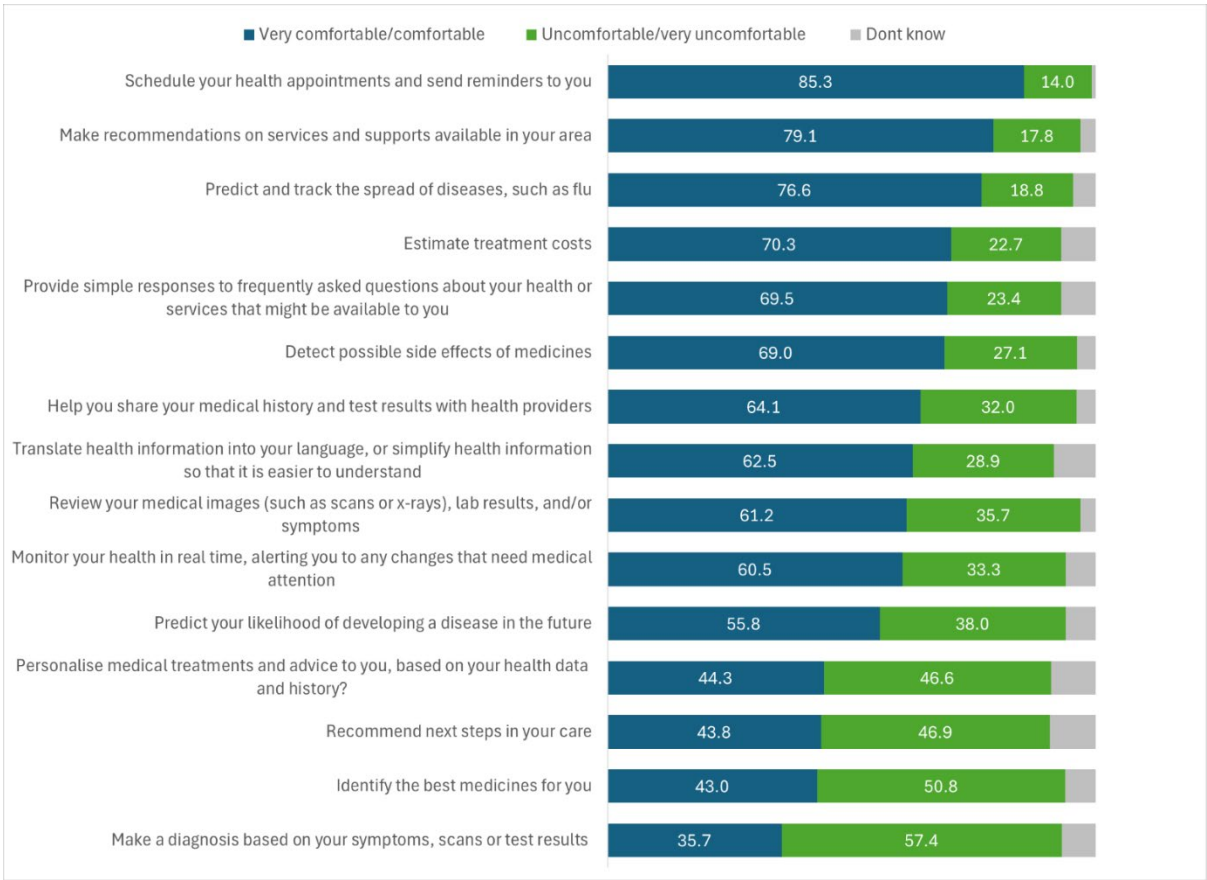

Figure 1 presents data on respondents' level of comfort with AI being used in various aspects of health care.

The majority of respondents (85.3%) reported being "comfortable or very comfortable" with AI being used to schedule their health appointments and to send reminders. Similarly, most respondents were comfortable with AI being used to make recommendations about services and supports in their area (79.1%), to predict the spread of diseases such as flu (76.6%) or to estimate treatment costs (70.3%).

Respondents appeared to be less comfortable with AI being used to personalise their own future care. Just 35.7% of respondents stated they would be "comfortable or very comfortable" with AI being used to make a diagnosis based on their symptoms, scans or test results; a minority (43.0%) were "comfortable or very comfortable" with AI being used to identify the best medicines for them, to recommend the next steps in their care (43.8%) or with AI being used to personalise their medical treatments and advice (44.3%)

Figure 2: Consumer trust in AI in Health Care (%) n=121-124

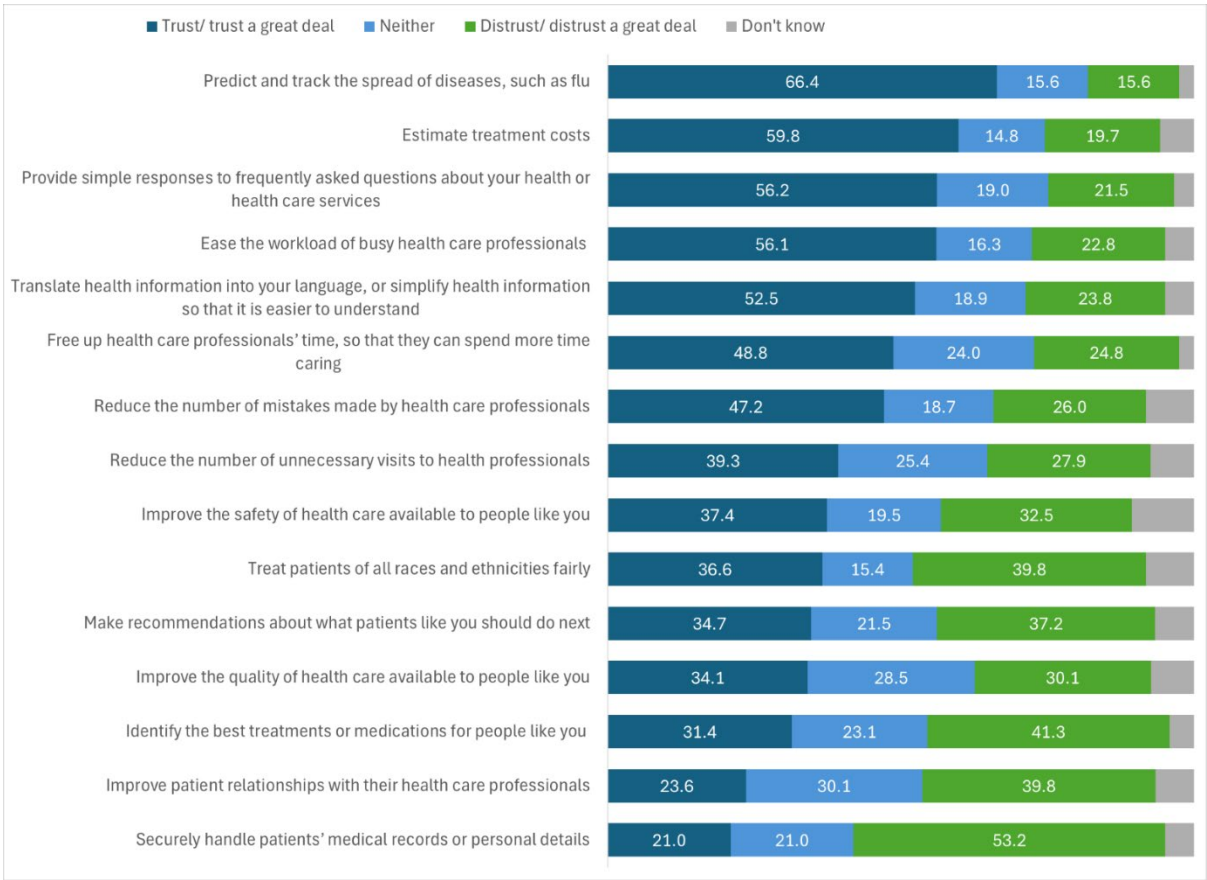

Respondents were also asked to report on their levels of trust in AI being used in various aspects of their health care (Figure 2). More than half of respondents reported that they either "trust a great deal" or "trust" AI to:

- Predict the spread of diseases such as the flu (66.4%)
- Estimate treatment costs (59.8%)
- Provide simple responses to frequently asked questions about their health or health services (56.2%)
- Ease the workload of busy healthcare professionals (56.1%)
- Translate health information into their preferred language, or simplify health information so that is easy to understand (52.5%)

Just 1 in 5 (21.0%) respondents stated they trusted AI to securely handle patients' medical records and information. Approximately half of respondents (53.2%) did not trust AI to do this.

Only around one third of respondents trusted AI to be used in the following aspects of health care:

- To improve the quality of care available to people like you (34.1% trusted AI)
- To identify the best treatments and medications for people like you (31.4% trusted AI)
- To make recommendations on what patients like you should do next (34.7% trusted AI)

Fewer than one in four respondents (23.6%) indicated that they trusted AI would improve their relationship with health care providers.

Figure 3: Consumer interest in having AI support their health care (%) n=124-126

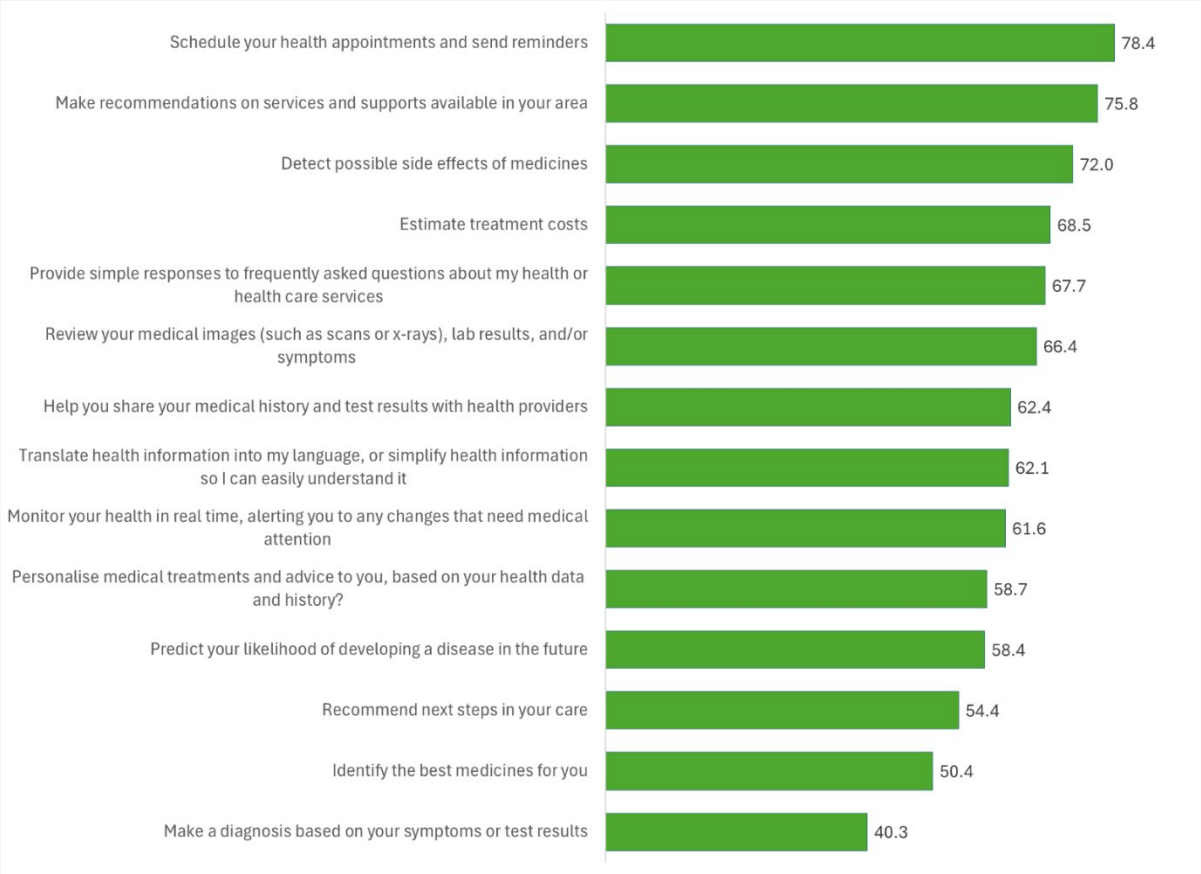

Figure 3 presents data on how interested the respondents were in having AI support their health care in various ways. The most common areas where respondents stated they were interested in AI supporting their health care were for scheduling appointments (78.4%), making recommendations on supports and services in your

area (75.8%), or detecting possible side effects to medicines (72.0%). Respondents were less likely to express their interest in AI supporting decision making to determine their future care, such as recommending the next steps in your care (54.4%), identifying the best medicines for you (50.4%) or making a diagnosis based on your symptoms or test results (40.3%).

Figure 4: Concern about the use of AI in Health care (%) n=119-121

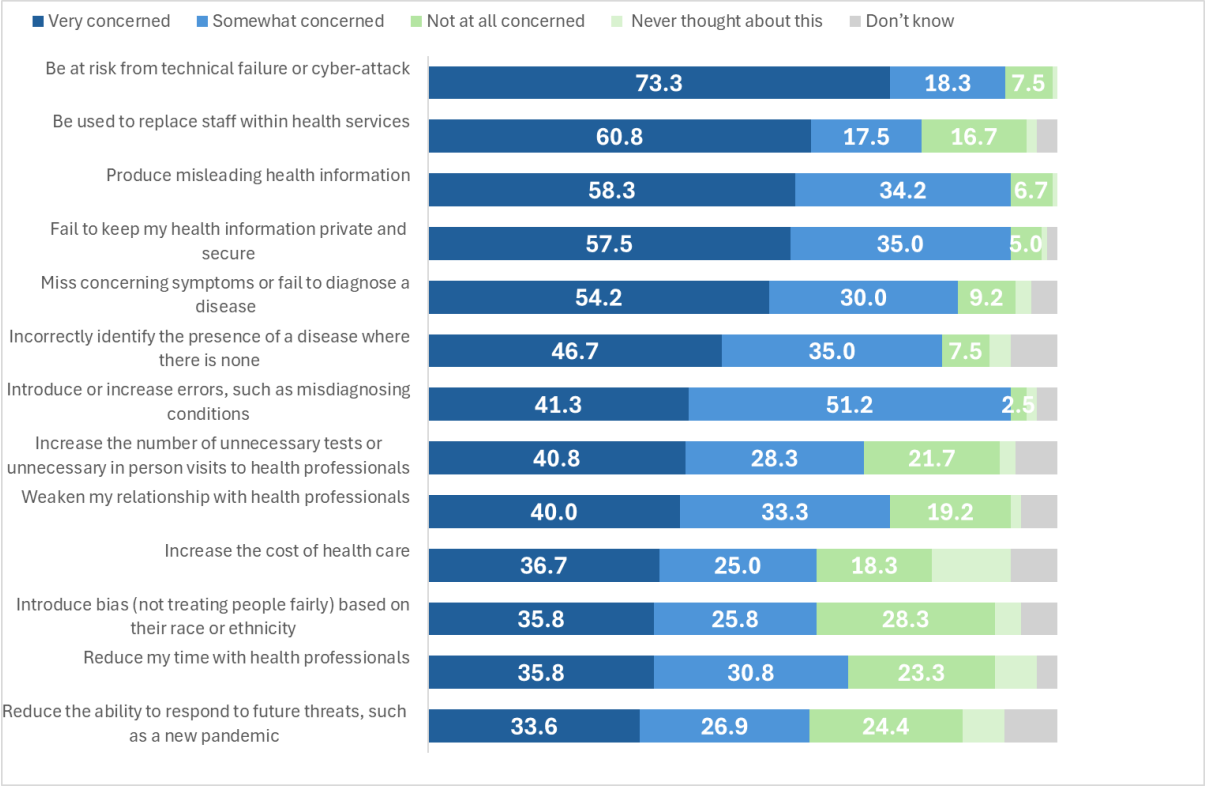

Respondents indicated their levels of concern with AI being used in aspects of their health care (Figure 4).

More than 90% of respondents stated being either "somewhat concerned or very concerned" that AI would

- Introduce or increase errors such as misdiagnosing conditions
- Produce misleading health information
- Fail to keep my health information private and secure
- Be at risk from technical failure or cyber-attack

A couple of areas appeared to be of heightened concern to respondents. The majority of survey respondents reported being "very concerned" that AI would be at risk of technical failure or cyber-attack (73.3%), whilst 60.8% were "very concerned" that AI would be used to replace staff within health services.

Respondents were also likely to express concern about the accuracy of AI when used in diagnoses:

- 84.2% were “somewhat concerned or very concerned” that AI would miss concerning symptoms or fail to diagnose a disease,
- 81.7% were “somewhat concerned or very concerned” that AI could incorrectly identify the presence of disease where there is none.

Figure 5: Comparisons between AI and trained health professionals (%) n=119 to 120

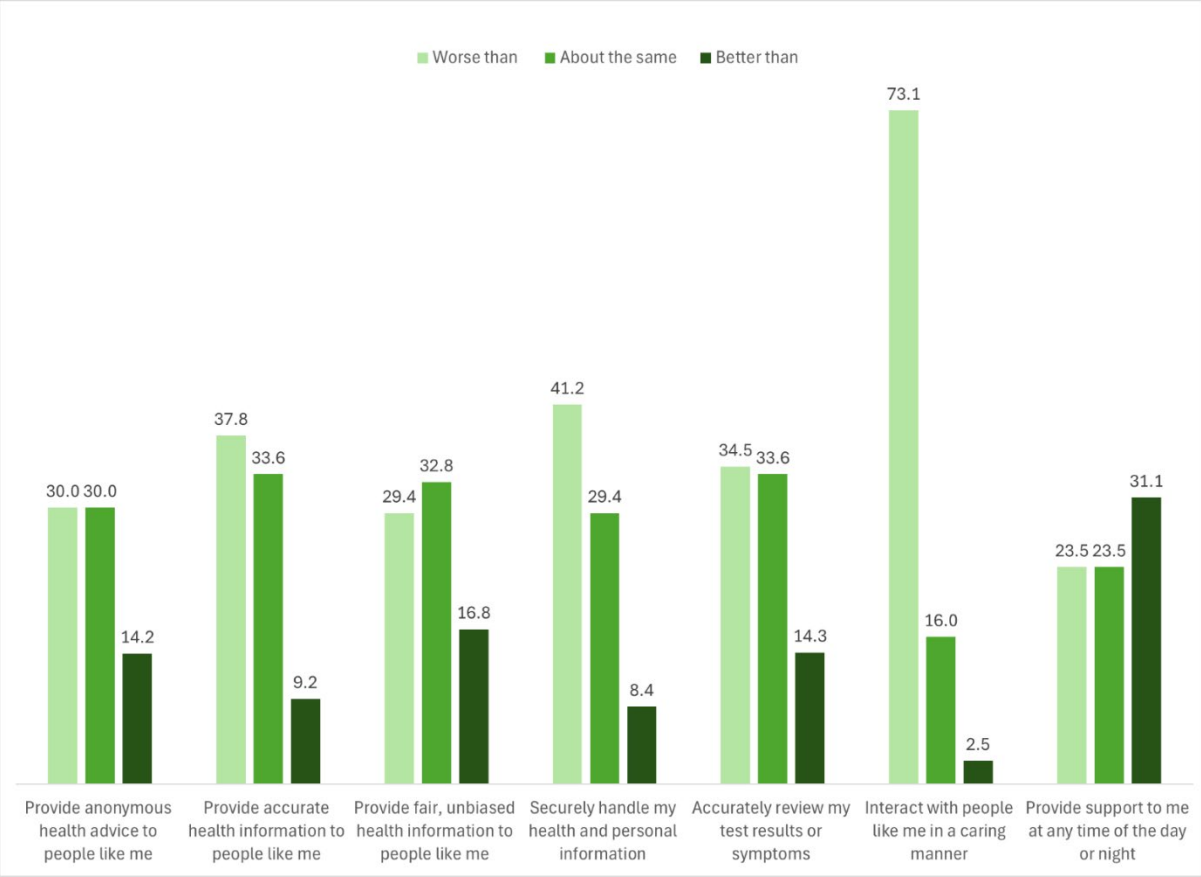

Respondents were asked their perceptions of how support offered by AI compares to support offered by trained health care professionals. On only one variable “provide support to me at any time of the day and night” were respondents more likely to rate AI as being better than trained health care professionals.

The majority of respondents (73.1%) indicated that they believed AI was worse than health care professionals at “interacting with people like me in a caring manner”.

On the remaining measures, respondents tended to rate AI as being slightly worse or about the same as trained health care professionals. These measures included, “accurately review my test results or symptoms”, “securely handle my health and

personal information", "provide fair unbiased health information to people like me", "provide accurate health information to people like me" and "provide anonymous health advice to people like me".

Figure 6: Ongoing use of AI in health care

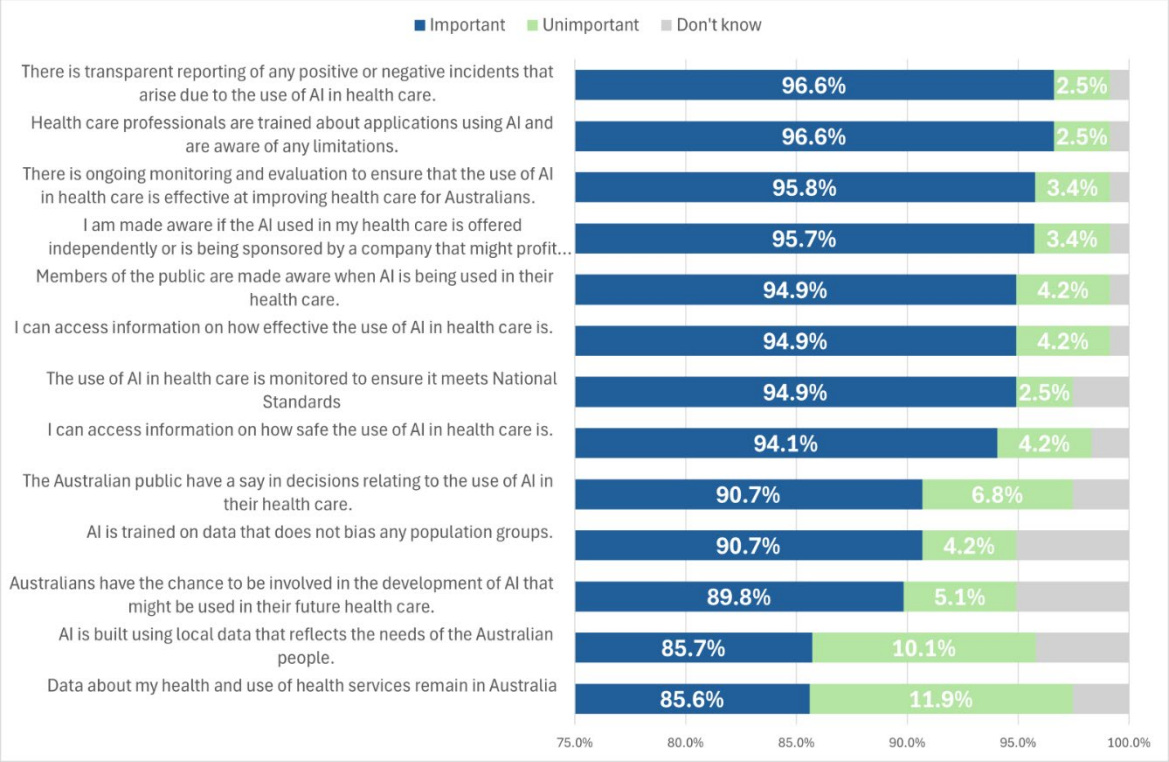

Respondents were asked to think about the use of AI in health care in Australia and respondents were highly likely to rate each statement as being "important or very important". Indeed, all of the above statements (Figure 6) were deemed to be "very important or important" for the ongoing use of AI in healthcare by at least 85% of the respondents.

The issues that respondents were most likely to rate as "important or very important" included: transparent reporting of positive or negative incidents that arise due to the use of AI in healthcare, that health professionals are trained about applications using AI and are aware of any limitations, and that there is ongoing monitoring and evaluation to ensure the use of AI in healthcare is safe.

The issues that respondents were slightly likely to rate as important, related to the use of data behind AI. The majority of respondents stated that it was important for data about health and use of health services to remain in Australia (85.6%) and that AI is built using local data that reflects the needs of Australians (85.7%).

AI Governance

Legislative Frameworks: Australia currently lacks an overarching regulatory framework for using AI. While specific components and uses of AI are regulated by existing legislation, many grey areas must be addressed. AI is mainly self-regulated, and the limited guardrails are designed by AI technology developers with conflicting priorities. On the one hand, developers are eager to avoid legal litigation, but on the other hand, they are driven by profits to increase the scope of use of their software.

It is, therefore, up to governments to ensure that processes are in place to audit AI software. Similar to what happens in cybersecurity – where penetration testing is a standard practice to find vulnerabilities in a computer system – AI auditing should be standard practice and a fundamental step in detecting weaknesses and features that can cause harm to consumers. Additionally, mandatory requirements for AI auditing should also come with clear standards for what counts as an impartial, comprehensive AI audit.

In the case of adverse effects produced by AI, there is currently no clear framework to establish where the responsibility ultimately lies. Should consumers be negatively affected by AI - for example, by receiving wrong health advice - there is no clear pathway for them to report the issue, nor a clear path for institutions to hold AI software developers accountable.

CHF sees a role for developers in creating mechanisms that guarantee safety and quality. Developers know their product better than anyone and should be included as essential stakeholders. But self–regulation alone is not enough to ensure that health consumer interests are protected. In an unregulated environment, the vested interests of AI developers would inevitably prevail over consumers. Without safety and quality standards, well-meaning actors trying to use generative AI for positive change would also run the risk of building unethical products. This could happen due to governance groups commissioning generative AI algorithms without a deep understanding of the algorithmic consequences their product may have and the generative AI companies they have hired to build the software defaulting to a profit-driven approach to coding. Government action is, therefore, paramount in ensuring the presence of an independent regulatory body. To be effective – and to ensure they reflect the needs of Australian society – the regulatory framework must be developed in partnering with consumers.

Other nations are at much more advanced stages of establishing such a framework. The EU Artificial Intelligence Act – for example – is a risk-based regulatory framework in which AI is regulated in a way that is directly proportional to the risk it brings to

consumers. The Act recognises that AI could harm the health and safety of healthcare consumers and focuses on the importance of preventing and mitigating the safety risks that AI technology may pose[4].

CHF calls for risk-based legislation to be adopted in Australia, in which healthcare applications of AI are classified as high-risk. Risk-based frameworks are routinely implemented in Australia. The PBS drug scheduling, for example, is a risk-based framework that sees drugs that pose higher risks to consumers being held to stricter regulations[5]. A risk-based framework would allow Australia to maximise AI technology's benefits while minimising risks.

Safety and Quality risks of AI

Regulatory grey areas: Currently, Australian regulations only partially protect consumers against the risks of AI use in healthcare.

For example, some technology and software used in clinical settings, such as note-taking software, are exempted from TGA approval[6] on the assumption that their only function is to collect patient notes. But what happens when AI is introduced? AI can listen to conversations and generate fully formed patient notes. But whose responsibility will it be if the software makes a mistake in generating medical notes? And how will risks related to privacy and security issues be mitigated if AI transmits health information through unknown servers?

CHF recommends that government bodies work cooperatively with consumers, clinicians, and developers to address critical regulatory grey areas.

Algorithmic opacity: Algorithmic opacity refers to the inability to fully observe or understand the inner workings of an AI model. How can an algorithm be audited if we don't know how an algorithm works? Algorithmic opacity can lead consumers to lose trust in clinicians and the healthcare system. Suppose we cannot understand a decision suggested by AI software because of the exceedingly high number of

[4] European Commission. "Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts." Brussels, Belgium: European Union, April 21, 2021 https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206

[5] Therapeutic Goods Administration. 2022. "Scheduling Basics of Medicines and Chemicals in Australia." Therapeutic Goods Administration (TGA). June 21, 2022. https://www.tga.gov.au/scheduling-basics-medicines-and-chemicals-australia.

[6] Therapeutic Goods Administration. 2022. "Is My Software Regulated?" Therapeutic Goods Administration (TGA). June 21, 2022. https://www.tga.gov.au/resources/resource/guidance/my-software-regulated.

parameters used to reach that decision. How can healthcare consumers feel safe and satisfied that a correct course of clinical action was identified if the deployers of the AI software and the clinicians using it cannot justify or explain how it came to a particular clinical decision? Algorithmic opacity may also lead to harm if the AI errors are not easily detected or understood. The lack of a clinician's full involvement in decision-making puts consumers at risk of inappropriate treatment. It reinforces the need for strong regulatory oversight and developers to disclose information about their AI algorithm to regulatory bodies, should they wish for it to be used in high-risk settings. Knowing that an independent regulatory body has reviewed and understood how an algorithm works will make consumers feel a lot more comfortable and safer about AI being one of the tools used in clinical management.

Bias: Research has shown that the output of AI is affected by gender and racial biases[7]. Adopting AI in healthcare without vetting algorithms for bias means that otherwise high-quality, inclusive services may become less so because of the introduction of AI software. Instead of closing the current gaps, health services risk exacerbating existing discriminatory practices with algorithms that are not equipped to cater to a diverse, multicultural population. Furthermore, AI output will likely not be explicit about their bias, presenting their output as plausible and evidence-based.

CHF calls for safety and quality consideration to include the assessment of algorithms, which must be vetted against bias. Standards must be created that will ensure the safety and quality of algorithms used in Australian healthcare, with a particular focus on ensuring that they don't reinforce commonly held biases and health disadvantages, especially towards Aboriginal, Torres Strait Islander and CALD communities.

The amplifying effect of algorithmic opacity should also be considered when discussing bias in AI systems. AI software may perpetuate these discriminatory practices without transparency, making them harder to identify and address.

AI hallucinations: Generative AI is known to "hallucinate," meaning that the language model will generate content that is not verified or accurate. The information that generative AI produces does not come with a warning advising the reader on what content has been drawn from highly reputable evidence and what content was simply made up: generative AI relies on the reader's ability to have enough expertise to discriminate between the two. This cannot be assumed when a health consumer

[7] Lyell, David, Ying Wang, Enrico Coiera, and Farah Magrabi. "More than algorithms: an analysis of safety events involving ML-enabled medical devices reported to the FDA." Journal of the American Medical Informatics Association 30, no. 7 (2023): 1227-1236.https://doi.org/10.1093/jamia/ocad065

reaches out to generative AI for health information or when a clinician searches online for a solution to a complex issue.

For example, Siontis et al. (2024) have shown how generative AI software Chat GPT (Version 3.5) has recommended the implantation of a subcutaneous defibrillator in a patient immediately after cardiac surgery. Such a step is not something that any cardiac surgeon would recommend, as it would not only interfere with surgery site healing but would also not be able to work appropriately so soon after surgery [8]. When questioned by clinicians on what evidence it used to recommend such a step, Chat GPT provided references to study results and journal articles that did not exist. This research highlights an enormous safety risk for consumers: not only could they be provided with inappropriate care, but they could also follow wrong, harmful health advice at home, which is presented to them as the fruit of rigorous scientific study.

Clinician over-dependence on AI tools: Clinicians may become overly dependent on AI, eroding their clinical judgment and decision-making skills. Relying on AI could make them more hesitant or unable to challenge AI-derived recommendations.

Lack of clarity in accountability: When the worst happens, and AI causes an error in clinical management and an adverse outcome for a healthcare consumer, it may be unclear whether the clinician or the AI system was ultimately responsible. Lack of clarity can lead to complex legal and ethical questions about liability and accountability, which will cause immense pain to affected consumers and clinicians.

Depersonalised healthcare experience: Overuse of AI could lead to a depersonalised healthcare experience, with health consumers feeling like data points rather than individuals with unique needs. The human connection, empathy, and understanding, which are fundamental elements of care provision, may become weakened and overlooked.

Distrust in the healthcare system: Concerns about algorithmic errors, transparency, and ethical considerations might lead to scepticism and disenfranchisement from the healthcare system and undermine the public's perception of legitimacy in healthcare institutions. Suppose people feel their health is managed by unaccountable algorithms rather than trusted human professionals. In that case, the social contract that supports democratic governance and legitimate authority in public services might be weakened.

[8] Siontis, Konstantinos C., Zachi I. Attia, Samuel J. Asirvatham, and Paul A. Friedman. "ChatGPT hallucinating: can it get any more humanlike?." (2024): 321-323. https://doi.org/10.1093/eurheartj/ehad766

Data, Privacy and Ethics

Data safety and privacy are of paramount importance for consumers. AI is bound to collect extensive amounts of data when utilised in clinical settings, and consumers have the right to know where and how this data is stored and used.

Data safety: Specific legislation that safeguards data collected and used by AI throughout its entire lifecycle, from data collection to storage to data elimination, needs to be implemented. Recent episodes of data breaches in Australia have highlighted how any weakness can result in vast amounts of identified, sensitive data being leaked. The stakes are exceptionally high with healthcare data. As such, CHF calls for the government to leverage protections already in place - such as data and consumer protection legislation – to create a comprehensive data safety framework for AI in healthcare that will protect consumers.

Data guardianship: Legislation must clearly state who can access data collected via AI and how data is collected, stored and used. Healthcare consumers' consent for data use must be informed and dynamic. Informed consent means that healthcare consumers must be notified if the data is about to be used in ways that depart from the original use to which they consented. Dynamic consent means that consumers must be able – at any time – to modify or withdraw their consent for their data being used. Clear procedures must be in place for data disposal, either by ensuring that the data is destroyed safely or ensuring that the entity the data is transferred to for custody and ownership also meets all safety and privacy requirements. To do all this, CHF calls for a comprehensive framework of independent data guardianship equipped with the tools to allow for swift, decisive action in case of data and privacy breaches.

Perverse profiting from acquired data: Legislation must also not avoid the challenges of legislating protections that prevent data from being profited from. For example, private health insurers must be denied access to this data, and it should be illegal for a private insurer to refuse a policy or a claim based on AI-collected data. This scenario sounds like a far-removed scenario, but it isn't. In the United States, several health insurers (Humana, Cigna, UnitedHealthcare) are already facing class actions for allegedly deploying AI technology to deny claims [9].

[9] "AI Lawsuits against Insurers Signal Wave of Health Litigation." n.d. News.bloomberglaw.com. https://news.bloomberglaw.com/health-law-and-business/ai-lawsuits-against-insurers-signal-wave-of-health-litigation.

Consumer complaints: Consumers must be able to raise complaints to an independent regulator when AI has caused safety and quality issues or harm.

Consumer Engagement

Raising consumer AI literacy: As discriminative and generative AI become increasingly established, consumers must be provided with the tools to learn what AI can and can't do for them.

Generative AI poses significant challenges and concerns, as it can churn out vast amounts of information, which is presented very confidently but is not necessarily accurate as it might be affected by biases and hallucinations described earlier in this submission. This poses a risk, particularly in health information, as consumers may be influenced to make decisions about their health on data that seems credible but is potentially false or even harmful.

As AI becomes increasingly popular, first-time users will attempt to use generative AI without knowing its capabilities and associated risks. Countries like Finland have started to build tools to minimise such risk: the University of Helsinki's "Elements of AI"[10] course – for example – is attempting to bridge that gap with six learning modules that explain what AI is, what issues it can solve, its real-world applications, how it works and how it mimics neural networks as well as its implications.

CHF recommends that a similar free course be developed in Australia or adapted from similar tools already developed overseas, with a specific focus on the benefits and limitations of AI in healthcare. This is necessary to ensure that the content is widely accessible to Australia's diverse communities and that its contents are tailored to Australian society. By providing accessible and understandable information about AI, consumers can navigate the digital landscape more confidently, making informed decisions about their health and well-being and avoiding major pitfalls.

If provided with the appropriate resources, an independent health consumer peak body such as CHF would be best placed to ensure that the themes of such a learning tool are relevant to consumers and written in accessible language.

Clinicians must understand AI and engage with consumers about AI: They must familiarise themselves with AI and how it can be used in healthcare. This would allow clinicians to use AI to benefit healthcare consumers by providing more accurate

[10] "A Free Online Introduction to Artificial Intelligence for Non-Experts." n.d. English. https://www.elementsofai.com/.

diagnoses, more streamlined decision-making processes, and more personalised care. Better AI literacy in clinicians would also positively impact the safety of AI: familiarity with AI will ensure clinicians can critically evaluate the recommendations and processes generated by the AI software. Just as with other technological advancements, clinician active engagement is paramount in ensuring that the benefit of new technology is utilised safely and effectively.

Healthcare providers will also play a pivotal role in bridging the gap between the complexities of AI and healthcare consumers' understanding. As trusted sources of advice, clinicians are best placed to inform healthcare consumers on how AI will impact their healthcare journey. Clinicians can empower consumers to navigate through these changes and allow consumers to actively participate in determining their treatment plans.

For all the above reasons, CHF recommends that efforts be made to ensure that generative AI is introduced under the stewardship of the Australian Commission on Safety and Quality in Healthcare. CHF recommends that the National Safety and Quality Health Service Standards (NSQHS)[11] be amended to reflect and explicitly address generative AI use and its implications.

Generative AI is bound to change how health services run, encompassing every facet of the NSQHS standards, from clinical governance to risk management to how health services communicate and partner with consumers. As such, the NSQHS standards would be the first port of call for health services wanting to find guidance on safely integrating generative AI into their processes.

Lastly, the amendment of the NSQHS standards must happen in partnership with consumers. Chapter 2 of the NSQHS establishes that the healthcare workforce must partner with consumers in planning, delivering, measuring and evaluating systems and services and that consumers must be partners in their care to the extent they choose. Therefore, establishing standards of use of generative AI in healthcare must also be consistent with this approach.

[11] Australian Commission on Safety and Quality in Health Care. 2021. "National Safety and Quality Health Service Standards (Second Edition) | Australian Commission on Safety and Quality in Health Care." Australian Commission on Safety and Quality in Health Care. May 2021. https://www.safetyandquality.gov.au/publications-and-resources/resource-library/national-safety-and-quality-health-service-standards-second-edition.

References

AI Lawsuits against Insurers Signal Wave of Health Litigation." n.d. News.bloomberglaw.com. https://news.bloomberglaw.com/health-law-and-business/ai-lawsuits-against-insurers-signal-wave-of-health-litigation.

"A Free Online Introduction to Artificial Intelligence for Non-Experts." n.d. English. https://www.elementsofai.com/.

Australian Commission on Safety and Quality in Health Care. 2021. "National Safety and Quality Health Service Standards (Second Edition) | Australian Commission on Safety and Quality in Health Care." Australian Commission on Safety and Quality in Health Care. May 2021. https://www.safetyandquality.gov.au/publications-and-resources/resource-library/national-safety-and-quality-health-service-standards-second-edition.

1 Dorricott, Pippa. n.d. "AI Can Revolutionise Healthcare but Only with a National Plan | Ai Health Alliance." https://aihealthalliance.org/2023/11/16/ai-can-revolutionise-healthcare-but-only-with-a-national-plan/.

European Commission. "Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts." Brussels, Belgium: European Union, April 21, 2021 https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206

"General Practice Data and Electronic Clinical Decision Support – Issues Paper Consultation | Consumers Health Forum of Australia." 2022. Chf.org.au. March 4, 2022. https://chf.org.au/publications/general-practice-data-and-electronic-clinical-decision-support-issues-paper

"GP Data and Electronic Clinical Decision Support | Consumers Health Forum of Australia." 2023. Chf.org.au. June 15, 2023. https://chf.org.au/publications/gp-data-and-electronic-clinical-decision-support.

Lyell, David, Ying Wang, Enrico Coiera, and Farah Magrabi. "More than algorithms: an analysis of safety events involving ML-enabled medical devices reported to the FDA." Journal of the American Medical Informatics Association 30, no. 7 (2023): 1227-1236.https://doi.org/10.1093/jamia/ocad065

Siontis, Konstantinos C., Zachi I. Attia, Samuel J. Asirvatham, and Paul A. Friedman. "ChatGPT hallucinating: can it get any more humanlike?." (2024): 321-323. https://doi.org/10.1093/eurheartj/ehad766

Therapeutic Goods Administration. 2022. “Is My Software Regulated?” Therapeutic Goods Administration (TGA). June 21, 2022. https://www.tga.gov.au/resources/resource/guidance/my-software-regulated.

Therapeutic Goods Administration. 2022. “Scheduling Basics of Medicines and Chemicals in Australia.” Therapeutic Goods Administration (TGA). June 21, 2022. https://www.tga.gov.au/scheduling-basics-medicines-and-chemicals-australia.

“Webinar: Automated Decision Making and Artificial Intelligence in Health | Consumers Health Forum of Australia.” 2022b. Chf.org.au. June 2, 2022. https://chf.org.au/events/webinar-automated-decision-making-and-artificial-intelligence-health