RACP
Specialists. Together
EDUCATE ADVOCATE INNOVATE

RACP submission to the Safe and Responsible Artificial Intelligence in Health Care – Legislation and Regulation Review

October 2024

About The Royal Australasian College of Physicians (RACP)

The RACP trains, educates and advocates on behalf of over 21,000 physicians and 9,000 trainee physicians, across Australia and Aotearoa New Zealand. The RACP represents a broad range of medical specialties including general medicine, paediatrics and child health, cardiology, respiratory medicine, neurology, oncology, public health medicine, infectious diseases medicine, occupational and environmental medicine, palliative medicine, sexual health medicine, rehabilitation medicine, geriatric medicine, and addiction medicine. Beyond the drive for medical excellence, the RACP is committed to developing health and social policies which bring vital improvements to the wellbeing of patients and the community.

We acknowledge and pay respect to the Traditional Custodians and Elders – past, present and emerging – of the lands and waters on which RACP members and staff live, learn and work. The RACP acknowledges Māori as tangata whenua and Te Tiriti o Waitangi partners in Aotearoa New Zealand.

Introduction

We thank the Department of Health and Aged Care (DoHAC) for the opportunity to provide feedback on the *Safe and Responsible Artificial Intelligence in Health Care – Legislation and Regulation Review.* This is an important consultation for the RACP, its members and their patient communities as artificial intelligence (AI) has significant potential to impact clinical practice, practice administration and patient care.

This submission draws on the perspectives of individual members from the Australasian Faculty of Occupational and Environmental Medicine (AFOEM), the Australasian Faculty of Public Health Medicine (AFPHM), the Australasian Faculty of Rehabilitation Medicine (AFRM), the RACP Paediatrics and Child Health Division (PCHD), the RACP Digital Health Advisory Group, the RACP Ethics Committee and the RACP Workforce Policy and Advocacy Advisory Group. Additionally, insights were drawn from a recent workforce issues survey of practising Australian members which included a component on AI and related topics.

What benefits will AI have for healthcare consumers and providers? (consultation questions 1-2)

AI applications offer potential for improving efficiency by saving time for direct patient care, streamlining clinical workflows, enhancing diagnostic and prognostic accuracy, identifying trends and patterns not immediately seen by human observation, optimising treatment selection and personalised care, empowering patient self-management and reducing costs.

The potential applications of AI in the clinical practice of physicians include, where appropriate, and in line with legal and ethical requirements:

- Automating administrative and clerical tasks
- Providing sophisticated clinical decision-making support
- Risk prediction
- Educating patients
- Remote monitoring devices to assess disease progress and optimise management
- Precision medicine – one RACP member identified particular ability to analyse genetic, environment and lifestyle factors to guide treatment options.

AI tools can have a significant impact on ancillary services such as radiology and pathology, on whose results much of physician decision-making relies.

The RACP's recent survey of practising Australian members also revealed that 78% believed that AI would add value to their work through 'note taking'. This was the most commonly identified purpose for AI tools. Other commonly identified purposes included the use of AI tools for 'patient/carer communication materials' (60%) and as a 'scheduling assistant' (57%).

Further individual feedback from RACP members garnered in the course of this specific consultation identified that AI could improve access to care in the form of telemedicine, diagnostic tools and virtual health assessments. These applications offer remote consultation, real-time monitoring and early detection of diseases. Situation specific indications can also be developed in consultation with physician and other healthcare specialties. Members also identified potential benefits from AI in helping to alleviate workforce pressures, including those posed by heavy workloads and long work shifts, through clinical decision and monitoring support. However, appropriate regulations and guardrails must be put in place.

Given the purposes of AI use members highlighted that regular validation of AI systems in real-world scenarios is essential. A testing and evaluation framework is needed that includes regular audits of AI performance, patient outcomes and cost-efficiency metrics. There is a need for agreed performance outcome measurements. Possible key performance indicators include diagnostic accuracy, wait times, patient satisfaction and healthcare costs. Stakeholder engagement should continuously guide the implementation of the testing and evaluation framework to assure that AI technology meets clinical needs and patient safety standards.

Regulatory standards are needed to provide clear provisions for appropriate cross-jurisdictional practice, allowing healthcare providers using AI to extend services across state and regional boundaries while allowing

for nationally consistent and unified monitoring of risk. Regulations that outline certification and accountability of AI systems would help build trust among patients and healthcare professionals.

In regulating the use of AI in healthcare emphasis must also be placed on transparency and explainability, ensuring that providers and patients understand how AI is used in care delivery and how it got to the output / answer.

Overall, regulatory frameworks should prioritise patient safety and equitable access to AI-driven healthcare solutions.

Are there specific low or high risks to using AI in healthcare and what criteria could be used to classify risk? (consultation questions 3-5)

The risks associated with the use of AI tools in healthcare are variable. Lower risk AI tools would involve those that automate routine administrative or clerical duties, whereas higher risk AI tools would involve assistive AI for image-based interpretation and augmentative AI for predicting disease risk and prognosis.

Given the variable risks of AI technology, the RACP expects that the implementation of AI will be staggered into phases according to risk category for appropriate refining, testing and technology integration. The criteria for characterising risk related to the implementation and use of AI tools should consider security, privacy, confidentiality and financial aspects. An RACP member suggested using the severity assessment coding system to categorise AI related health care risks as this is an established measure within Australian healthcare.[1]

Even in cases where AI is used for lower risk functions, consumers should be informed as some may hold negative attitudes towards AI being used in health care, limiting its use.[2] An option is the development of a Patient (or Consumer) Charter for AI, one that details the person-centric requirements for responsible and ethical use of AI in clinical care, particularly patient rights to understand where AI is used in components of their healthcare and the legal, professional and ethical requirements for consent.

Key concerns around the use of AI in clinical practice centre on:

- Amplification and propagation of bias in datasets used to train AI models
- Model inaccuracies or 'hallucinations' (incorrect or misleading outputs)
- Lack of transparency and accountability in AI model development and testing.

Other sources of concern are:

- Physician medicolegal liability in the event of patient harm arising from AI errors or misuse
- Automation bias
- Increasing overdiagnosis
- Data breaches and infringements on patient privacy
- Patient distrust of AI.[3,4]

Disruptions to clinical workflows and opportunity costs incurred by AI training programs, the expense of maintaining software and hardware operations, and the environmental impact of AI tools consuming water and energy must also be considered.[5,6]

In our recent survey of practicing Australian members 86% indicated that the greatest risk related to the use of AI tools in clinical practice would be 'accuracy'. Other commonly identified risks included the risk of AI use on 'data security' (76%) and on 'confidentiality' (61%).

---

[1] NSW Health. Incident Management Policy.https://aci.health.nsw.gov.au/__data/assets/pdf_file/0011/319718/severity-assessment-code-pd2014-004.pdf

[2] Young AT, Amara D, Bhattacharya A, Wei ML. Patient and general public attitudes towards clinical artificial intelligence: a mixed methods systematic review. Lancet Digit Health 2021; 3: e599-e611.

[3] Senevirathna P, Pires DE, Capurro D. Data-driven overdiagnosis definitions: A scoping review. J Biomed Inform 2023;147:104506.

[4] Capurro D, Coghlan S, Pires DE. Preventing digital overdiagnosis. JAMA 2022;327(6):525-6.

[5] Dhar P. The carbon impact of artificial intelligence. Nat Mach Intell 2020; 2: 423-425.

[6] Bloomfield PS, Clutton-Brock P, Pencheon E, et al. Artificial intelligence in the NHS: climate and emissions. J Clim Chang Health 2021; 4: 100056.

When assessing the ability to implement AI into regional, rural and remote (RRR) settings and related operational risks the virtual and digital infrastructure of the area needs to be considered. Often RRR settings have more limited access to healthcare and adequate technological and digital infrastructure that can facilitate AI implementation. Considerations such as improving internet access, training and system robustness in RRR settings and ensuring accessibility and affordability should all be considered.

Model bias, or misrepresentation of population data by models, is often disproportionately distributed to under-served populations with poorer health, reinforcing the need for representative training and input data. This also highlights the need for cultural safety in AI practice and appropriate design for use, especially for our First Nations communities, culturally and linguistically diverse peoples, complex patient populations and lower-income patient communities, particularly given the health disparities between our metropolitan and RRR areas. However, one RACP member identified that AI could help overcome biases where its data is limited to verified information.

The use of AI tools in healthcare will be heavily influenced by consumer perceptions and ultimately, by legal consent and privacy requirements.

RACP members emphasised the need for choice by physicians, other healthcare providers and patients in the use of AI. However, providers may not have an explicit choice in the use of AI as part of their practice as it can be implemented upstream, by third parties, or in services external to their control, such as radiology and pathology. As such, options for managing AI tools and training / education to guide the use of AI tools are required. Physicians should be informed fully in advance where AI is deployed or used. This call would be reflective and sensitive of some uncertainties expressed by Australian respondents in our recent survey, which showed that 20% of respondents remained uncertain about whether AI had a role to play in their practice. Decisions around where it is expected to be used, further training, information and education should be given to the physician, especially given patient privacy and consent implications.

Is regulation the answer, what are the possible unique regulatory changes needed for AI in healthcare, and could we learn from overseas? (consultation questions 6-10)

In healthcare, any AI tool must address a pressing real-world operational or clinical problem with clear, actionable and evidence-based outcomes. AI decision support applications must enhance current decision-making for commonly encountered but challenging scenarios, such as early detection of sepsis.[7] To be effective AI applications must also perform better than current well accepted, high performing but simpler decision or prediction rules. AI developers must involve clinician end-users in acquiring a deep, mutual understanding of the clinical task, the datasets being targeted and why, their amenability to AI, current clinical decisional performance, and end-user needs and goals, ideally expressed as measurable targets.

For classical machine learning (ML) systems, users should be able to interact with intuitive, user-friendly human-computer interfaces that quickly provide advice or alerts that are visually presented in a standardised format, are easily interpretable, and allow clinician discretion in accepting or over-riding them is critical.[8] Evidence suggests some clinicians supporting AI use prefer graphic or numeric displays of probabilities or alert thresholds, often colour-coded according to seriousness or acuity (e.g. low, moderate, high), for a diagnosis or event, coupled with links to relevant, consistent recommendations for tests or treatments. As for generative AI, such as ChatGPT, clinicians would require prompts (inputs) to contain all relevant information in enabling the AI to generate responses as accurate and as relevant as possible.

AI applications must blend seamlessly into clinical workflows, avoid creating workarounds and alert fatigue, allow customisation of alert thresholds to local populations, and prevent both cognitive overload and over-reliance on AI outputs. Involving clinician end-users is critical in detailing current operational context, pre-empting training and support needs, and raising awareness of how incorrect use by clinicians, such as

[7] Sanders S, Doust J, Glasziou P. A systematic review of studies comparing diagnostic clinical prediction rules with clinical judgment. PLoS One 2015;10:e0128233

[8] Asan O, Choudhury A. Research trends in artificial intelligence tools in human factors health care: mapping review. JMIR Hum Factors 2021;8:e28236.

inputting data errors, misinterpreting information displays or clicking wrong options, can lead to patient harm or other adverse implications.[9]

Using AI tools, especially large language models (LLMs), to produce evidence syntheses, clinic letters and discharge summaries may free up cognitive time and space for clinicians to engage more in empathetic, person-centred shared decision-making (SDM). However, the true impacts of AI applications on clinician-patient interactions in different contexts, application designs that best support each step of SDM, and best methods for obtaining patient consent to AI being used to assist SDM all remain to be defined.[10]

Other considerations include the need to engage with consumers to account for ethical issues as well as issues with data security, algorithmic bias and fairness, patient safety, privacy and informed consent as outlined above.

These unique considerations may be addressed through targeted regulatory changes in clinical validation standards, data privacy and security regulations, bias mitigation guidelines, explainability and transparency requirements and liability frameworks.

The need for safeguards for and continual evaluation of the use of AI tools is greater for healthcare as compared to other sectors. Factors that may play a role in the need for regulation of AI in healthcare compared to other sectors include the direct impact of physician practice on human lives, the complexity and sensitivity of patient data, the rigid and complex healthcare regulatory environment, the need for high accuracy and low error margins, ethical considerations, patient consent, explainability, trust and integration with physician expertise.

In adopting AI, regulatory frameworks must clearly determine when liability for errors and resultant patient harm from AI use lies primarily with users and their personal indemnity insurer (e.g., negligent or other inappropriate use), their employing organisation, and / or the application developers, vendors and other entities. Liability could conceivably extend to 'failure to use' if using a specific AI application becomes a practice standard for certain clinical scenarios. Such frameworks remain works in progress for the Therapeutics Goods Administration (TGA) and other regulatory bodies attempting to balance regulation with innovation and aligning approval procedures with evolving governance procedures but are greatly needed.

More autonomous tools or those directly impacting critical clinical decisions will require greater regulatory oversight and higher levels of safety evidence for approval.[11] Ongoing monitoring of application performance, auditing processes and in-built self-improvement feedback loops will be needed in ensuring the application remains resilient to dataset shifts, external 'noise' / distractions and cyberattacks.[12] While formation of a dedicated AI in healthcare body is one option, there is a great number of existing bodies and organisations that need to collaborate to ensure AI is implemented into healthcare efficiently and ethically; this requires further consultation. RACP members emphasised a need for oversight mechanisms to be healthcare specific.

In adopting AI, implementation fundamentals must be met at inception and throughout the life cycle of an AI application for assurance.[13] In particular, strong data governance policies must ensure data privacy, security and standardisation. Training datasets must be representative, reliable and accessible. Clear, shared definitions for clinical terms and data elements are necessary for consistency and effective communication.

The governance framework must also include clinical governance that mandates clinician and executive oversight and leadership, prioritisation processes for ranking and selecting AI applications for development, determining if changes are required to digital infrastructure and clinical workflows, undertaking workforce training, and deploying or decommissioning an application if it proves valueless, reduces safety, undermines patient ethical and legal requirements, is not implementable or scalable, fails in prospective evaluations or leads to potentially unsafe over-reliance.

Key legislative and regulatory principles would be to:

[9] Salwei ME, Carayon P. A. Sociotechnical systems framework for the tool of artificial intelligence in health care delivery. J Cogn Eng Decis Mak 2022;16:194–206
[10] Sauerbrei A, Kerasidou A, Lucivero F, et al. The impact of artificial intelligence on the person-centred, doctor-patient relationship: some problems and solutions. BMC Med Inform Decis Mak 2023;23:73.
[11] Bitterman DS, Aerts H, Mak RH. Approaching autonomy in medical artificial intelligence. Lancet Digit Health 2020;2:e447–9.
[12] Liu X, Glocker B, McCradden MM, et al. The medical model audit. Lancet Digit Health 2022;4:e384–97.
[13] Reddy S, Allan S, Coghlan S, et al. A governance model for application of AI in healthcare. J Am Med Inform Assoc 2020;27:491–7.

- Clearly define and stabilise regulatory requirements by adopting a quality management system accreditation approach, in collaboration with the TGA.
- Use a pragmatic, tiered approach to regulation based on the sensitivity of data used in training and model inputs.
- Evaluate data security measures required in accordance with data risk category.
- Incorporate ethical considerations and processes.
- Create sandbox environments, taking on an iterative approach to the development of regulatory guidance on the basis of new knowledge.
- Establish indemnity arrangements in which clinician, employers and model developers and vendors share the liability risk depending on the relative attribution of error to each party and the function (assistive, augmented, autonomous) of the application.
- Align liability and incentives so that the individual(s) or entity(ies) best positioned to know the AI system risks and to avert or mitigate harm do so through AI design, development, validation, and implementation.
- Provide clinician protections from liability when they do not know or have no reason to know the quality and safety concerns of an AI-enabled technology.
- Consider no-fault compensation payments for patients harmed by AI-related errors.

Recent high-profile corporate data breaches, system outages and cyberattacks (some associated with ransomware) suggest there will be raised anxiety that sensitive private data can be accessed and used by unauthorised parties, anxiety that digital systems on which AI relies may go down, and AI applications may be corrupted and rendered harmful by adversarial cyberattacks. Key guiding principles for data protection would include embedding strong data privacy protections, perhaps similar to the US Health Insurance Portability and Accountability Act (HIPAA) or the EU General Data Protection Regulation (GDPR), including:

- Employing federated learning, localised architectures and, where appropriate, synthetic datasets in model development.
- Ensuring back-up systems are in place to maintain high-risk, care-sensitive applications.
- Prioritising AI applications using data contained within firewall protected, health service-controlled servers over commercial or third-party software requiring data transmission to external servers.
- Implementing several robust AI algorithms rather than rely on a single one.
- Overlaying predictive algorithms that can detect and suspend corrupted or compromised algorithms.
- Cross-examining algorithms and LLMs for risk of data breach using penetration tests for adversarial attacks.
- Developing benchmarks to evaluate the privacy vs utility trade-off.
- Implementing validation frameworks for multimodal AI evaluation.

We reiterate that a nationally coordinated governance and regulatory framework encompassing local, state and national levels is needed that protects both physicians and patients to use AI applications in their clinical care. Physicians will continue to carry primary responsibility for patient safety through their duty of care and be held accountable for patient harm, until such frameworks are established with binding legislation.

A governance framework, encompassing both data and clinical governance for AI tool development and evaluation, and a regulatory framework for scrutinising and approving AI tools for deployment, must be in place in ensuring patient safety and protecting physicians from legal liability for AI tool-induced harm.

Overseas a number of position statements and roadmaps for AI use in healthcare have been released, which Australia could source guidance from:

- Guidance relating to use of AI in healthcare has been issued from the World Health Organisation.[14]
- In the US, the American Medical Association has released updated principles for the development, deployment and use of healthcare AI,[15] and the American College of Physicians has released a policy position paper on the use of AI in the provision of healthcare.[16]

[14] World Health Organization. Ethics and governance of artificial intelligence for health: WHO guidance. Geneva: WHO, 2021. https://www.Who.int/ publications/i/item/9789240029200
[15] American Medical Association. Principles for augmented intelligence development, deployment and use. Accessed 8/6/24 at: www.ama-assn.org/system/files/ama-ai-principles.pdf
[16] Daneshvar N, Pandita D, Erickson S, et al. Artificial intelligence in the provision of health care: An American College of Physicians policy position paper. Ann Intern Med 2024;177:964-967.

- In the UK, the Royal College of Physicians has released a position statement that urges industry to address real-world challenges, doctors to appraise the technology and regulators to develop guidance and evaluation methods.[17]
- In Canada, the Canadian Medical Association has submitted various recommendations to the House of Commons Standing Committee on Industry and Technology on how AI should be used in providing healthcare.[18]

While the RACP has not aligned itself with any international approaches there are a number that look to tackle issues with AI implementation in healthcare in settings similar to Australia. International approaches worth reviewing include:

- The EU's General Data Protection Regulation (GDPR)[19] and AI Act[20]
- The US's Food and Drug Administration report 'Artificial Intelligence and Medical Products: How CBER, CDER, CDRH and OCP are Working Together'[21]
- The UK's National Health Service AI lab[22] and NICE framework[23]
- Canada's Pan-Canadian AI Strategy[24]
- Singapore's Model AI Governance Framework[25]

Is AI always right and what is the role of humans interfacing with AI? (consultation questions 11-13)

RACP members believe AI-enabled technologies should augment but not usurp clinical decision-making of physicians and that humans should be able to overrule AI findings or decisions.

The training, observations and reasoning of physicians must remain the central tenet of patient care. This requires a commitment by physicians not to overly rely on AI tools in their decision-making.

In all stages of development and use, members noted that AI tools must involve user-centred design and engagement in ensuring these AI tools reduce physician and patient burden in support of patient care and are seamlessly integrated into clinical workflows. This requires multidisciplinary collaboration between AI tool developers, data and ICT specialists, physicians and patients in defining the problem to be addressed by the AI tool and designing it in ways that render it easy to use.

AI developers and implementers must expect, monitor and mitigate biases and errors in AI tools that threaten patient safety and risk exacerbating inequities in care delivery.

As outlined above, any deployed AI tool must be subject to a process of continuous validation, feedback and optimisation over its life cycle in ensuring clinical safety and effectiveness. The real-world performance of the AI tool must be closely monitored using standardised measures that can detect any drift in accuracy. Users of the AI tool must be able to report bias or errors they perceive as degrading tool performance to developers and regulatory authorities. Standard clinical governance and continuous quality improvement principles or risk database management systems can be used to monitor quality of AI driven care.

RACP members were also supportive of errors made by AI being reported, potentially through using similar approaches to existing healthcare adverse outcome reporting, such as centralised databases.

How should healthcare data, privacy, consent and transparency be managed for AI use? (consultation questions 14-19)

[17] RCP Position Statement on AI. Accessed 9/6/24 at: Artificial intelligence (AI) in health | RCP London
[18] Canadian Medical Association. Submission to the Standing Committee on Industry and Technology (INDU): An Act to enact the Consumer Privacy Protection Act, the Personal Information and Data Protection (cma.ca)
[19] EU General Data Protection Regulation. 2016. https://gdpr-info.eu/
[20] The EU Artificial Intelligence Act. 2024. https://artificialintelligenceact.eu/
[21] U.S. Food and Drug Administration. Artificial Intelligence and Medical Products: How CBER, CDER, CDRH and OCP are Working Together'. 2016. https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device
[22] National Health Service England. The NHS AI lab. https://transform.england.nhs.uk/ai-lab/
[23] National Institute for Health and Care Excellence. Evidence standards framework (ESF) for digital health technologies. 2022. https://www.nice.org.uk/about/what-we-do/our-programmes/evidence-standards-framework-for-digital-health-technologies
[24] Government of Canada. Pan-Canadian Artificial Intelligence Strategy. 2022. https://ised-isde.canada.ca/site/ai-strategy/en
[25] Personal Data Protection Commission Singapore. Model AI Governance Framework. 2020. https://www.pdpc.gov.sg/help-and-resources/2020/01/model-ai-governance-framework

In building trust in AI tools on the part of both patients and physicians, the development, testing and use of AI for patient care must be transparent, accountable and collaborative.

Transparency means the ability, as much as possible, to know when AI tools are being used to inform care (with physicians and patients having the right to opt out of using them), to know how personal information is being collected and used, and to know the basic facts of how the model works in generating its outputs.

Accountability means there is proper oversight of tool development and performance, existing and future AI-related policies and guidance are enforced, and mechanisms are in place for rapidly identifying and addressing errors or adverse events resulting from the use of AI tools.

The predictions, recommendations or other outputs of the AI tool, classical or generative, must be able to be readily understood and actioned.[26] However, these outputs do not necessarily have to be fully explainable in how they were derived, particularly for complex, deep learning models which, in many instances, are more accurate than simpler, more inherently transparent models such as decision trees and regression equations. With classical ML prediction models, various techniques (such as LIME [Local Interpretable Model-agnostic Explanations] and SHAP [(SHapley Additive exPlanations] values for structured data-base models; heat maps for image-based models) seek to identify, in retrospect, salient input features strongly associated with the output. But these post-hoc explanations do not necessarily reflect true causal relationships.[27] In the case of LLMs, rationales for outputs can be provided using chain-of-thought, prompt-based reasoning techniques. In general, clinicians and patients will likely trade-off model explainability for greater accuracy, as full explainability is, in many instances, neither possible nor necessary for both clinician and patient acceptance which instead is more dependent on robust clinical validation.[28]

AI tool developers, implementers and researchers must ensure, as much as possible, the privacy and confidentiality of patient and clinician data collected and used for AI tool development and deployment and disclose where risk exists for comprehensive informed consent. Privacy means no unauthorised access to data; confidentiality means respecting patient privacy so that they will continue to seek care and discuss their problems candidly without fear of discrimination on the basis of their medical conditions. Where possible, data should be de-identified/anonymised and aggregated, contained within protected firewalls, and protected by techniques that minimise the risk of data breaches. Outlined above are suggested principles for privacy and confidentiality risk management in AI implementation and use in health care.

In general, using patient data to develop products or selling patient data should only be permissible when explicit consent has been sought from patients and either the appropriate regulatory oversight is in place or appropriately clear disclaimers are communicated as part of the consent process. This is particularly important given the category of information relied upon will either be personal or personal health information. Regulation of patient data for use in AI should be considered similarly to other technologies whilst also considering the inherent complexity and sensitivity of the data.

RACP members also noted that currently many IT systems function on a global scale implying that international sharing of data may be inevitable. They recommended that in the instance of international data sharing, it is important that transparency and consent regulations are maintained. One disadvantage of international data sharing is that shared data is more vulnerable to privacy breaches as it is hard to ensure data security between jurisdictions.

An additional area of consideration is the requirement to safeguard the health data of groups of people who have complex health conditions, those with intellectual or other disabilities / impairments who cannot provide their own consent, and children.

Concluding remarks

We thank DoHAC for this opportunity to provide feedback on this important matter.

[26] Tschandl P, Rinner C, Apalla Z, et al. Human-computer collaboration for skin cancer recognition. Nat Med 2020;26:1229–34.
[27] Ghassemi M, Oakden-Rayner L, Beam AL. The false hope of current approaches to explainable artificial intelligence in health care. Lancet Digit Health 2021;3:e745–50.
[28] van der Veer SN, Riste L, Cheraghi-Sohi S, et al. Trading off accuracy and explainability in AI decision-making: findings from 2 citizens' juries. J Am Med Inform Assoc 2021;28:2128–38

The RACP welcomes continued engagement with DoHAC and other key stakeholders to ensure that any regulatory and legislative changes are effective and supportive of physicians and that our various physician specialties can contribute their important specific expertise around legislation and regulation of AI tools in health care.

Please contact Christian White, Policy & Advocacy Officer, for questions or comments about this submission via email: policy@racp.edu.au