Australian Privacy Foundation

https://www.privacy.org.au

Secretary@privacy.org.au

https://privacy.org.au/about/contacts/

September 19, 2022

Draft RANZCR position statement – Ethical Principles for Artificial Intelligence in Medicine

The Australian Privacy Foundation (APF) is the nation's pre-eminent civil society body concerned with privacy. We appreciate this opportunity to contribute to the draft Royal Australian and New Zealand College of Radiologists (RANZCR) position statement about the Ethical Principles for Artificial Intelligence in Medicine.

Given the complex mix of factors at the heart of Artificial Intelligence (AI) development, including promising but immature and often inscrutable machine learning and algorithmic tuning parameters, we are pleased to note that people are placed at the heart of your statement. Unintended outcomes can easily proliferate in real life as a consequence of the automation of human information processing by what may be opaque algorithms and incomprehensible and large data sets used in machine learning processes.[1-2] Differences between AI training data sets and real world populations, and poorly understood modes of failure or bias, mean that clinicians should only use AI within well understood and prescribed bounds, and "monitor performance and intervene when it fails."[1]

APF feedback is formatted in accordance with that used for the position statement. It is interspersed with numbered lines of text, as our legend illustrates, in the draft document below.

Yours sincerely,

David Vaile, Chair
Australian Privacy Foundation

Dr Juanita Fernando, Vice Chair & Chair, Health Committee,
Australian Privacy Foundation

APF feedback legend

- highlighted word in draft -> highlighted word: APF comment
- *Added text*
- ~~*Deleted text*~~

The Royal Australian and New Zealand College of Radiologists

Ethical Principles for Artificial Intelligence in Medicine

The nine ethical principles outlined below guide the development of professional and practice standards regarding the research and deployment of machine learning (ML) systems and artificial intelligence (AI) tools in medicine.[3] These tools *and systems* should at all times reflect the needs of patients, their care and

> patients: *The APF ask you to consider also adding a reference to the patient's extended family. This is because the data sets and capacity to extract traits and insights from the ML system may be able to extrapolate from one person to another, especially where they are closely related. This means the impact or implications of involvement with a large scale ML system may extend beyond the individual, especially if attempts at de-identification or 'differential privacy' fail or can be easily defeated by use of other advanced analytics or data sets.*

their safety, and they should respect the clinical teams that care for them. Within this document, the term "AI tools" includes all variations of simple machine learning and complex deep learning acting as AI in clinical decision support. A full list of definitions is included in Appendix One.

> list of definitions: *The definitions are unhelpful. If you are using a mixture of ISO and other definitions, we suggest identifying into which category each definition fits, and the source document and pinpoint reference for each, so the reader can follow the trail back to the source.*
>
> *Some of these definitions are opaque or impractical in the context of unresolved concerns and philosophies about the nature, implications and failure modes of intrinsically complex and error prone ML technologies under constant development and revision. For instance, why does 'bias' mean 'systematic deviation from the truth'? What sort of truth is meant here? How is the moral, ethical or logical concept of truth useful when at issue is the degree of reliability, similarity, or parametric closeness of traits or associations derived from a range of incomplete, non-overlapping data sets related to samples of populations which may or may not be similar enough to be useful for a given technical or categorisation purpose?*
>
> *Or, for example, in which setting does 'variance' imply 'random deviation' from something?*
>
> *If these are established or ISO definitions then knowing that, and their source, may assist readers or authors to understand the context intended for their use; and to consider the degree to which the assumptions they contain can be adopted uncritically, or may warrant some caution. We are concerned that the definitions prove to be appropriate to the target audience, many of whom are not ML/AI technical experts.*

These principles are intended to guide all stakeholders involved in research or deployment of AI tools including developers, health service executives and clinicians. They are also designed to

> developers, health service executives and clinicians: *Consider adding "patients", 'data subjects" and "others potentially affected by the output of ML systems" as stakeholders, as they have a stake in some of the potential effects of such systems. For similar reasons, also consider adding Human Research Ethics Committees and similar ethical oversight bodes, as they are traditionally expected to be able to look objectively at the potential for unexpected or unintended impacts on persons known and unknown from all sorts of research involving humans.*

complement existing medical ethical frameworks (see appendices), which do not address the issues likely to emerge from use of AI in medicine.

In order to bridge this gap, the Royal Australian and New Zealand College of Radiologists (RANZCR) has developed nine ethical principles specifically to guide the following:

> nine ethical principles: *The guide should incorporate a 10^{th} principle about data protection and privacy. Principle 10 would specify adherence to data protection and privacy issues, including collection, use, disclosure, store, deletion and ostensible de-identification of all personal information utilised.*

- development of standards of practice for research in AI tools
- regulation of market access for AI tools
- development of standards of practice for deployment of AI tools in medicine
- upskilling of medical practitioners in AI tools, and
- ethical use of AI tools in medicine.

All stakeholders should take heed of all the ethical principles for AI in medicine, noting that some will have greater applicability to them.

Principle One: Safety

Although AI tools have enormous potential, a range of new risks will emerge from AI tools or through their implementation.

> implementation: *We suggest that all results must be referred to human experts for ratification. We cannot rely only on results of algorithms when identifying human disease manifestations that vary widely among people.*

The first and foremost consideration in the development, deployment or ~~utilisation~~ *use* of AI tools must be patient safety and quality of care, with ~~the~~ *an intelligible*[3] evidence base to support this.

Principle Two: Privacy and Protection of Data

Healthcare data is amongst the most sensitive ~~data~~ *information* which can be held about an individual. Patient data must not be transferred from the clinical environment[4] at which care is provided without the patient's

> transferred: *We query whether "transferred" is too limited a concept here. The Foundation suggests you consider using both transferred and "accessed from outside" the clinical setting, as this hints at the potential for API-type remote access that may not involve the wholesale "transfer" of large volumes of data.*

consent~~,~~ and approval from an ethics board or where otherwise required or permitted by law. Where data is

approval from an ethics board: The APF is not comfortable with putting "approval from an ethics board" as an alternative to consent - in many cases it would be additional to consent, with the ethics board setting out the nature of the consent required. In this instance, the word "and" should come before "approval from an ethics board".

Secondly, where there is access to identified clinical records done without informed consent, it is questionable whether it is appropriate to frame this in the context of 'privacy' for the patient because it is not based on their knowledge and acceptance, it does not fit the common usage of privacy. (It is within the legal usage of "privacy", but in Australia this is often about the ways which privacy can be breached. This document should aim higher, at an unimpeachable ethical level and avoid too-ready reliance on legal technicalities.)

permitted by law: "Permitted by law" is potentially a very broad category, especially in light of the experience of the issues and problems identified during the development, implementation and operation of the federal government My Health Record. The use of legal technicalities and opaque generic clauses are common ways in which the need for patient awareness and consent can be avoided, and should be treated as potentially hostile to patient interests.

transferred or otherwise used for *ML systems/*AI research, it must be de-identified such that the patient's identity cannot *ever* be reconstructed.

AI research: We ask what the term "AI research" means here? Is this synonymous with an ML system or an AI tool, or does it imply some different, perhaps narrower, use?

de-identified such that the patient's identity cannot ever be reconstructed: It would be useful here to acknowledge that de-identification is and will be increasingly reversible, given sufficient access to different ML tools and other data sets by which to match traits and identifying characteristics. The risk is only likely to grow over time as the advanced analytic tools and big data sets proliferate globally: the implications for the level of risk thinking and "informed consent" disclosure that are required need to be rethought — no-one can promise that any given de-identification method can withstand the attentions of "a motivated intruder" with access to other advanced tools.

Every effort must be made to store a patient's data securely*, according to differential privacy algorithms* and in line with relevant laws and *technical* best practice.

Every effort: We query whether 'every effort' is a strong enough expression. It is no longer realistically possible to protect digital networked data from "a motivated intruder" with advanced tools, so the potential for data breach and failure of security must be dealt with directly. Will the ML system/AI tool operator or host be liable when the system is breached? If so, then this is a strong incentive to apply the highest level of effort. If not, then 'good intentions' will see the risk shifted to the data subjects, patients, or others.

It should be clear from these principles whether you intend that those who host ML system data that is breached should be considered both ethically and legally responsible for the ensuing consequences, or whether you will side with them against the data subjects and in effect, accept 'a reasonable effort'/ good intentions as a means to project the risk onto the patients alone.

Data on this scale is potentially a "toxic asset" as security expert Bruce Schneier says, and this has implications such as, for example, the adoption of 'data minimisation' as the preferred means of protection, rather than 'data maximisation plus a nice fence'.

This principle needs to differentiate between the data/information used to train the AI or develop useful algorithms and the data/information obtained once the system is in use for diagnostic purposes. Unless the system is being trained in "real times" which is unlikely for most "in service" radiology systems and equipment, we have assumed that patient data is held separately from the system data that was used for training and related operations. I think we need to clarify the distinction and special sensitivity of patient data. Patient health data is especially sensitive and they should have the oversight of and final say on all health decisions.[3]

Drawing ML system designers, operators and users' attention to this necessary trade-off, and the flaws in assuming perfection in any extant data security technique, would be a useful goal of these issues and principles.

---

[1] RANZCR has adopted definitions of AI and ML (see Appendix 1)

[2] A clinical environment is any area relating to patient treatment or diagnosis and may include physical or secure virtual environments.

Principle Three: Avoidance of Bias

AI tools are limited by their algorithmic design and the data they have access to making them prone to bias. As a general rule, AI tools trained on greater volumes and varieties of data should be less biased.

> bias: *Might it be useful to have a more nuanced discussion about what 'bias' means here, what its implications might be, and how difficult it may be to avoid. The definition of systemic deviation from truth is unhelpful for encouraging an imaginative inquiry into the many ways bias and similar flaws can manifest in a given ML system used for a given purpose.*
>
> greater volumes and varieties of data: *See note above: this is an uncritical invitations to adopt somewhat-dated 'data maximalism', which trades off security for hoped for bias benefits. At the very greater volumes of data should be identified as not an unalloyed good but as something with a cost in terms of data security and privacy protection risk.*
>
> *There is also scope instead to encourage more frugal but sophisticated ways to address bias attributed to small training set size. This would help point the profession in a useful direction, rather than uncritical reliance on maximalism as the cheap and universal cure for bias.*

Moreover, bias in algorithmic design should be minimised by giving conscious consideration to avoiding

> Moreover: *"Moreover" is redundant. We suggest this sentence is put before the previous one, so that sophisticated, thoughtful design and deep understanding of the roots of bias are the focus, with data volume mentioned as one among many mitigating options, and one that, by encouraging maximalism, comes at a heavy price in terms of risk.*

bias and involving a range of perspectives and skill sets in the design process.

The data on which AI tools are based should be representative of the target patient population on which the system or tool is being used. The characteristics of the training data set and the environment in which it was tested must be clearly *and intelligibly*[3] stated when marketing an AI tool to provide

> marketing: *This seems too late a stage to address marketing. We suggest you consider a focus on the stage where sophisticated stakeholders assess or evaluate an ML system for a particular purpose instead. Any marketing or other information provided must be required to reveal all the types of information that a critical and deep assessment would need.*

transparency and facilitate implementation in appropriate clinical settings. Particular care must be taken when applying an AI tool to a population, demographic or ethnic group for which it has not been proven effective.

> population, demographic or ethnic group: *Would you consider also mentioning 'disease or syndrome' or similar? Some ML systems may be marketed as quite wide in scope of application but their zone of high reliability (false positives and negatives) may be more limited, and this effect may be as relevant as demographic 'bias'.*

To minimise risk of bias, the process, training data set and outcome measures used during development must be transparently stated.

> **during development**: *Is it only 'during development' that bias is relevant? Some systems may undergo continuous self modification using various sources, and this may result in changes over time to bias and other parameters of operation.*

Principle Four: Transparency and Explainability

AI tools can produce results which are difficult to interpret or replicate. When used in medicine, the

> difficult to interpret: *This point may be may prove erroneous. This might be the place at which you could usefully introduce the notion that ML and DL tools are in many cases effectively about 'correlation not causation', and that this may sometimes produce a run of more or less 'correct' answers but be reliant on data-derived associations, which are brittle or otherwise not reliable, and so prone to unexpected or misleading anomalies.*

medical practitioner must be capable of interpreting the basis on which a result was reached, weighing up the potential for bias and exercising clinical judgement regarding findings.

AI tools should ideally employ explainable AI (XAI) techniques *that have been previously published or documented by peers* in a way that is understandable to humans.

When designing or implementing an AI tool, consideration must be given to how a result that can impact patient care be best understood and explained by a medical practitioner.

Principle Five: Application of Human Values

The development of AI tools for medicine should ultimately benefit the patient and society. ML and AI are programmed to operate in line with a specific world view, however the use of AI tools should function

> programmed to operate in line with a specific world view: *We are not certain that "programmed to operate in line with a specific world view" is the right way to look at this issue.*
>
> *ML systems are implemented using a series of assumptions, many of which may not really be linked to "a specific world view" but may invite user over-reliance on what are inevitably flawed and imperfect attempts to categorise, analyse and identify features of clinical relevance and draw diagnostic inferences and observations from these features.*
>
> *The indirect, iterative, complex and massive-scale traits of these systems make it hard for even experts to understand how the associations were made and features extracted — so unintended, subtle, intangible or sporadic bias and other error effects (or their consequences for individual or public health care) are intrinsically hard to detect, even with XAI techniques strapped on.*

without unfair discrimination and not exacerbate existing disparities in health outcomes. Any shortcomings or risks in AI tools should be considered and weighed against the benefits of enhanced decision making for specific patient groups.

The medical practitioner must apply humanitarian values (from their training and the ethical framework in which they operate) to any circumstances in which AI tools are used in medicine, but must also consider the personal values and preferences of their patient in this situation. Entities developing AI tools must demonstrate an understanding of ethical principles*, respect for patient autonomy* and human values.

Principle Six: Decision-Making on Diagnosis and Treatment

Fundamental to quality healthcare is the relationship between the medical practitioner and the patient. The medical practitioner is the trusted advisor on complex medical conditions, test results, procedures and treatments who then communicates findings to the patient clearly and sensitively, answers questions, *explains risks and uncertainties* and ~~agrees~~ *seeks active and informed agreement* on the next treatment steps.

While AI tools can enhance decision-making capability, final decisions about care are made after a discussion between the medical practitioner and the patient taking into account the patient's presentation, history, options*, level of understanding, properly informed consent* and preferences. *The output of ML systems/AI tools should not be over-emphasized or deferred to and the potential for various limitations on their reliability or accuracy should be acknowledged where relevant.*

Principle Seven: Teamwork

ML and AI in research and medicine will need new skillsets and teams. It is imperative that in ~~the application of~~*, the deployment of ML systems/AI tools, human supervision and accountability (human warranty) is in place, both downstream and upstream if algorithms, by patients and clinicians.*[3] ~~a~~*A*ll team members *must* know *and understand* each other's strengths, capabilities and integral role in the team.

In order to deliver the best care for patients, each team member must understand the role and contribution of their colleagues and leverage them through collaboration and *human warranty*.

Principle Eight: Responsibility for Decisions Made

Responsibility for decisions made about patient care rests principally with the patient and the medical practitioner.

Medical practitioners need to be aware of the limitations of *ML systems/*AI tools, *the importance of human warranty* and must exercise solid clinical judgement at all times. However, given the multiple potential applications of *ML systems/*AI tools in the patient journey, there may be instances where responsibility is shared between:

> responsibility is shared: *The notion of 'shared responsibility' is potentially interesting but open to the clinician being tempted to rely on the output of an ML system they do not really understand or that is not understandable because they assume someone else understands or takes responsibility for it. This is problematic. The emphasis should probably be in explainable output, comprehensible to even a relatively literate patient, as an essential system/tool feature not a "nice to have".*
>
> *The clinician should over time be able to expect XAI to be the routine benchmark and not merely the ideal, and that 'tools' created for this sort of critical use must be assessed as needing the maximum level of explainability and transparency. Earlier generations of AI and ML, and those tools used for trivial matters like ad-targeting, got away without this, and without being held to a high standard of reliability.*
>
> *It should be clear from these principles that shared accountability is not really an acceptable level of patient care. When the emphasis and expectation of explainability has its effect of raising this standard, the notion of 'shared responsibility' will be less or not relevant. The manager's role will clearly be to insist on and supply XAI systems which enable clinician and patient to understand and assess the conclusions, the manufacturer's role will be to understand what is needed functionally to deliver this outcome, and the clinician's role will be to work with the patient to reach clinical decisions which the patient is happy to live with and the clinician treats as best clinical practice. The shared model embodied in the principal is a legal morass for patients looking to sue for damages related to wrongful use or incorrect diagnoses.*
>
> *Another way of putting this is that the concept of shared responsibility may, wrongly, seek to blur responsibility for the fact that many systems at present would not give either clinician or patient comfort that they knew and trusted about the way the ML system's conclusions were reached, and how reliable they were. Fudging and sharing responsibility for this is not the right approach, better to emphasise what it will take for the clinician and patient both to be happy in their central decision-making role, and that the XAI is transparent in explaining to them the basis for its conclusions.*
>
> *Accountability needs to be generally defined so that the manager is accountable for 'duty of care' in acquisition and maintenance, the operator is accountable for safe usage according to manufacturer's instructions and the clinical user of the output is accountable to ensure the quality of system output and interpretation of results.*

- The medical practitioner caring for the patient;
- The hospital or practice management who took the decision to deploy the systems or tools; and
- The manufacturer that developed the ML system or AI tool.

Although the prime responsibility regarding patient care remains with the medical practitioner, when using *ML systems/*AI tools, the responsibility is also shared by the managers of the healthcare environment and the manufacturers and developers of *ML systems/*AI tools. This potential for shared responsibility when using *ML systems/*AI tools must be identified, *practised by,* recognised by the relevant party and recorded upfront, *including in human warranties*[3], when researching or implementing *ML systems/*AI tools.

Principle Nine: Governance

ML and AI are fast moving*, immature* areas with the potential to add great value but also to do harm. The implementation of *ML systems/*AI tools requires consideration of a broad range of factors including how the ML or AI will be adopted across a hospital or practice and which patient groups will be affected and how it might align with patients' goals of care.

A hospital or practice using or developing *ML systems/*AI tools for patient care applications must have accountable governance to oversee implementation and monitoring of performance and use, to ensure the practice is compliant with ethical principles and standards. ***Device autonomy, the extent to which the device performs automated tasks independent of clinicians, and human warranty, must underpin accountable governance.***[3-4]

Broader Ethical Frameworks

Other ethical frameworks cover the expected approach and behavior of medical practitioners when delivering care to patients and provide general guidance relating to the development and adoption of new technologies in medicine.

Medical practitioners in Australia are expected to practi~~s~~*c*e in accordance with the Medical Board of Australia's Good Medical Practice: A Code of Conduct for Doctors in Australia [i] and the Australian Medical Council's Good Medical Practice.[ii]

Medical practitioners in New Zealand are expected to practi~~s~~*c*e in accordance with the New Zealand Medical Council's Good Medical Practice[iii] and the Code of Ethics set by the New Zealand Medical Association. [iv] Medical ~~P~~*p*ractitioners in New Zealand must also comply with the Code of Health and Disability Services Consumers' Rights.[v]

RANZCR has also developed a more explicit Code of Ethics for clinical radiologists and radiation oncologists.[vi]

Appendix One — **Definitions**

Technical definitions for artificial intelligence are available from the International Organization for Standardisation (ISO)[vii], general definitions are included below.

definitions: Please refer to our previous comments about the introductory sentence beneath line 8.

The draft position statement refers to ML and AI. But ML techniques can differ for AI, which is sometimes more static. So some parts of the statement may relate more or less to ML or to AI or to both. The definitions make it clear there are differences in ML, AI and other techniques. However the draft RANZCR position statement does not clarify what principles apply to a specific technique or incorporate other technqiues defined in the appendix. It is confusing. Perhaps the use of the generic term AI is more appropriate throughout this position statement?

Artificial Intelligence

"An AI system is a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. AI systems are designed to operate with varying levels of autonomy."[viii]

Explainable Artificial Intelligence (XAI)

"A set of processes and methods that allows human users to comprehend and trust the results and output created by machine learning algorithms".[ix]

Algorithm

"A series of instructions for performing a calculation or solving a problem, especially with a computer. They form the basis for everything a computer can do and are therefore a fundamental aspect of all AI systems." [x]

Bias

"A systematic deviation from the truth." [xi]

Variance

"A random deviation from the truth."[xi]

Expert system

"A computer system that mimics the decision-making ability of a human expert by following pre-programmed rules, such as 'if this occurs, then do that'. These systems fuelled much of the earlier excitement surrounding AI in the 1980s, but have since become less fashionable, particularly with the rise of neural networks. ~~Error!**Bookmark not** defined.~~

Machine learning

"One particular form of AI, which gives computers the ability to learn from and improve with experience, without being explicitly programmed. When provided with sufficient data, a machine learning algorithm can learn to make predictions or solve problems, such as identifying objects in pictures or winning at particular games, for example. ~~EFFOBro!okmark not defined.~~

Supervised Machine Learning

"A type of ML for which the algorithm changes based on data with known labels. In clinical radiology to evaluate medial images, supervised ML is a repetitive process to match images to existing labels." [xi]

Unsupervised Machine Learning

"In supervised ML the algorithm is fed an unlabelled dataset (i.e. without answers). In this case the algorithm groups the image findings into clusters based on one or more features it "learns".[xi]

Deep learning

"A more recent variation of neural networks, which uses many layers of artificial neurons to solve more difficult problems. Its popularity as a technique increased significantly from the mid-2000s onwards, as it is behind much of the wider interest in AI today. It is often used to classify information from images, text or sound." [X]

Neural network

"Also known as an artificial neural network, this is a type of machine learning loosely inspired by the structure of the human brain. A neural network is composed of simple processing nodes, or 'artificial neurons', which are connected to one another in layers. Each node will receive data from several nodes 'above' it, and give data to several nodes 'below' it. Nodes attach a 'weight' to the data they receive and attribute a value to that data. If the data does not pass a certain threshold, it is not passed on to another node. The weights and thresholds of the nodes are adjusted when the algorithm is trained until similar data input results in consistent outputs. ~~E o! Bookmarnkot defined.Error! Bookmark not defined.~~

APF Definitions

Differential privacy
A dataset preparation mechanism "that uses random noise to ensure that publicly visible information doesn't change much if one individual in the dataset changes." [3]

Human Warranty
"Evaluation by patients and clinicians in the development and deployment of ML/AI technologies … requires the application of regulatory principles upstream and downstream of the algorithm by establishing points of human supervision. If something goes wrong with ML/AI technology there should be accountability."[1]

Properly Informed Consent
Explicit, revocable, specific, and not unduly 'bundled', coerced or mandatory consent. 'Properly informed' implies that not only benefits and processes are adequately revealed in ways that are understood by the particular person being asked, but also so too are the risks, whether common but minor or rare but serious, and whether manifesting over short or long future timeframes.

i Medical Board of Australia (2014). [Internet] [Cited 2022 March 31]. Available from: https://www.medicalboard.gov.au/Codes-Guidelines-Policies/Code-of-conduct.aspx
ii Australian Medical Council (2009). [Internet] [Cited 2022 March 31]. Available from: http://www.amc.org.au/about/good-medical-practice
iii New Zealand Medical Council (2016). [Internet] [Cited 2022 March 31]. Available from: https://www.mcnz.org.nz/about-us/publications/good-medical-practice/
iv New Zealand Medical Association (2014). [Internet] [Cited 2022 March 31]. Available from: https://www.nzma.org.nz/publications/code-of-ethics
v Health and Disability Commissioner (1996). [Internet] [Cited 2022 March 31]. Available from: https://www.hdc.org.nz/your-rights/about-the-code/code-of-health-and-disability-services-consumers rights/
vi RANZCR Code of Ethics. [Internet] [Cited 2022 March 31]. Available from: https://www.ranzcr.com/documents/3958-ethics/file
vii ISO (International Organization for Standardisation) and IEC (the International Electrotechnical Commission). [internet] [cited 2022 March 31]. Available from: ISO/IEC DIS 22989(en), Information technology — Artificial intelligence — Artificial intelligence concepts and terminology
viii The Organisation for Economic Co-operation and Development (OECD). [internet] [cited 2022 March 31]. Available from: https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449
ix IBM Watson Explainable AI. [internet] [cited 2020 March 31]. Available from: https://www.ibm.com/watson/explainable-ai
x Select Committee on Artificial Intelligence. AI in the UK: ready, willing and able? [Internet] [cited 28 March 2019]. Available from: https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf
xi J. Raymond Geis, Adrian P. Brady, Carol C. Wu, et al, Ethics of Artificial Intelligence in Radiology: Summary of the Joint European and North American Multisociety Statement. *Canadian Association of Radiologists Journal* 2019 Nov;70(4):329-334

APF References

1. *Zhanglong Ji, Zachary Chase Lipton & Charles Elkan. Differential Privacy and Machine Learning: a Survey and Review. CORR abs/1412.7584 2014. https://doi.org/10.48550/arXiv.1412.7584. Cited 2022 August 15*
2. *Challen R, Denny J, Pitt M, Gompels L, Edwards T, Tsaneva-Atanasova K. Artificial intelligence, bias and clinical safety. BMJ Quality & Safety 2019;28(3):231-37 doi: 10.1136/bmjqs-2018-008370 Cited 2022 August 15.*
3. *WHO. Ethics and Governance of Artificial Intelligence for Health, WHO guidance. World Health Organisation (WHO), 2021. Cited 2022 August 13, https://apps.who.int/iris/bitstream/handle/10665/341996/9789240029200-eng.pdf.*
4. *Lyell D, Coiera E, Chen J, Shah P, Magrabi F. How machine learning is embedded to support clinician decision making: an analysis of FDA-approved medical devices. BMJ Health Care Inform. 2021 Apr;28(1):e100301. doi: 10.1136/bmjhci-2020-100301. PMID: 33853863; PMCID: PMC8054073. Cited 2022 August 13.*