Australia's National Science Agency

AI trends for healthcare

March 2024

Citation and authorship

CSIRO (2023) AI trends for healthcare.
CSIRO, Herston. This report was authored by David Hansen, Denis Bauer, John Grimes, David Silvera, Liesel Higgins, Naomi Stekelenburg, Pierrick Bourgeat, Priya Ramarao-Milne, Hang (Hollie) Min, Rosita Shishegar, Teresa Wozniak, David Ireland and Filip Rusak.

Copyright

Important disclaimer

AEHRC

The Australian e-Health Research Centre (AEHRC) is the largest digital health research program in Australia with over 150 scientists and engineers and a further 50 higher degree research students. As CSIRO's national digital health research program, the AEHRC has offices across Brisbane, Sydney, Melbourne, Canberra, Adelaide, and Perth. AEHRC is unique world-wide in covering the full value chain in health care, from basic science all the way to delivering technology and services into the healthcare system.

Acknowledgements

CSIRO acknowledges the Traditional Owners of the lands that we live and work on across Australia and pays its respect to Elders past and present. CSIRO recognises that Aboriginal and Torres Strait Islander peoples have made and will continue to make extraordinary contributions to all aspects of Australian life including culture, economy, and science.

The project team is grateful for the time and input of the stakeholders from industry, government and academia who were consulted throughout this project. The team is particularly thankful to Dana Bradford, David Hansen and Bevan Koopman who were the authors of the first AI Report (2019), from which parts of the 'Primer' section of this report are taken.

Special thank you to our colleagues at CSIRO who provided invaluable contributions to the report including draft reviews.

Contents

Introduction

CSIRO’s Australian e-Health Research Centre (AEHRC) plays a pivotal role in creating, analysing, and implementing Australia’s digital health infrastructure.

In 2023 artificial intelligence (AI) technologies have started to move into mainstream use. Tools such as ChatGPT have provided a way that almost everyone can start to interact with AI technologies – and many people are finding new and novel ways of using the technology. In many ways this stage of the use of AI in society mirrors the arrival of Google in the late 1990s, which really marked the beginning of widespread use of the internet.

AI use across many industries is not new – and that includes healthcare. However, so far, such use has been limited and specific. Three factors have led to the inevitability of more broad use of AI in everyday tasks:

- the escalation in the amount of digital data
- the surge in compute power through cloud computing
- availability of AI tools that allow the reimagining of task management.

We originally published our AI in Health Report in 2020. One of the key motivations for the report was to detail how we were mobilising AI in our health and medical applications. Across 34 different case studies we detailed the many different AI techniques we used to solve health and medical challenges.

Over the past three years our use of AI has only increased, and the emergence of new techniques has been incorporated into many of our technologies.

A key difference between the use of AI in healthcare and the use of AI in other industries is where AI provides decision making for diagnosis, prevention, prediction, prognosis, monitoring or treatment. In these cases, the AI is considered a medical device and is currently regulated as such – ‘software as a medical device (SaMD)’. This feeds into a bigger discussion of the use of AI across healthcare, for clinical and non-clinical purposes, and ensuring Australia is ready for its use. The healthcare consequences of the rise of generative models are rapidly unfolding and the national discussion about how to regulate AI is gaining pace.

Australia's largest
digital health research program

Full health and
biomedical informatics
research program

Est.
2003

>150
scientists and
engineers

50
higher degree
research students

Australia is undoubtedly set to benefit enormously from AI in healthcare – we are already seeing AI algorithms performing tasks like chatbots and image processing. However, we also need to ensure our nation's preparedness for the unintended consequences of AI in the health sector.

We have identified four trends in digital health – each being driven by the increasing digitisation of society and the increased willingness of doctors, researchers, and patients to interact using digital tools. All support the use of AI and machine learning in everyday life, including healthcare.

- Interoperability: ensuring that data can be safely exchanged between systems to support patient care and increased system performance.
- Cloud: cloud computing is increasing being used by health systems for secure data storage and enabling data exchange as well as for high performance computing.
- Apps and personalisation: the availability of patient data is supporting a more personalised approach to providing healthcare.
- Data analytics as a service: as well as the availability of data on cloud computing is enabling algorithms to be 'brought to the data' rather than the data being shared.

To fully harness AI and machine learning (ML) we need not to just let it happen, but rather plan for its introduction into healthcare. This means we will be able to benefit properly from AI by ensuring the frameworks are firmly in place for ethical implementation and that the safety, quality and monitoring guidelines are established as we strive to create newer and better AI based digital tools.

Rather than an update on our 2020 report, this report provides an update on specific areas of the use of AI across the AEHRC and then provides a deeper treatment of how we use AI through seven (7) case studies. As such this report is an adjunct to our 2020 report rather than a replacement. We trust that it provides valuable information to the reader and prepares clinicians and consumers alike for the next phase of the digital health transformation.

David Hansen
CEO and Research Director,
CSIRO's Australian e-Health
Research Centre

1

Use of AI and ML by AEHRC Research Groups

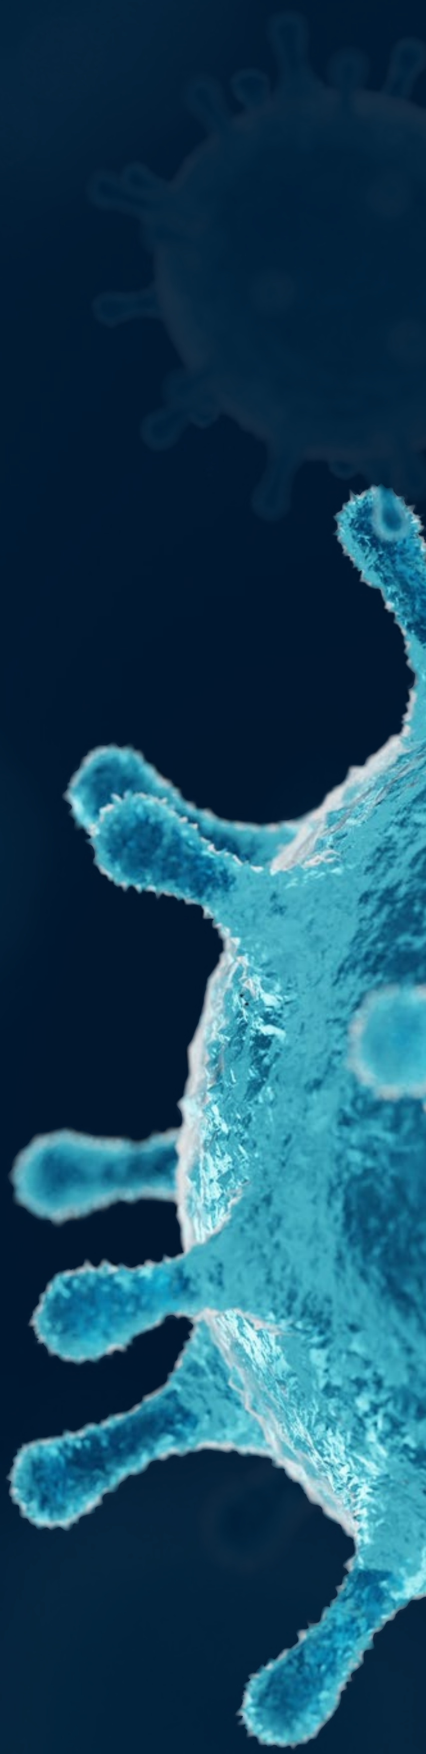

AEHRC's 150 scientists and engineers and over 50 research students work across five research groups. The groups have each developed multiple platform technologies and digital health solutions that enable them to work together and with stakeholders and collaborators to tackle the key challenges of contemporary healthcare.

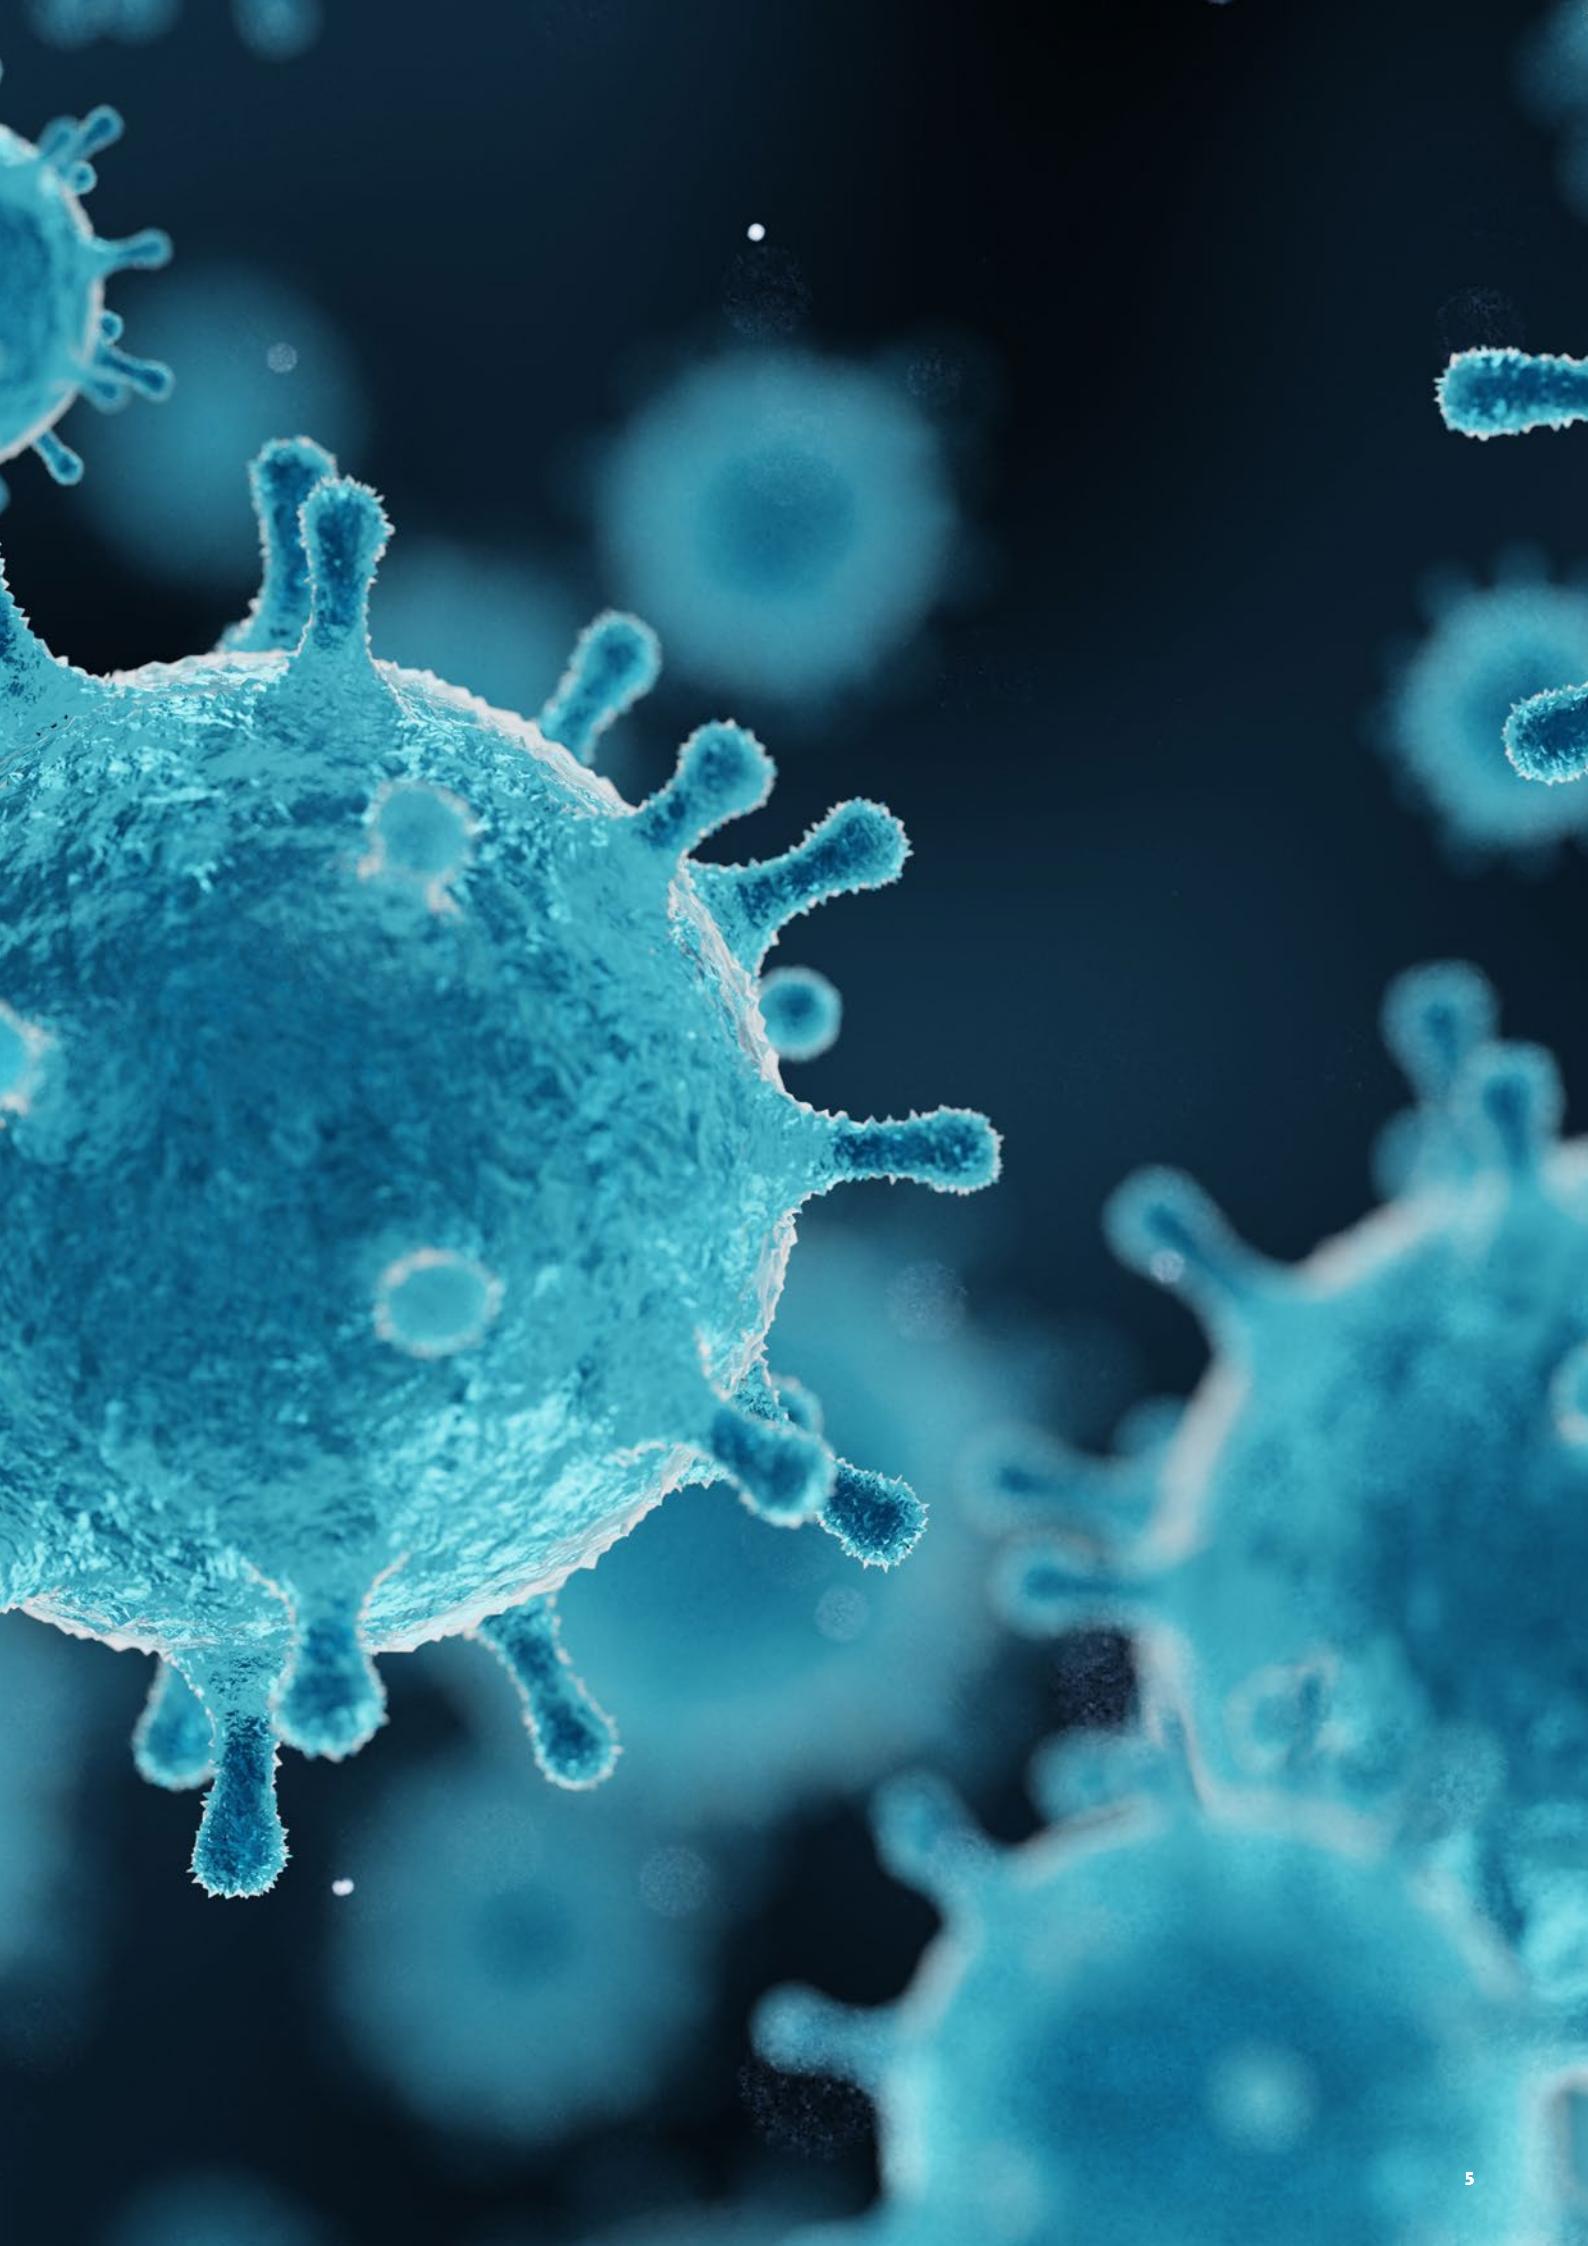

Health System Analytics Group

The Health Systems Analytics group develops AI and ML based software tools to increase health system productivity, improve patient safety, and deliver higher quality care. The software tools optimise patient, clinician, and resource level flows, and simultaneously provide intelligent decision support.

Health Data Semantics and Interoperability Group

The Health Data Semantics and Interoperability group improves patient outcomes and health system performance and productivity through data collected in EMRs and other clinical information systems.

The group works closely with two key international standards, SNOMED CT – the international standard clinical terminology – and HL7's FHIR® to improve interoperability across digital health systems. The group also uses natural language processing to process medical narratives such as pathology reports or medical literature to support interoperability, clinical decision making and clinical research.

Transformational Bioinformatics Group

The Transformational Bioinformatics group are world leaders in cloud-native genomics research, using the latest in ML and bioinformatics to drive innovation in the use of genomics in the health system. The group applies AI and ML across two main genomics disciplines – genomics insights and genome therapeutics.

Biomedical Informatics Group

The Biomedical Informatics group develops innovative medical technologies for the discovery and communication of meaningful patterns in biomedical data. Especially valuable in areas rich with recorded information such as medical images or genomics, these technologies rely on the simultaneous application of statistics, computer programming, and applied mathematics. The group also develops techniques to report and visualise complex biomedical information for clinical diagnosis and screening and to communicate insights to clinicians and clinical researchers.

Digital Therapeutics and Care Group

The Digital Therapeutics and Care group takes advantage of emerging sensor systems, digital technologies, data access, and analytics to improve healthcare for people who are chronically ill , older people and those living with disability. The multi-disciplinary group combines expertise in clinical research, tele-medicine systems and AI for medical image and data analysis to develop cutting-edge digital platforms.

A big win for AEHRC at MEDINFO202

This year at one of the world's biggest digital health conferences, AEHRC's Aida Brankovic received a special accolade for her research, which she presented at the conference.

The MedInfo 2023 Best Paper Award was presented to Aida, Wenjie Huang, David Cook, Sankalp Khanna, and Konstanty Bialkowski for their paper 'Elucidating Discrepancy in Explanations of Predictive Models Developed Using EMR'. The award recognised the researchers' paper as the highest ranking of the 281 papers presented at the conference.

Aida's most recent project has shown the potential benefits of using an AI algorithm to improve engagement and health outcomes from digital health programs. The study, published in The Journal of Medical Internet Research, used a CSIRO-developed algorithm that uses AI to predict when a person will drop out of an online weight loss program.

Award winner Aida Brankovic contemplates her bright future in AI.

2

Breaking down AI basics

Despite its prevalence, the term AI is still often misused and misunderstood. That's why before we get into unfolding some of the future directions and trends for AI in healthcare, a short guide to some of the vocabulary and concepts of AI may be beneficial.

Given the omnipotence of AI in the healthcare sector, AI literacy is an increasingly essential part of the clinician, researcher, and industry partner toolkits.

Scientific domains of AI

AI techniques can be described across various components or scientific domains. In this report we focus on the four different domains highlighted in the figure below.

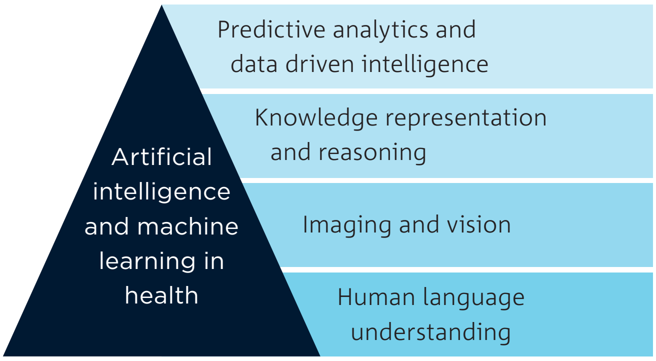

Domains of artificial intelligence techniques.

Predictive analytics and data-driven intelligence

Data often exists in very large sets, often too big to be analysed by humans. Predictive analytics and data driven intelligence use the compute power of machines to extract insights from existing data. In this field, the intent is for insights to be bottom-up. This means trends are identified from low-level data. The advantage of this is a hypothesis-free, unbiased examination of the data to identify underlying patterns. An example is CSIRO's VariantSpark tool, which processes millions of genetic variables and extracts both linear and non-linear information to identify variables that explain complex phenotypes.

Knowledge representation and reasoning

Computers can help humans infer knowledge by solving complex tasks. Before this can happen, humans need to represent or classify information about our world in a form that can be read by a computer system. In healthcare, this is typically about representing medical concepts (such as diseases) and their properties and relationships in a machine readable and understandable form.

In many instances, solving the knowledge representation problem is the pivotal challenge. Once the data is represented in the right form, the problems become 'tractable'. That is, they're able to be processed using compute power in an appropriate timescale.

Imaging and vision

This domain harnesses the power of images and videos for insight into the cause and impact of medical conditions. Computer vision and image processing are two areas that have been transformed by new AI methods, particularly deep learning-based methods.

Deep learning allows us to 'train' a machine to 'read' an image and in some instances to produce its own. Machine created medical imagery can act as a baseline for measurements when no other baselines are available.

Medical image captioning lays the groundwork for multimodal medical image analysis tools that can assist with clinical documentation. It has the potential to lead to medical imaging tools that can maintain and improve the consistency, quality, and efficiency of clinical reporting, produce a rich textual description from medical images, provide a fast, inexpensive second reader, and help reduce teaching time.

Human language understanding

One of the main forms of data from humans comes as language – or 'natural language'. Although in medical settings we aim to standardise data and make it machine readable, the AI we use needs to be adept at extracting meaning, searching, summarising, and classifying human language.

Developing AI able to do this is one of the key objectives of research in human language understanding.

In the same way that AI is categorised in terms of its objectives or domains, it can also be categorised in terms of its technologies, which generally falls in one of the two contrasting approaches – symbolic and statistical. Both paradigms play an equally important role, especially for AI solutions developed for domains such as healthcare.

AEHRC post-doc wins international imaging competition

Interpreting and summarising information in medical images is a time-consuming task and a significant resource burden on the health system.

Recently, a post-doc in AEHRC's medical imaging analysis team, Aaron Nicolson, won an international ImageCLEF award for his medical imaging captioning work, where he developed the groundwork for a multi-modal image analysis tool that can assist with clinical documentation.

Medical imaging tools like this will, once approved and deployed, maintain and improve the consistency, quality, & efficiency of clinical reporting, produce a rich textual description from medical images, provide a fast & inexpensive second reader, and help reduce teaching time.

> What was the competition about?
> imageclef.org/2023/medical/caption

ImageCLEFmedical is an international competition focussed on the development of machine learning methods for several important computer aided diagnostic tasks. One such task is automatic medical image captioning, where participants are tasked with generating a synthetic description of a given medical image, where the image could be one of many modalities, e.g., radiography, ultrasonography, computed tomography (CT), magnetic resonance imaging (MRI), etc. The generated captions were evaluated using several automatic metrics.

Team Name	Run	BERTScore	ROUGE	BLEURT	BLEU	METEOR	CIDEr	CLIPScore
CSIRO	4	**0.643**	0.245	0.314	0.161	0.080	0.203	**0.815**
closeAI2023	7	0.628	0.240	**0.321**	0.185	0.087	**0.238**	0.807
AUEB-NLP-Group	2	0.617	0.213	0.295	0.169	0.072	0.147	0.804
PCLmed	5	0.615	0.253	0.317	0.217	0.092	0.232	0.802
VCMI	5	0.615	0.218	0.308	0.165	0.073	0.172	0.808
KDE-Lab Med	3	0.615	0.222	0.301	0.156	0.072	0.182	0.806
SSN MLRG	1	0.602	0.211	0.277	0.142	0.062	0.128	0.776
DLNU CCSE	1	0.601	0.203	0.263	0.106	0.056	0.133	0.773
CS Morgan	10	0.582	0.156	0.224	0.057	0.044	0.084	0.759
Clef-CSE-GAN-Team	2	0.582	0.218	0.269	0.145	0.070	0.174	0.789
Bluefield-2023	3	0.578	0.153	0.272	0.154	0.060	0.101	0.784
IUST NLPLAB	6	0.567	**0.290**	0.223	**0.268**	**0.100**	0.177	0.807
SSNSheerinKavitha	4	0.544	0.087	0.215	0.075	0.026	0.014	0.687

CLEF scoreboard, the primary metric used to rank participants is highlighted in grey.

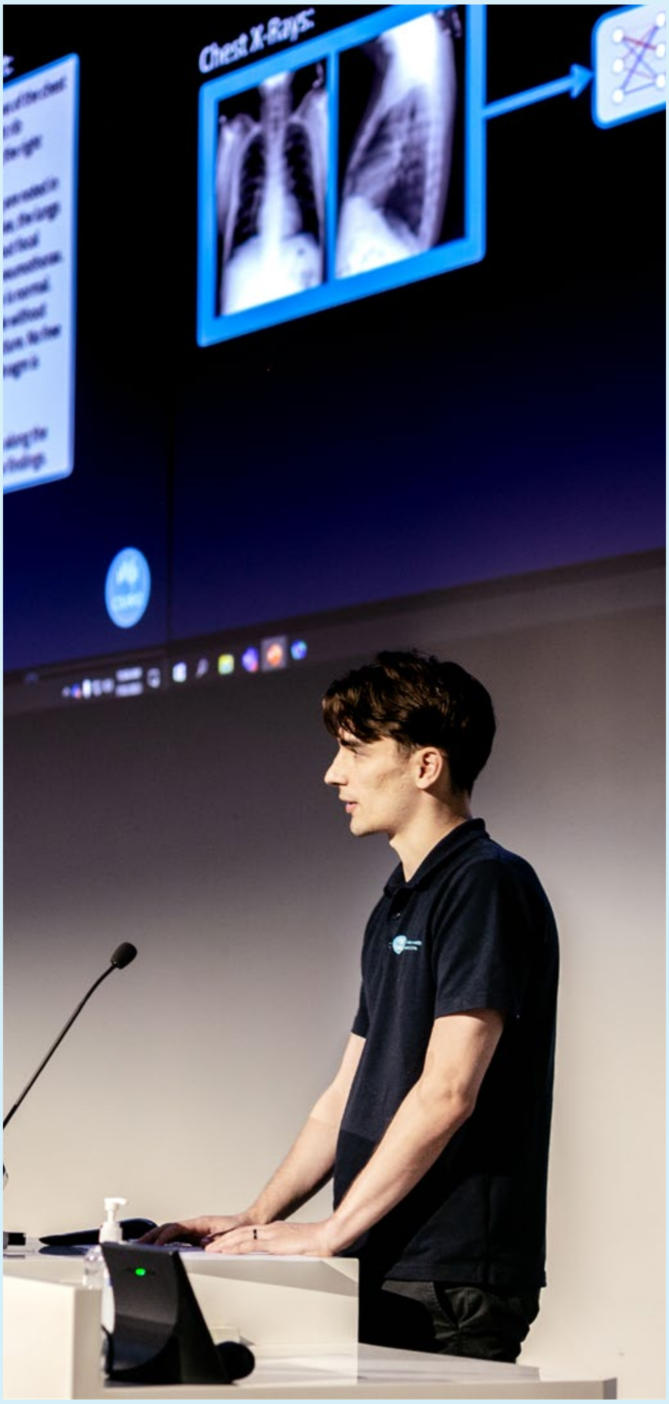

Dr Aaron Nicolson, winner of the ImageCLEF competition 2023.

Symbolic or statistical artificial intelligence?

In AI, there have traditionally been two schools with contrasting approaches – symbolic and statistical.

Symbolic AI

Symbolic AI refers to the representation or encoding of human knowledge into a form of known facts and/or rules, typically known as ontologies. These facts or rules are then used with data to reason, or infer, other facts from the data.

Statistical AI

Statistical AI takes the opposite approach – rather than predefining the knowledge and rules, it 'learns' these from data. This approach uses existing data and evidence along with computational techniques to extract patterns and insights to reason about the world.

Data, data models, quality, and standards

Machine learning

The 'intelligence' component of AI depends on data.

A simple analogy is to consider how you might go about training your dog. You could provide the dog multiple opportunities to perform the task. These opportunities would exist as data points in the dog's learning. The dog would use the data points to decide how to respond to a command you give.

In the same way, the intelligence component of AI uses data points to make decisions based on commands.

Using a dog training analogy, if you were consistent in your response to the opportunities for learning, the dog would have a strong set of data points on which to base their 'intelligence'. For example, if every time your dog sat you gave her a treat, she would learn the connection between sitting and getting a treat. If on the other hand you gave your dog poor opportunities for learning, for instance your communication was not clear or non-existent, you were inconsistent in your responses or you used unpredictable stimulus for the learning opportunities, your dog may not learn as quickly or as well.

Similarly, the quality of the data used to either train AI models or for AI based analysis has a direct impact on the quality of outputs and downstream tasks. Quality data for healthcare is defined by data that is accurate, complete, timely, and fit for purpose.

Machine learning uses data to give computers the ability to learn without being explicitly programmed.

> Traditional ML includes techniques such as regression, decision trees, support vector machine, Bayes network, K-nearest neighbour, principal component analysis. These techniques analyse various types of data – numerical, categorical, binary, time series, text, image, audio, and video data – to identify relationships in the data.

Classification vs regression

There are two main ML tasks: classification and regression.

Classification involves using a ML model to 'classify' some data according to a finite set of categories; for example, classifying the type of cancer found in a pathology report: breast, lung, etc. The simplest case being a binary classification – yes/no, true/false, cancer/not cancer, etc.

Regression, in contrast, involves using a ML model to predict a continuous value rather than a category. For example, predicting length of stay for a patient given their condition.

Most ML models perform either regression or classification; however, there are models that can be implemented to cover both.

Three of the main machine learning techniques are:

- statistical ML aims to find some type of predictive function from the training data
- reinforcement learning approaches provide AI algorithms with ‘rewards’ or ‘penalties’ based on their problem-solving performance
- deep learning approaches make use of artificial neural networks.

Health data

All models depend on health data.

Health data comes in various forms. As many AI approaches are linked intrinsically to the data type, the table below outlines some of the more common types of health data (Table 1).

Table 1: Categories of health data

DATA TYPE	DESCRIPTION	FORMAT AND STANDARDS	AVG. SIZE/ PATIENT
Clinical Data	A wide range of data used by health organisations: • patient records and is what large electronic health record systems would store. • laboratory data (e.g. pathology and radiology reports).	Includes unstructured (free-text narratives) or highly structured (e.g. intensive care unit (ICU) observations) data. There is a strong drive to develop standards for clinical data: • HL7 for messaging/transfer of clinical data • FHIR to replace many older HL7 for better patient data interchange • SNOMED CT – an ontology for medical knowledge representation • International Classification of Diseases (ICD) – a classification for diagnostic coding • Human phenotype ontology for describing features associated with genetic disease	100s MB
Genomics data	Growing source of data as precision medicine becomes more commonplace. Includes single gene tests, panel tests, whole exome sequencing, whole genome sequencing. With precision medicine, the boundary between clinical and genomic data will narrow.	Genomic data can be provided in multiple forms of processing stages each with its own set of standards or uniquely structured data, ranging from raw sequence data such as FASTA (consensus genome sequence) and FASTQ (sequencing reads), processed data such as SAM/BAM (aligned reads), to tertiary derivatives such as VCF (genomic variant information) or unique structures for annotated VCFs, methylation, transcription etc. This is more dominant in the non-human domain such as metagenomics or pathogenomics.	138 GB (whole genome, FASTQ) 10s-100s GB/ cohort (VCF)
Imaging	Medical images typically arising from radiology. Imaging results from MRI, CT, X-ray, PET etc.	Standards such as DICOM exist to capture medical images; however, these are still pictures (or 3D images) represented with pixels.	Modality dependent, 100s MB–GB
Administrative data	Data not associated directly with a patient's medical condition. Includes billing, insurance, financial data, and efficiency metrics and patient flow numbers.	A mix of structured and unstructured. Use of some terminology systems such as ICD-10.	100s MB
Sensor and wearables data	Data provided by sensors. Sensors vary widely from ICU observations through to wearables in the home. Rapidly growing areas with the Internet of Things (IoT).	Often numeric, structured and time-series based. Can be image, sound, number, or text.	Varies greatly

Supervised vs unsupervised learning

ML models learn from data, either supervised or unsupervised.

Supervised learning means that the computer is provided with the correct answers or labels along with the data. For example, the is provided with lung X-rays labelled with either ‘cancerous’ or ‘not cancerous’ and then learns characteristics of the image that indicate these classifications.

Unsupervised learning means the computer is not directed by any manually provided answers or labels. An example would be clustering algorithms that group people’s genome by their ethnic backgrounds.

Deep learning

Deep learning is an approach that uses artificial neural networks for either classification or regression, both supervised and unsupervised.

The name deep learning comes from the fact that their architecture has many ‘deep’ layers of neural networks. These can capture richer and more high-level features in the data they are provided; for example, deep learning for facial recognition will capture both coarse grained location of the eye within the face as well as more fine-grained parts of the eye.

Deep learning shows impressive results in a variety of domains but is particularly well suited to image processing and speech processing.

Predictive analytics and data driven intelligence

Predictive analytics and data driven intelligence refer to the analysis of data and the discovery of patterns that provide novel insights to inform workflow and decision making in the chosen business domain.

These are often used in the context of big data and employ AI and ML techniques to extract meaning and knowledge from this data to discover relationships and trends, forecast more accurately and reliably, guide informed decision-making, and optimise business operations.

Predictive analytics and data-driven intelligence might use health data to help improve capacity management and patient flow through the health system, develop patient-centred evidence-based models of care, address the burden of chronic disease, health monitoring and management of home-based care.

Knowledge representation and reasoning

Knowledge representation and reasoning is a core tenant of AI and typically an example of symbolic AI. An intelligent system needs to represent information and knowledge in a form that computers can understand to solve complex problems such as diagnosing a medical condition or understanding natural language. Once knowledge is appropriately represented, reasoning systems can then apply this knowledge in new situations, to acquire or infer novel knowledge.

From a knowledge representation perspective, health is quite unique. This is because considerable effort has been made to capture and represent medical knowledge in a standardised and machine-reasonable manner. Today, some of the largest domain-specific knowledge representation systems are health based.

CSIRO's AEHRC has a long history of development of tools that support the adoption and use of SNOMED CT. This includes a reasoning engine able to use SNOMED CT to infer new knowledge, a terminology service that allows third parties to easily look up, retrieve and leverage SNOMED CT concepts, and support for analytics of clinical data captured as SNOMED CT codes in electronic health records or other clinical information systems.

Example

While there several terminologies for representing medical knowledge, one that adheres the formal logic mentioned above is the SNOMED CT ontology. In SNOMED CT, knowledge representation is achieved through three core components:

- Concepts: a specific medical concept identified through a unique numerical code e.g. 74400008 is the code for appendicitis.
- Descriptions: textual descriptions for a concept, for which there may be multiple e.g. the descriptions 'heart attack' and 'myocardial infarction' both pertain to concept 22298006.
- Relationships: which connect concepts together e.g. 'appendectomy' can actually be represented as a relationship between the 'appendix' and 'excision' concepts.

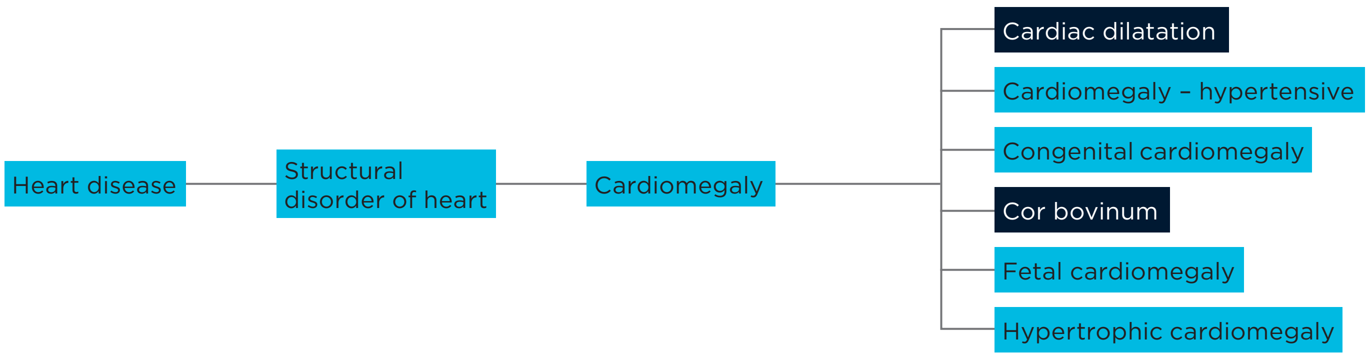

A sample portion of the SNOMED hierarchy as visualised by the CSIRO Shrimp browser for cardiomegaly.

Information models

The easiest way to think of an information model is of a simple form for collecting data. In this case the form's fields provide the information model while the data collected in each field might be from a standard terminology, or free text etc.

For instance, there might be a field to capture diagnosis and the data captured in the form should come from the 'clinical finding' hierarchy of the SNOMED CT terminology. In healthcare, HL7 is the international standards body for standards supporting the transfer of clinical and administrative health data. Their standards, including HL7 v2, CDA (Clinical Document Architecture) and FHIR (Fast Healthcare Interoperable Resources), are the bedrock information models used for sharing information in healthcare. The information model – whether it is a HL7 standard or one of many others – provides the 'meta-data' for any data used in AI or ML applications.

Human language understanding

Humans communicate in natural language, which is ambiguous, does not follow strict syntactic or semantic rules and is therefore difficult for machines to understand. In response, there has been a strong emphasis to avoid using natural language and adopt standard terminologies (e.g. SNOMED CT and ICD).

However, for certain tasks – such as communication between humans – natural language remains the most effective format. In response, computational methods for natural language processing (NLP) and, more generally, natural language understanding have emerged.

With the advent of new deep learning models for NLP, considerable improvements have been made in the field.

Imaging

Imaging and vision are key applications of AI – used to extract information from images and to make decisions. Imaging applications can range from using images taken by regular cameras through to images acquired by advanced imaging machines typically used in healthcare. In the case of visual applications, images are processed in real time and used for applications in robotics, which are increasingly deployed in medical applications from surgical robots to social robotics.

QUICK SCIENCE

From mobile health to wearable devices

AEHRC has been at the forefront of mobile health research since 2014, when we conducted the world's first validation of a mobile health enabled cardiac rehabilitation program. Under- utilisation of cardiac rehabilitation results in adverse patient outcomes and increased hospital admissions. We aimed to improve rehabilitation participation by providing an alternative to standard in-clinic care through mobile app delivered cardiac rehabilitation – including physical exercise, counselling, health monitoring and education support. The trial, with 120 participants, clearly demonstrated the acceptability of this option, with greater uptake, adherence, and a 30% increase in completion of cardiac rehabilitation compared to an in- clinic delivered program.

Since then, the AEHRC has undertaken several trials in mobile health using a similar approach, while also expanding our research to include other sources of data beyond mobile phone apps. This includes sensors in homes and integration with medical devices to monitor and support the health and wellbeing of aged people living alone. We have developed wearable sensors for babies to monitor movement for early detection and diagnosis of conditions such as cerebral palsy. We have also investigated chatbots to provide monitoring and support for a wide range of physical and mental health conditions.

These innovations and integration in mobile health, connected devices, and wearables shift the paradigm of tele-health and tele-medicine using video conferencing to 'virtual care', in which access to healthcare and data is pervasive, and can be delivered to the patient at home or in the community.

Increasingly these models of care use AI to analyse the data being captured in real time. The novel analytics developed using AI span monitoring and diagnosis to prediction of likely events at individual level, as well as providing an aggregated view of healthcare.

3

Safe and Responsible AI

With the rapid uptake of AI systems by health SMEs, hospital systems and clinicians, it is vital for the health sector, perhaps more than any other, to ensure responsible and ethical implementation.

Examples of frameworks for responsible AI

Among the multiple examples of frameworks for responsible AI is the Australian Government's Department of Industry, Science and Resources has published an artificial intelligence ethics framework comprising 8 principles to help ensure AI is safe, secure, and reliable.

Examples of AI Frameworks:

AU: https://www.industry.gov.au/publications/australias-artificial-intelligence-ethics-framework/australias-ai-ethics-principles

EU: https://artificialintelligenceact.eu/

USA: https://www.nist.gov/itl/ai-risk-management-framework

Canada: https://ised-isde.canada.ca/site/innovation-better-canada/en/artificial-intelligence-and-data-act-aida-companion-document

As the application and technological advancements of AI continue to expand, a range of ethical concerns has surfaced regarding its implementation, especially in human-sensitive environments such as healthcare.

AI is revolutionising research and industry by expanding into previously human-exclusive areas, bringing breakthroughs and potential harms to individuals and groups including discrimination, misinformation, polarisation, deepfakes, scams and cyber-attacks. Such technological disruption calls for governance to ensure human-safe and responsible usage of AI.

Responsible AI encompasses a developing interdisciplinary field that delves into the ethical implications of AI, aiming to identify, comprehend and regulate the ethical obligations associated with this fast-evolving technology. The focus of responsible AI lies in the conscientious development, deployment, and operation of AI systems that adhere to ethical standards, promote transparency, and establish governance and accountability. By doing so, responsible AI seeks to mitigate biases and foster fairness, equity, and equality, with the ultimate goal of ensuring the safe usage of AI systems and the protection of the human rights of the users involved.

Responsible AI at CSIRO

CSIRO's National AI Centre (NAIC) is at the forefront of developing and encouraging the adoption of responsible AI practices for industry. This work is being carried out principally by NAIC's Responsible AI Network. Together with their partners and collaborators, the Responsible AI Network provides curated guidance and advice for the development of a world-class capability in Responsible AI. This is done via seven actionable pillars:

Law | Standards | Principles | Governance | Leadership | Technology | Design

Over the past decade, considerable efforts have been made to reconcile ethical considerations with the transformative potential of AI, resulting in various AI ethics guidelines globally. These ethics guidelines are mainly constructed around human rights and risk mitigation for their violation. While various ethical AI guidelines were proposed by reputable international organisations such as UNESCO, World Economic Forum (WEF), Organisation for Economic Co-operation and Development (OECD) and IEEE, as well as national governments including EU, USA and Canada, the operationalisation of such guidelines is still unclear and yet to be defined. The operationalisation of responsible and ethical AI guidelines requires AI regularisation, the establishment of standards, appointments of advisory boards and ethics officers, the formation of processes for trustworthy AI assessments as well as raising awareness, community participation and building capabilities. While the operationalisation of the existing frameworks is not landed yet, it is important to create inclusive, interdisciplinary, and ongoing discussions across sectors to shape the regulatory processes and to size the potential of AI in a safe way.

While AI has immense potential to improve wellbeing, quality of life and grow our economy, the current regulatory framework likely does not sufficiently address known risks presented by AI systems, particularly in high-risk settings. It is crucial to ensure the design, development and deployment of AI systems in Australia in legitimate, but high-risk settings, is safe and can be relied upon, while ensuring the use of AI in low-risk settings can continue to flourish largely unimpeded. It is recognised that we need to design modern laws for modern technology. Harms can be prevented from occurring through testing, transparency, and accountability measures, clarifying and strengthening laws to safeguard citizens, working internationally to support the safe development and deployment of AI.

While the operationalisation of the existing frameworks is not landed yet, it is important to create inclusive, interdisciplinary, and ongoing discussions across sectors to shape the regulatory processes and to size the potential of AI in a safe way.

Australia’s AI Ethics Principles at a glance

Human, societal and environmental wellbeing

AI systems should benefit individuals, society and the environment.

Human-centred values

AI systems should respect human rights, diversity, and the autonomy of individuals.

Fairness

AI systems should be inclusive, accessible, and should not involve or result in unfair discrimination against individuals, communities, or groups.

Privacy protection and security

AI systems should respect and uphold privacy rights and data protection and ensure the security of data.

Reliability and safety

AI systems should reliably operate in accordance with their intended purpose.

Transparency and explainability

There should be transparency and responsible disclosure so people can understand when they are being significantly impacted by AI and can find out when an AI system is engaging with them.

Contestability

When an AI system significantly impacts a person, community, group or environment, there should be a timely process to allow people to challenge the use or outcomes of the AI system.

Accountability

People responsible for the different phases of the AI system lifecycle should be identifiable and accountable for the outcomes of the AI systems, and human oversight of AI systems should be enabled.

4

Challenges and opportunities

The digital transformation of the healthcare system carries with it multiple opportunities and challenges. We cover some of them in this section.

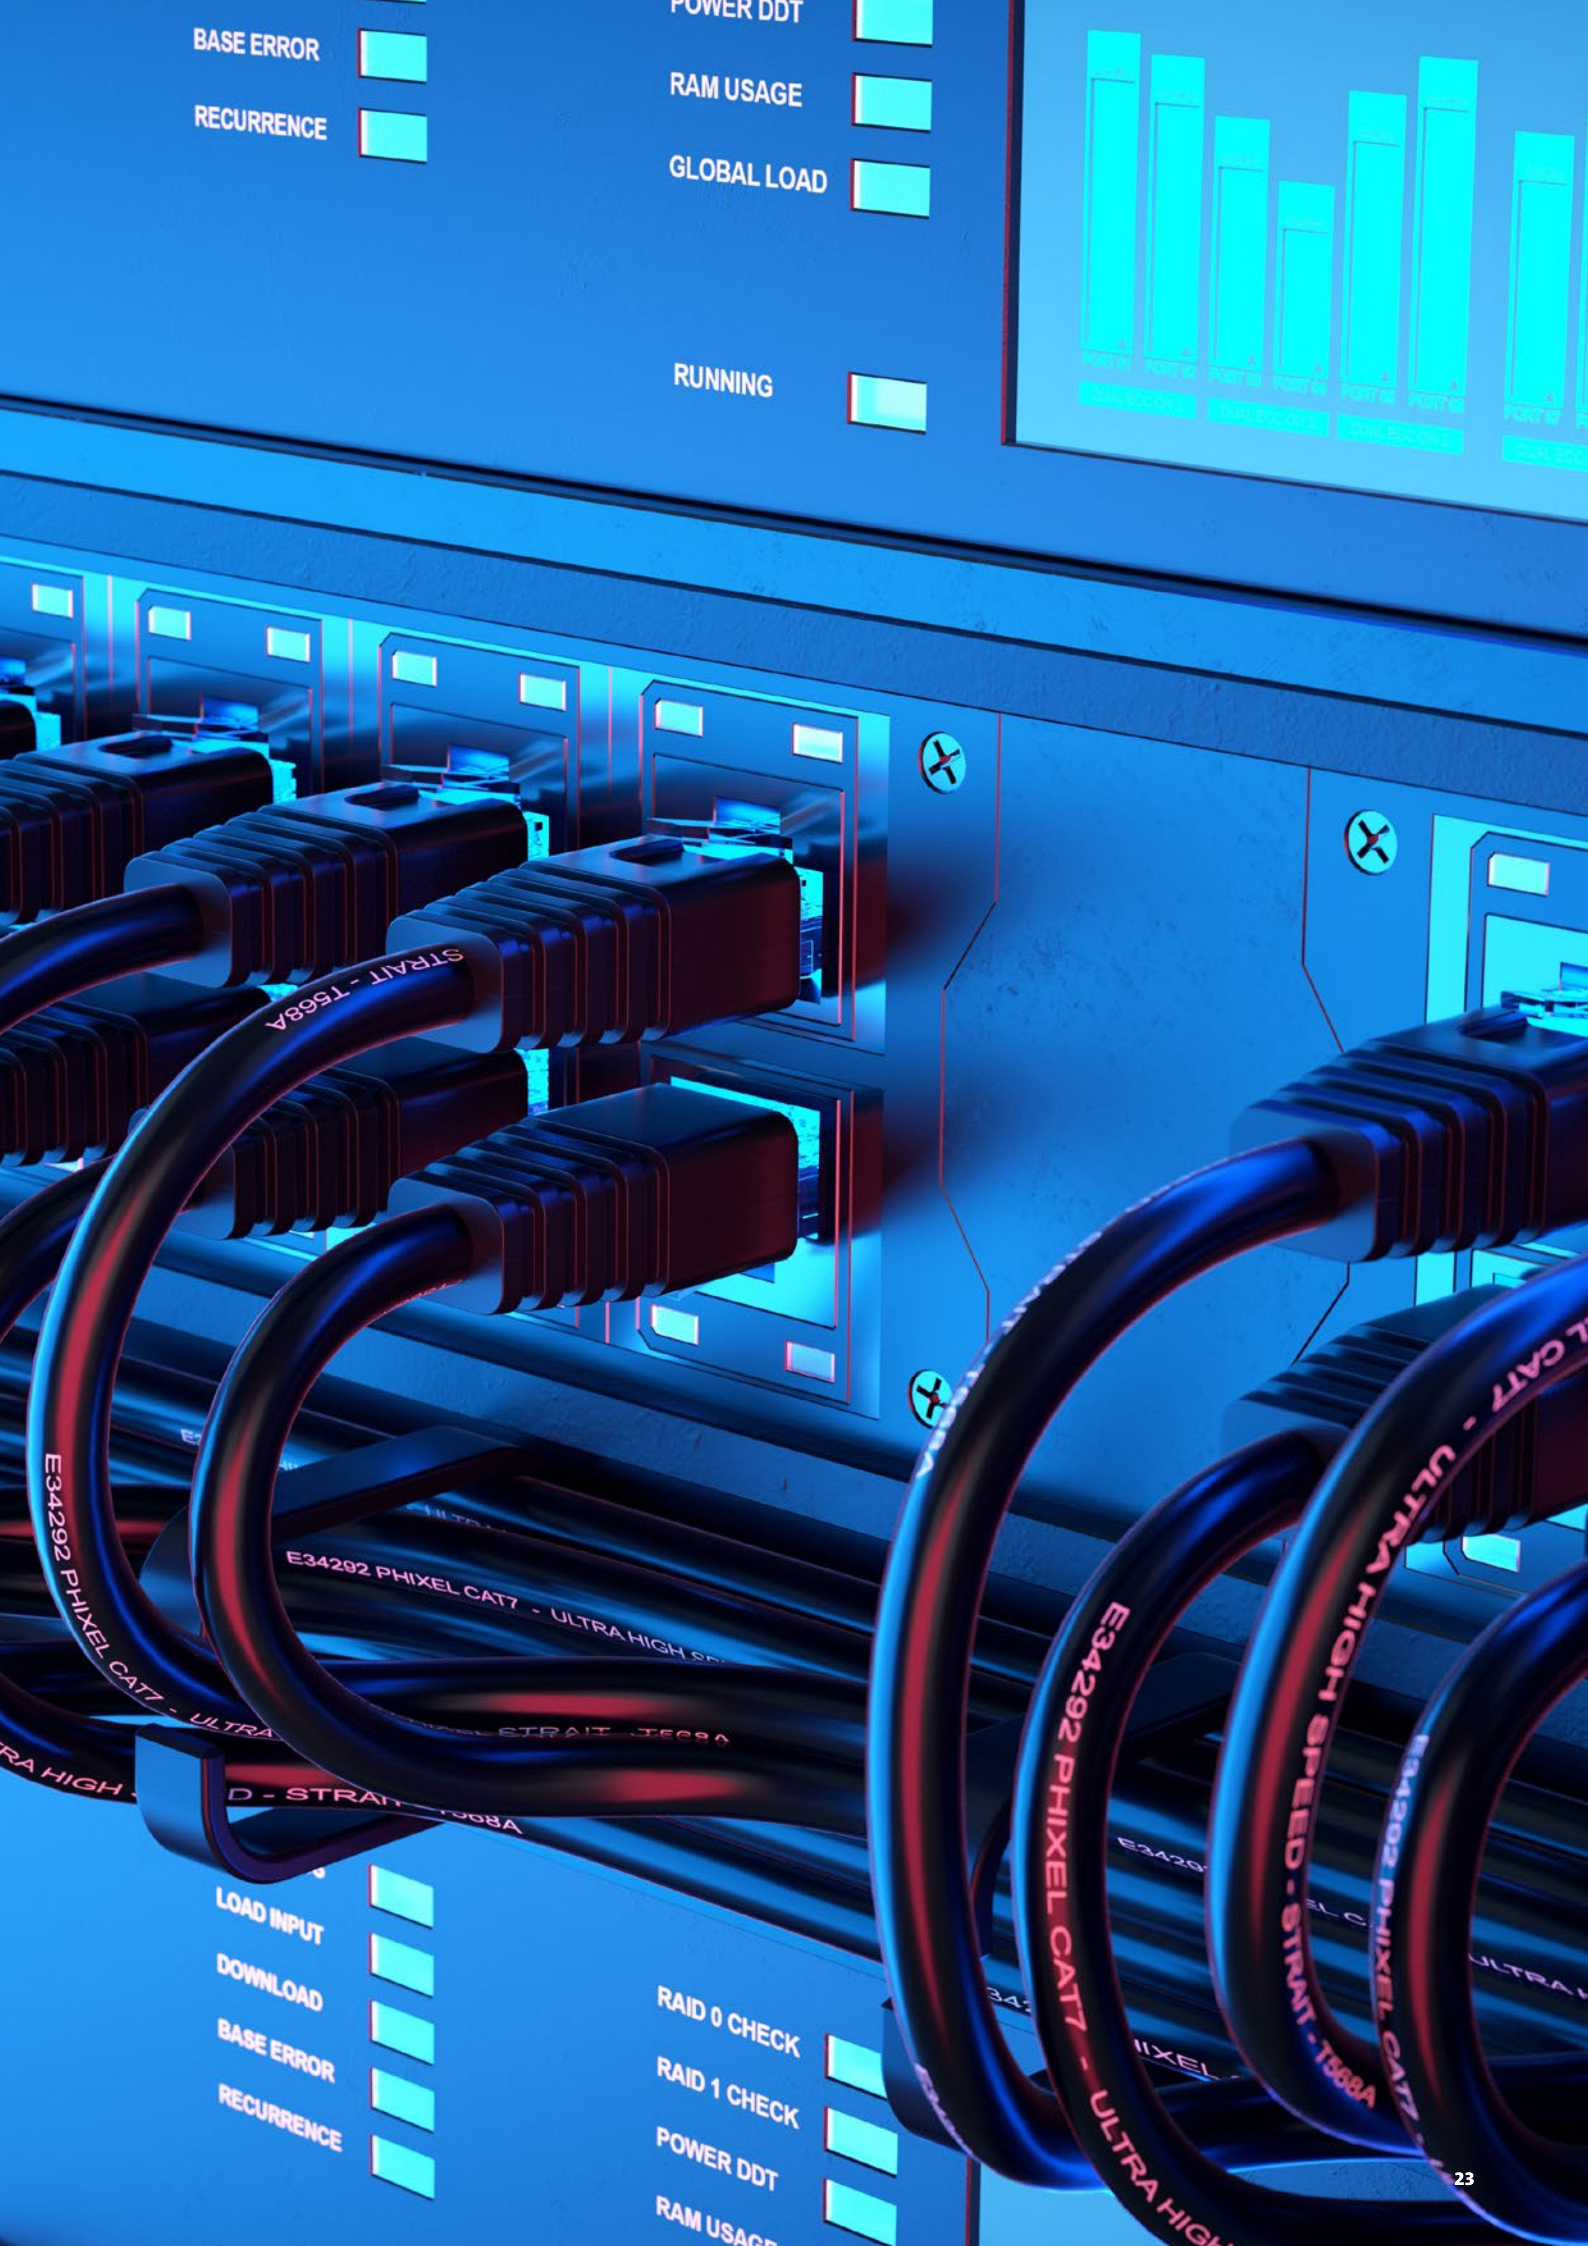
BASE ERROR
RECURRENCE
RAM USAGE
GLOBAL LOAD
RUNNING
E34292 PHIXEL CAT7 - ULTRA HIGH
LOAD INPUT
DOWNLOAD
BASE ERROR
RECURRENCE
RAID 0 CHECK
RAID 1 CHECK
POWER DDT
E34292 PHIXEL CAT7 - ULTRA HIGH
ULTRA HIGH SPEED - STRAIT - T568A

Convergence of chatbots and voice assistants

The advent of conversational agents (e.g., Siri, Google Assistant, Alexa, etc.) and custom chatbots has meant that natural language interfaces to data are on the rise. This requires the ability to successfully interpret, reason, and respond in natural language, using NLP.

Contributing to significant advancement and growth are the new generative language model technologies (OpenAI's GPT-4, Google's BERT, etc.), the availability of vast amounts of training data, and the exponential growth in computing power (including specialised processors called GPUs).

Digital tools using this technology for medicine are emerging – leveraging modern digital health standards such as FHIR to access data in electronic medical records (EMRs) and the SMART on FHIR application framework to embed practical applications in those EMRs.

An early example is Suki AI, a voice based medical assistant that understands the context of the doctor's practice and can learn their preferences. The system is so advanced that it determines intent (a command to review existing notes versus dictation of new ones) and accurately selects from similar terms (e.g., peroneal vs. perineal) to create clinically accurate notes. The technology leverages FHIR and SMART to integrate with the major EMRs, allowing doctors to use their iOS, Android, web, and Windows devices to access patient information, such as medications or vital signs and to create notes, including ICD-10 codes.

CSIRO's Ontoserver is now levering NLP and the GPT-3 generative language model to build OntoGPT – supporting wider adoption through its simpler, less technical user interface to standardised terminology that, like Suki, could become a part of more facile digital physician charting while increasing the medical content that is structured into ontologies like SNOMED CT. This will contribute to the reuse of medical data by SMART on FHIR based clinical tools for purposes such as decision support as well as for analysis and research.

As illustrated by Microsoft's new Bing, chatbots and voice assistants are converging. This means people will be able to control more sophisticated chatbots with speech enabled virtual assistants to help them with far more complex tasks than checking the weather and potentially even in their jobs.

These technologies represent an extraordinary epoch in medicine, where machines will be able to lighten the administrative load for clinicians, offer therapeutic support through chatbots that can take histories and offer education to patients, and enable improved clinical decision making. This together with their ability to enhance patient care when they are linked to electronic records via SMART on FHIR represents a potentially rich and impactful research and implementation opportunity for CSIRO.

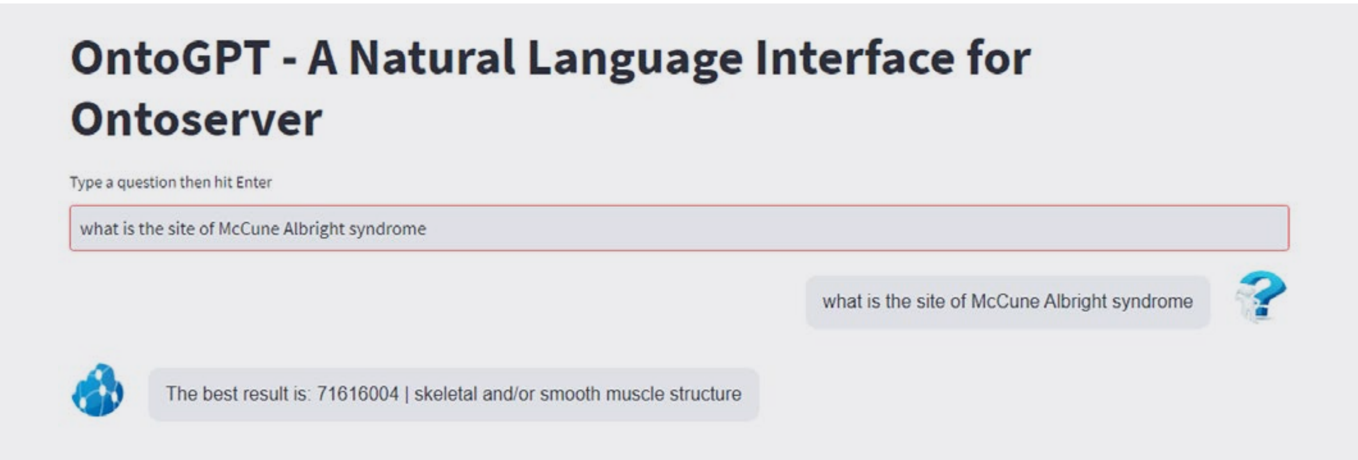

OntoGPT navigates SNOMED CT's complex structure to make it easy for its user to find the body structure affected by a rare disease.

Compute

The biggest contributor to the rise of machine learning in the last decade is the significant increase in computational power.

Ever since *graphics processing units* (GPUs), yes – the tools used for gaming – were leveraged for training deep artificial neural networks, an arms race has been in play. The paradigm in ML is that deeper and larger neural networks provide more performance, which in turn requires more and more powerful GPUs. This means that the development of machine learning technology goes hand in hand with the development of GPU technology and the number of GPUs available to the data scientist. For example, the recent rise of LLMs, such as ChatGPT, were made possible only with large clusters of GPU servers.

However, this kind of computational power requires significant investment, potentially resulting in large inequalities between research groups in terms of available compute.

The lab with the bigger computer can train larger deep neural networks on more data and thus better solve a problem than other labs.

While innovative machine learning techniques can subvert this narrative, research groups are at risk of falling short if they don't have the compute available to adequately develop machine learning solutions. When applying this to digital health, not having enough computational power could prevent the development of machine learning solutions targeted at improving patient care.

QUICK SCIENCE

Increasing interoperability

We have collaborated with government and industry stakeholders over several years to implement Fast Healthcare Interoperability Resources (FHIR) for primary care practice-to-practice exchange of health records.

As an extension of this collaboration, AEHRC was funded in the 2023-24 Federal health budget to support the Department of Health and Aged Care, Australian Digital Health Agency and HL7au to provide coordination and subject matter expertise to accelerate the development of national FHIR standards.

Regulation

AI and ML are increasingly used in medical software and devices. For clinical translation, several of our AI projects are considered Software as a Medical Device and will be subject to regulatory approvals (for example, TGA, FDA or CeMark approval). The pathway is well defined for AI where the model in each release is static, however there are both challenges and opportunities for deployed AI models which can learn from real world use and experience to improve their performance.

Medical imaging

Medical imaging (along with genetics) has become one of the key elements in precision medicine in advanced healthcare systems.

Since the first use of X-rays in 1896, the fields of radiology and nuclear medicine have helped revolutionise the diagnosis and treatment of a whole range of health conditions. There is currently a long list of available modalities including X-ray radiography, magnetic resonance imaging (MRI), ultrasound, endoscopy, elastography, tactile imaging, thermography, medical photography, and nuclear medicine functional imaging techniques such as positron emission tomography (PET) and single-photon emission computed tomography (SPECT).

In healthcare, outcomes are dependent on the acquisition technology, interpretation, and communication of the medical images. Traditionally, medical imaging is usually interpreted qualitatively by trained experts. One of the opportunities for advancement in this area of AI and health is the development of technologies that extract and analyse quantitative imaging biomarkers for use in screening, risk stratification, diagnosis, and treatment for various clinical and research applications. The developed technology turns images into information that is used for earlier detection of diseases and improved diagnostic accuracy.

A large component of this work involves the use of AI and in particular ML and deep learning approaches that can perform or improve image-based tasks such as image acquisition, reconstruction, quantification (segmentation) and analysis. This enables machines to interpret images using clinical scoring (of pathology), diagnosis and prognosis.

Not all imaging requires expensive medical imaging machines. Other applications of AI use images taken with regular cameras – for example when examining skin lesions or burns – or with cameras that photograph parts of the eye for retinal image analysis. Techniques include automated methods for registration of retinal images that are collected over time (longitudinally), or obtained using different retinal imaging devices/modalities, or are captured from different angles.

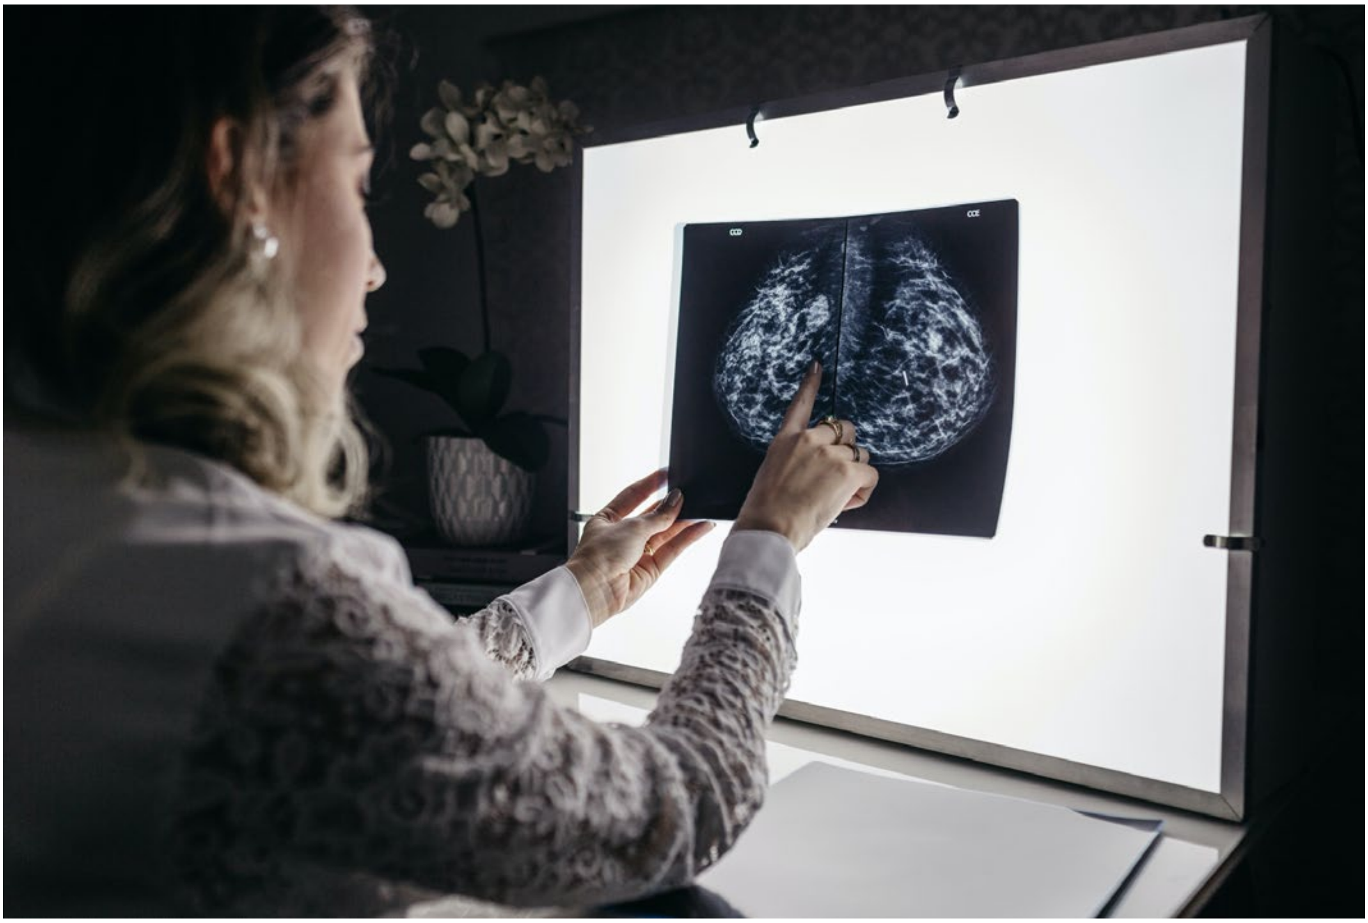

There are six main challenges/opportunities in medical imaging analysis.

AEHRC is currently undertaking research in these areas.

1 **Registration**
Mapping imaging data between time points or different patients.

2 **Synthesis**
Generating a new type of medical image from a input image, for example generating synthetic CT for radiation therapy planning.

3 **Reconstruction**
Combining data from different sensors to generate 2D or 3D images, for example generating 3D volumes from a set of ultrasound images taken from different positions.

4 **Multi-modal prediction**
Combining different type of data, such as imaging and text, to generate new insights.

5 **Segmentation**
Identifying the anatomical boundaries in images, for example quantifying different substructures in the heart.

6 **Classification**
For example, determine whether there is fracture present in a wrist x-ray? Or identify whether a patient is likely to respond to cancer treatment.

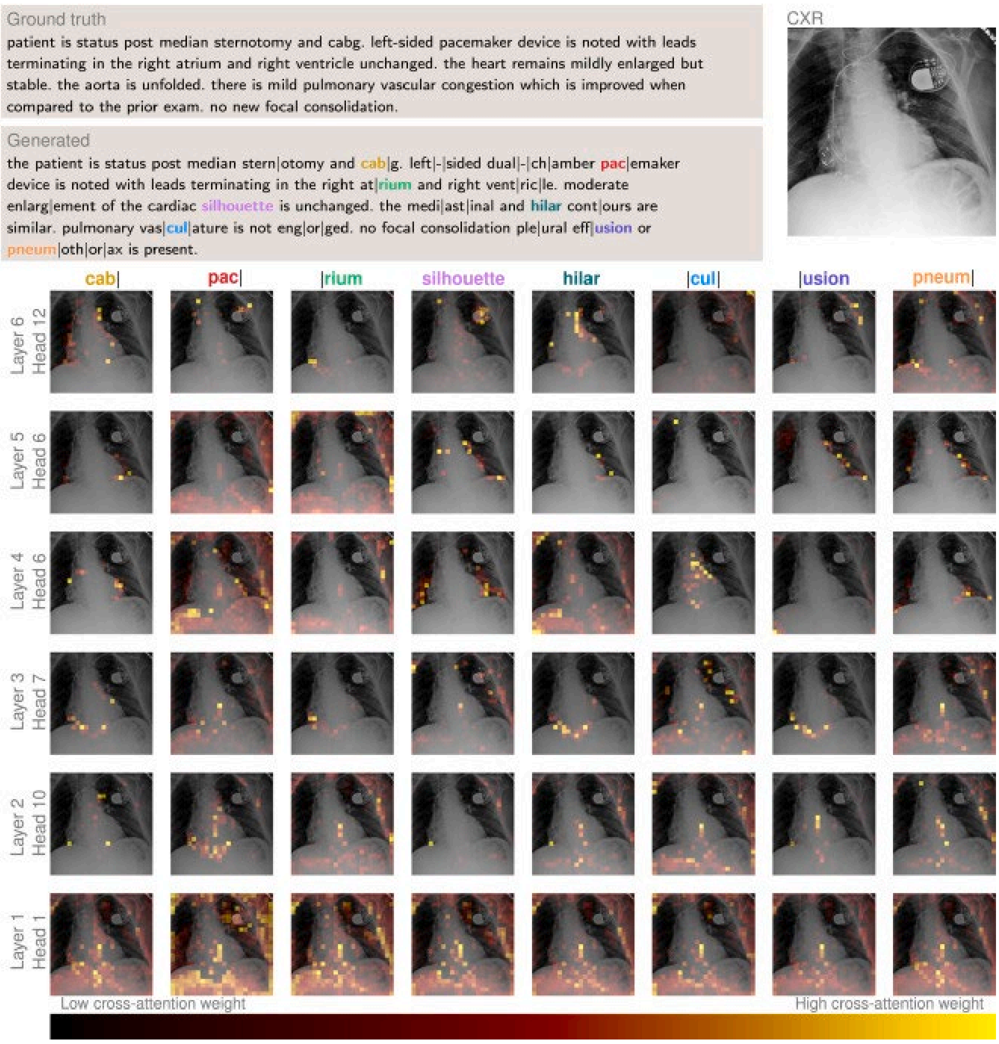

Case study comparing a 'ground truth' radiologist's report with a machine learning generated report of a chest X-ray. Taken from: https://doi.org/10.1016/j.artmed.2023.102633

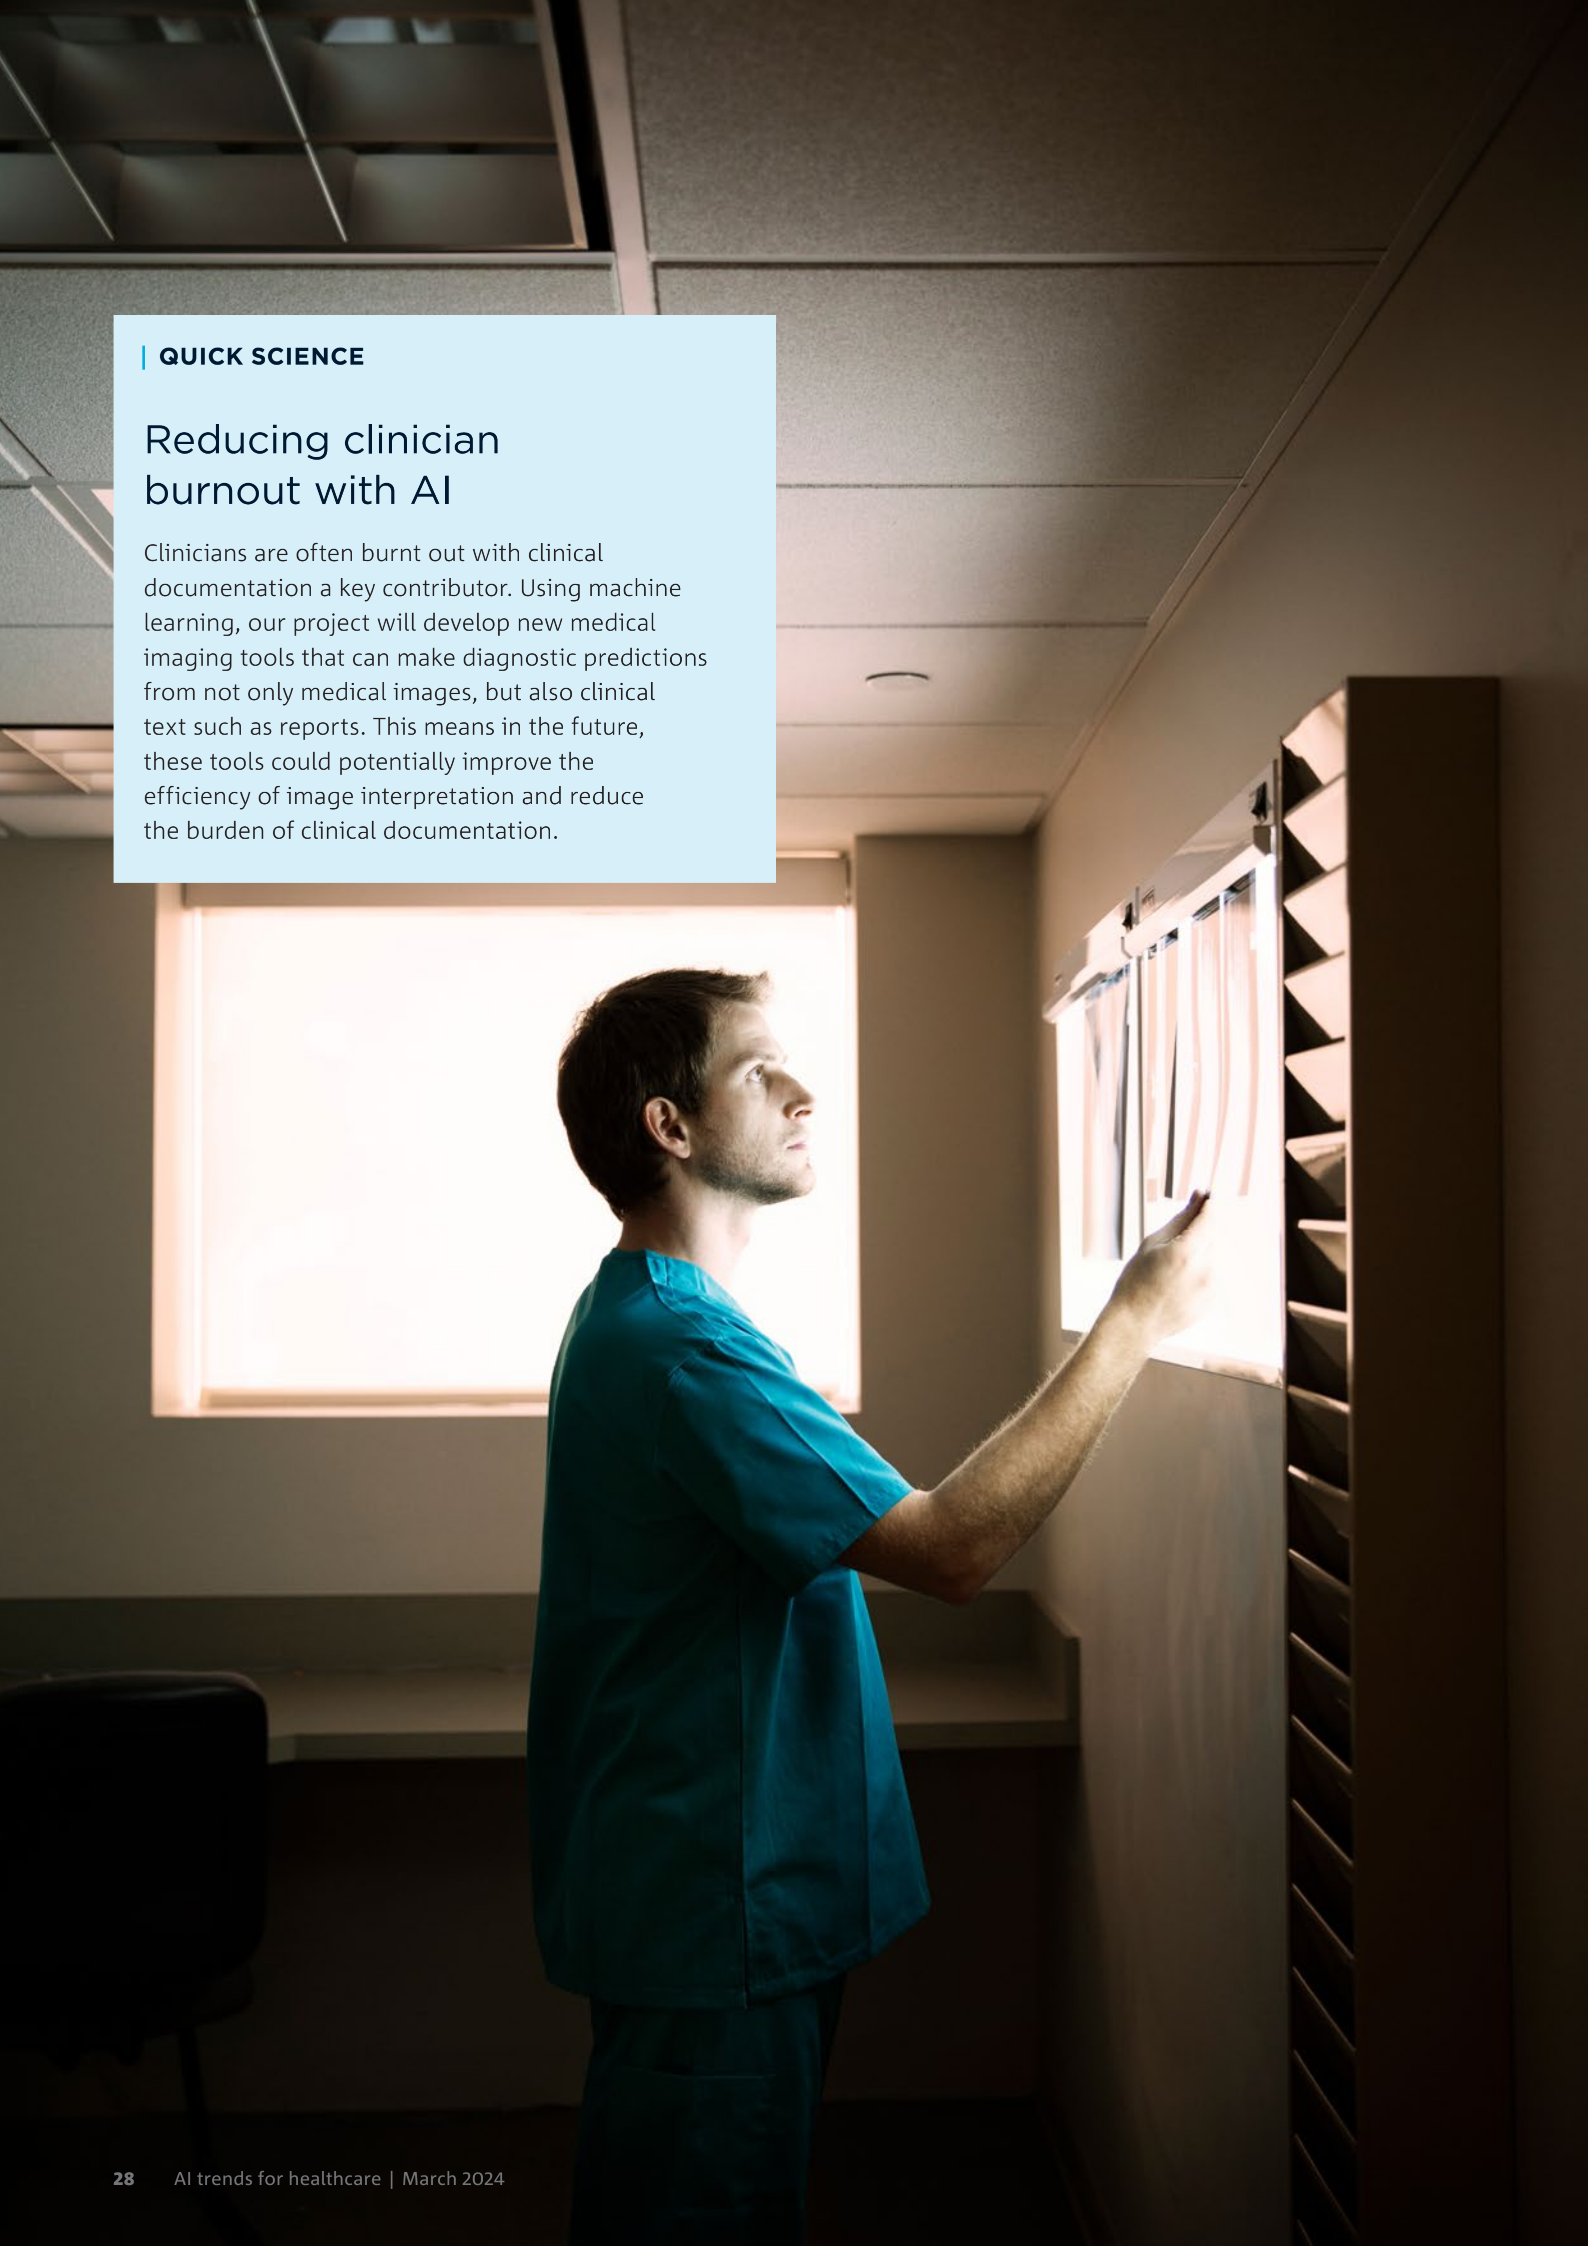

QUICK SCIENCE

Reducing clinician burnout with AI

Clinicians are often burnt out with clinical documentation a key contributor. Using machine learning, our project will develop new medical imaging tools that can make diagnostic predictions from not only medical images, but also clinical text such as reports. This means in the future, these tools could potentially improve the efficiency of image interpretation and reduce the burden of clinical documentation.

AI and workforce

There are multiple health workforce challenges locally and globally. These include growing demands due to more complex disease profiles combined with an ageing population and increased consumer expectations.

In addition, there are also various new workforce pressures that accompany changes in the labour market, including increased burden of administration, resulting in clinician burnout and challenges to recruit and retain staff.

The use of AI can ameliorate/buffer these challenges through supporting the health ecosystem at each layer:

Systems level

AI enables the use of data for reporting and identification of system efficiency and improvements, thereby reducing administrative burden on clinician, freeing up clinician's time for more direct patient care or leadership/strategic work.

Clinician level

AI enables appropriate data to be provided to clinicians at point of care, reducing clinician's time on time-consuming work that are comparatively 'low value' (or less complex/ require less cognitively demanding parts), such as liaising with other clinicians on complex care, basic triage, streamlining diagnosis and treatment process.

Patient level

Reduces demand on the system and clinician, patients are afforded higher quality treatment, personalised medicine reduce the diagnostic odyssey and improved patient experience and outcomes.

Research level

Discovery of new treatment and interventions required – reduce system demand and improved patient experience and outcomes.

Developing the AI workforce

AI is already in use by the healthcare workforce. The use of SNOMED CT is an excellent example of symbolic AI in healthcare. The use of SNOMED CT enables FHIR and the use of embedded SMART apps in electronic medical records, which can help reduce the burden of administration for clinicians and other health staff.

There are still more actions needed to fully digitise Australia's healthcare workforce.

To harness the capabilities already at our disposal within the health ecosystem appropriately and responsibly – it is critical that our workforce (from clinicians, scientists, system administrators etc) are equipped to identify both the opportunities, challenges, and potential pitfalls as we navigate these new paths together.

To this end AEHRC is involved in several initiatives aimed at educating the workforce to prepare them for a digitised work experience. Our engagement through FHIR connectathons educated the Australian healthcare industry and community to active, enthusiastic, and confident adoption of standards. To date, we have held over 30 events and community forums, each involving over 120 clinicians, software engineers and policy makers. We are also working with the university sector to provide FHIR training courses, as well as developing opportunities with health bodies such as Queensland Health to develop staff on FHIR aligned activities. Examples such as these help to build a robust workforce ready for the digitisation of health.

Aged care

There is an increasing interest in how AI might impact care provision and enhancement of day-to-day function for those ageing at home or residential care. With an ageing population and a dwindling aged care workforce, the evolution of AI is bringing hope and curiosity.

Aged care, both residential care and community-based care, is actively looking for solutions to assist in providing enhanced care while also implementing the Australian Government's Aged Care Digital Transformation Strategy. This strategy is predicted to set a standard by which aged care providers can be guided, to some degree, in the acquiring and using technology to support the delivery of care. This should act to enhance opportunities in using AI in all areas of the aged care industry.

> Opportunities for AI hold promise for more freedom for older Australians, as well as better support and clinical vigilance. Some key technologies for this include the development of new innovative assistive technologies to support day-to-day functioning of people living at home and in residential care facilities, wearable medical devices, algorithms for clinical decision support, and preventative risk management.

Opportunities for AI hold promise for more freedom for older Australians, as well as better support and clinical vigilance. Some key technologies for this include the development of new innovative assistive technologies to support day-to-day functioning of people living at home and in residential care facilities, wearable medical devices, algorithms for clinical decision support, and preventative risk management. These technologies may allow for consensual tracking of activities, collection of clinical data, support in clinical decision making and the enabling of agentic, preventative health behaviours based on feedback from technological devices. However, opportunities, especially given considerations of the age cohort and the resource demands of the industry, bring risk and challenge, and these must be considered in the ongoing introduction of AI into aged care.

While many technologies and IoT devices are already available for use, key challenges for the introduction of AI into the aged care area include the lack of rigorous evidence of the impact of AI-enabled products and services including physical, emotional, and ethical considerations; education of users to know what they are engaging with and how to use it appropriately and safely. Further research is needed, and should be actively sought out and supported, to ensure this type of technology is thoroughly understood, challenged, is reliable and effective. AEHRC is working to engage in multi-disciplinary research so that a broad lens is used in testing and validating AI in aged care.

Current AEHRC AI/aged care research

- Smart homes for independent living mobile health applications and clinician platforms to support chronic conditions (e.g., secondary prevention of stroke)
- Eye health diagnosis and prevention
- Falls prevention
- Responsible use of AI in aged care

Consent

As our society becomes more digitally literate and we start to appreciate the power of data, our demands for data privacy and security will increase. Specifically, patients and research study participants will want to have full ownership and control of their information. One aspect of this is dynamic consent, where access to personal information can be given and revoked depending on the context the information is used in.

Historically, individuals provided written consent, for example a blanket consent to all future medical research on their data. This was needed because returned to the individual to ask for informed consent for every new research use of their information was technologically and logistically impossible. This has stifled participation and limited the usefulness of the collected data as certain entities (commercial) were automatically precluded under this provided consent.

Advancement in digital consent and the communication through personal devices (phones, tablets) now enables efficient two-way interactions with participants and holds the promise for true dynamic consent.

Besides enabling individuals to control their data, another aspect of future-proofing consent is for individuals to have full ownership of their data. This can be achieved through several approaches, such as distributed storage and homeomorphic encryption of data, self-sovereign identity for management of credentials, and tamper-proof decentralised dynamic consent objects.

> '...the key for all future-ready consent platforms is to put the human at the centre of the design.'

However, the key for all future-ready consent platforms is to put the human at the centre of the design. Specifically, the sociocultural angle of molecular and medical data management must be considered as it encompasses the social license for research and medical applications. This in turn impacts participation, especially of Indigenous populations and, with that, our ability to provide the best care for the diverse populations of our countries.

Disability (NDIS Assessment Framework)

Provided it is developed in true consultation with the disability community, digital technology has the potential to transform the lives of people with disability. This is particularly the case when AI is embedded within assistive technologies. In such cases, AI can promote better functioning and greater independence and dignity.

One of the challenges with AI-enabled assistive technology is ensuring that products and services are fit for purpose for the people who use them. Not all technology is well suited to everyone, and there is often limited guidance on the identification and use of assistive technology for people with disability.

Working with people with disability, carers, and industry to guide market development and support improved matching of technologies to individual needs and preferences is one way to help surmount the challenge of harnessing the full potential of AI-enabled assistive technology.

There are opportunities for AI to assist people with disability in their daily lives. We've already begun to see the benefits of mobile technology to support people who are blind or partially sighted, and a range of digital products to support people living with epilepsy, mobility, and audio-visual impairments. But the expansion of the assistive technology space must be informed by people with lived experience of disability to ensure that products are safe, appropriate, and beneficial for individuals.

An example of the collaborative work needed in this space is our project with the National Disability Insurance Agency (NDIA), in which we developed a principles-based evaluation framework for AI-enabled assistive technology.

Our approach

The Framework draws on current frameworks, guidelines, and academic research as well as multiple rounds of stakeholder consultation with people with disability, their carers, industry representatives, peak bodies, researchers, service providers and government departments. This rounded and inclusive collaboration resulted in a person-centric approach for assessing AI-enabled assistive technology that accounts for an individual's needs and preferences, as well as the context (environmental, social, and cultural factors) in which the AI-enabled assistive technology will be used.

It also acknowledges the unique capabilities, preferences, and goals of end-users. The principles-based framework is guided by a set of six core domains. Each domain encompasses a principle, and two or more critical measurement areas.

QUICK SCIENCE

sBeacon

Increasingly an individual's genetics can be used to aid diagnosis or treatment decisions. This is referred to as precision medicine – the diagnosis and treatment of disease based on an individual's genes. And it's also true on a public health level. Two distinct challenges with using informative genetic information at the public health level involve the lack of diversity in the genetic databases, as well as challenges in accessing and analysing the data.

sBeacon offers a solution to these difficulties by creating a federated learning solution to allow exchange of specific information without sharing an entire database. The tool is also relatively cost-effective, allowing less wealthy nations the ability to contribute population data to data banks. On the other end of the spectrum, sBeacon also supports mega-biobanks by scaling to 3 billion genomic locations and 40 million individuals and is the only production-ready federated data exchange implementation of the GA4GH Beacon v.2 specification.

5

Case studies

VariantSpark for identification of SARS-CoV-2 variants

As COVID-19 spreads, the virus that causes COVID-19, SARS-CoV-2, accumulates mutations or changes in its genome, resulting in variants that result in more severe disease.

The current method of searching for dangerous SARS-CoV-2 variants is to look at the spike protein for single-mutation changes that might indicate more harmful effects in the human host. This method is limited in its capacity to detect groups of mutations, or mutation 'signatures' that can modulate disease risk.

By using a variant analytics pipeline, researchers can analyse the genome of the whole variant, identifying mutations and signatures that affect disease prognosis, allowing for a more wholistic understanding of their potential functional impact. The dataset used for this study consisted of too many datapoints to be analysed by traditional methods used for association studies. So, machine learning represented a solution for the analysis of large amounts of data to make these vital assessments of harmful SARS-CoV-2 mutations.

We used machine learning to develop a cost-effective and accurate machine learning solution to identify severe COVID-19 causing mutations in SARS-CoV-2 viral genomes.

We developed a faster and more comprehensive way to identify emerging and dangerous COVID 19 variants, by analysing the genome of the whole variant, rather than the current method of monitoring changes to the spike protein.

We used the analytic capability of a powerful machine learning tool we developed called VariantSpark. Using this tool, we were able to analyse the genomes of 10,520 SARS-CoV-2 samples, which is the largest number of samples ever analysed in this way.

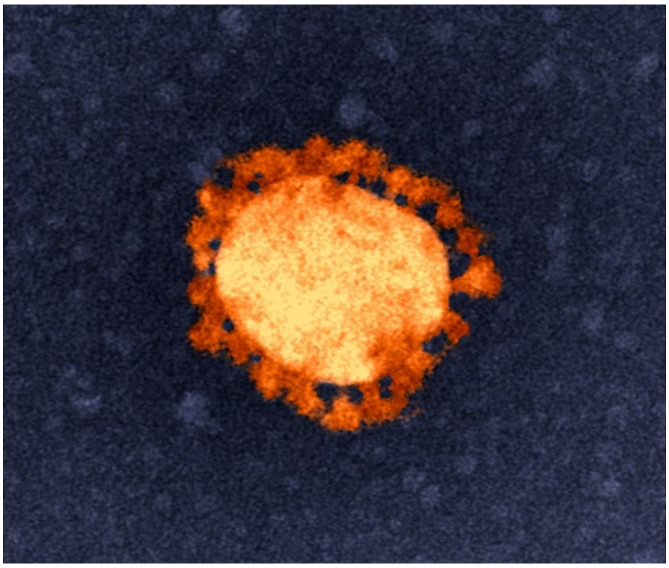

This method could help inform an early warning system that can determine which variants will be the deadliest to humans. Our approach was able to identify variants that could be monitored a week before they were flagged by health organisations.

> 'We can also apply this approach to other viruses – in fact it has the potential to become the international standard of disease surveillance,' said Transformational Bioinformatics group lead and Research Scientist, Dr Denis Bauer.

CSIRO worked with RONIN, whose cloud-based system supported the analysis, and Intel on the study, which is the largest of its kind in the world.

HOTspots surveillance and response program

Anti-microbial resistance (AMR) is increasing around the world and becoming a major threat to human and animal health in Australia and elsewhere. Monitoring outbreaks of AMR bacteria and infections is an increasingly important line of defence against possible bacteria-borne epidemics. Outbreaks in Australia are likely to come from the tropical northern region and can travel quickly through human and animal vectors if not contained.

CSIRO Research Scientist and Research Team Lead for the Digital Solutions for AMR group Teresa Wozniak, says ‘We know AMR is a global problem. In Australia, the problem is hidden because national surveillance activities don’t capture the most vulnerable Australian populations.’

The HOTspots program delivers critical data on antimicrobial resistance (AMR) to clinicians and policy makers to reduce the threat in Northern Australia. A digital platform offers secure and interactive spatiotemporal AMR analytics and data visualisation to practitioners and health decision-makers.

Currently over 200 sites contribute data to the HOTspots program. Data comes directly from pathology providers, hospitals, community clinics and GP practices. Clinicians can use the data in clinical decision-making about antibiotic prescriptions, including updating local antibiotic guidelines, stewardship education and program development.

The HOTspots program contributes data to national surveillance activities via the Antimicrobial Usage and Resistance in Australia program to fill the gap in the surveillance of regional and rural settings in Northern Australia.

HOTspots is currently focussed on expanding geographically to other states and territories and to include data from animal and environmental sector.

HOTspots team members Teresa Wozniak, Lorraine Bell and Majella Murphy on a research trip in Darwin for HOTspots.

An artificial intelligent system for prostate cancer diagnosis in whole slide images

1 in 5 men in Australia experience prostate cancer.

Radiation therapy is one of the main methods for treatment for prostate cancer.

Anyone who's experienced radiation treatment for cancer will understand the level of care and attention given to the measurements of the site receiving radiation. The reason for this vigilance is because clinicians need to ensure they focus the radiation onto the cancerous tumour while simultaneously avoiding radiation toxicity to surrounding areas.

Clinical research trials make a vital contribution to the development of standards and protocols that guide cancer treatment. These best practice guidelines are especially important when cancerous tissue is small and surrounded by other organs, as is the case with the human prostate. The method to identify the boundaries of these volumes ('volume delineation') needs to be accurate and precise as it is used to guide the deposition of radiation in patients receiving radiation treatment.

Imaging via CT and MRI is one of the main sources of diagnosis and assessment for treatment in both clinical and research settings and is currently being used in image guided radiotherapy treatment trials. During these trials, which often take place across various hospitals and treatment centres, radiation oncologists plan treatment (including volume delineations) for patients according to a proposed clinical study protocol. They then send these plans to the lead study investigators for review. If the plans do not meet the trial protocol they are rejected (for example, if organ boundaries have not been correctly identified). Once the plans meet the trial protocol the patient is treated, and the treatment outcomes are recorded and eventually published.

While there are delineation guidelines for radiation therapy clinical trials, these can sometimes differ from local guidelines. Protocol adherence can also be different across sites. This can lead to the potential for uncertainty which may degrade the accuracy and quality of reporting from the trials. One of the challenges faced in these clinical trials is that the volume delineation quality assurance is usually performed manually by radiation oncologists who are also the trial principal investigators. When done manually, this is time consuming and labour intensive and dependent on human factors such as attention and vigilance for its accuracy. In practice this means that only a small portion of the plans for trial participants are reviewed.

Machine learning methods are well suited to provide automated quality assurance of volume delineation in radiotherapy. Recent years have seen a few attempts to develop automatic radiotherapy delineation quality assurance tools on CT, many of which operate on 2D slices. Compared with CT, MRI has higher soft tissue contrast without using ionising radiation so is increasingly used with pseudo-CT in prostate cancer radiotherapy planning.

> 'Machine learning has the potential to improve the efficiency of volume delineation quality assurance and reduce the uncertainty in radiotherapy planning.'
> Hang (Hollie) Min, CSIRO Research Scientist

Currently, there is still a lack of research on automatic delineation quality assurance for MRI-based radiotherapy.

In this world-first study, we investigated the efficacy of automatic radiotherapy delineation quality assurance on prostate MRI in a multicentre clinical trial. We developed and validated an automated machine learning system which flags clinical target volume and organs of risk delineations that may not meet trial protocol and therefore require further review and revision by radiation oncologists.

CSIRO Research Scientist, Hang (Hollie) Min, said machine learning has the potential to improve the efficiency of volume delineation quality assurance and reduce the uncertainty in radiotherapy planning.

To do this, we trained a machine learning model to assess the volume delineation and determine whether it is acceptable or a violation based on its similarity to the machine generated benchmark delineation. The model was then tested on data from a multicentre trial and generated a pdf report for clinicians to review.

Our outcomes showed the AI model can identify the delineations that did not meet trial protocol and require further revision. This can also be extended to radiation treatment for other cancers, and on both MRI and CT data.

CSIRO Principal Research Scientist and author on the published work, Jason Dowling, says that automated quality control is an increasingly important research field which may lead to improved patient treatment and outcomes. The method can be used for training, validating contours in real time, or for auditing data from retrospective radiation oncology databases.

Hollie Min's research is reducing uncertainty in radiology planning.

Data harmonisation for Alzheimer's research

Large observational studies of patients with Alzheimer's disease (AD) across the globe often involve different cognitive tests and scales. While the studies are all in the same area, the data from the studies are not consistent, making analyses and conclusions drawn from the data difficult. This difficulty often means insights about AD are not gleaned from existing data and, importantly, opportunities to know more from the massive amount of data that exists are missed.

To combat this challenge, the method used is to harmonise the data, which essentially means to bring all the various data into alignment so it can be analysed. We did this using an AI-based method for harmonising cognitive data across large observational cohorts in AD, which have different cognitive tests and scales.

> CSIRO Research Scientist, Rosita Shishegar explains the importance of this data set to the discovery of prevention and treatment of Alzheimer's, 'If we can use digital technologies to find ways to prevent or even delay the onset of disease and give people more time with their families, that's huge.'

We used a multiple imputation approach to create a common cognitive measure based on the available cognitive data from the Australian Imaging, Biomarkers, and Lifestyle (AIBL) and Alzheimer's Disease Neuroimaging Initiative (ADNI) cohorts.

The imputation model used age, education, sex, and the apolipoprotein E (APOE) ε4 genotype as predictors of missing data.

The harmonised cognitive measure was evaluated for its ability to predict clinical outcomes, including cognitive decline and conversion to mild cognitive impairment (MCI) or AD.

The results showed that the harmonised cognitive measure performed well in predicting clinical outcomes, and it was more predictive than using the original cognitive measures from each cohort separately.

Since the initial results were published we have extended the method to other cohorts with different cognitive tests or scales, in an international, multi-centre research project to track cognitive decline over time, and to identify individuals at risk of developing MCI or AD.

Our proposed method allowed us to create the largest dataset of Alzheimer's disease in the world as part of this international research.

The method can also be used to harmonise other types of data across cohorts, such as imaging or biomarker data.

Pathling for scalable query of FHIR and clinical terminology data

One of the major challenges in the digital transformation of healthcare is the management of data so it can be sent, queried, analysed, and the analysis easily understood.

Pathling is a CSIRO developed tool that enables query and transformation of bulk and streaming sources of FHIR data. Pathling also works with a FHIR terminology server such as Ontoserver to facilitate query of clinical terminologies such as SNOMED CT across large datasets. Not only are these functions central to the ability to use data, they also represent a foundational capability that enables construction of the standards-based data pipelines to train the next generation of AI models within the healthcare sector. In other words, Pathling is an essential tool for future proofing Australia's healthcare system.

Pathling uses

- Data preparation within Python and R data science workflows
- Creation of views for self-service analytics and business intelligence
- Providing a FHIR API for patient cohort identification and data extraction services
- Streaming analytics from real-time data sources such as medical devices
- Preparation of training data for machine learning and artificial intelligence

Recently the Federal Government commissioned Australian Genomics to develop recommendations around the implementation of a national approach to genomic information management. Pathling was selected to participate in a prototyping exercise to identify solutions to deliver this national infrastructure.

Pathling was integrated with the software that facilitates access to genomic research data to provide a connected solution. This is a federated system that allows access to genomic and phenotypic study information to researchers who have been authorised, while allowing the data to remain safely within the control of the custodians responsible for it.

This solution received excellent feedback from the international panel that reviewed it, which positions it well for inclusion in the next phase of development within the next funding cycle.

Pathling has also been a key driver behind the international 'SQL on FHIR' initiative, which is an international standards collaboration aimed at making FHIR easier to use for analytic use cases. This new Health Level Seven standard builds on the work of our researchers to deliver standards that will enable 'push button population health', standards-based quality indicators, and further remove barriers to information blocking.

John Grimes
CSIRO Principal Research Consultant and creator of Pathling. John is a leading expert in FHIR, clinical terminology and health analytics.

Measuring amyloid levels in Alzheimer's disease

Dementia is the second leading cause of death of Australians and is likely in future to become the leading cause as our society ages. Despite this, there is no existing cure. Alzheimer's disease is the most common form of dementia, estimated to make up 60-70% of cases (WHO, 2023, who.int/news-room/fact-sheets/detail/dementia).

One suggested route for cure possibly involves the eradication of amyloid and tau plaques from the brain, as the presence of these biomarkers are among the hallmarks for the diagnosis of Alzheimer's.

Before we can consider how to develop medicines and other tools to cure Alzheimer's, work is needed to find evidence of how the brain changes throughout the progression of the disease, which can often take decades.

AI, with its capacity to analyse large amounts of data, is adept at gathering evidence about the progression of Alzheimer's via imaging studies.

> 'Before we can test the efficacy of amyloid related drug treatments, we need to be able to accurately measure amyloid in the brain. The impact of this research is vital to one of the potential treatments currently being developed in the Alzheimer's space, called anti-Aβ therapy.' Pierrick Bourgeat

Recent developments in medical imaging have allowed the in-vivo examination of brain pathology associated with Alzheimer's disease, such as Aβ plaques, glucose metabolism, cortical atrophy and more recently, tau tangles. In vivo is a term to describe something inside an alive body as opposed to outside of a dead body. It's important to look at the brain when it is still alive because once a person dies the brain immediately begins to shrink, which can sometimes affect the accuracy of methods that involve measurement. Two important tools for in vivo imaging are MRI and PET images.

MRI is a useful tool for measuring neurodegeneration but as it focusses on cortical atrophy, it lacks the capacity to differentiate between a range of neurogenerative diseases that along with Alzheimer's disease cause shrinkage of the cortex. PET, on the other hand, measures amyloid and tau proteins, which are specific to Alzheimer's, making it an excellent tool for diagnosis and monitoring of the disease.

CSIRO Research Scientist and lead author on the study Pierrick Bourgeat said, 'Before we can test the efficacy of amyloid related drug treatments, we need to be able to accurately measure amyloid in the brain. The impact of this research is vital to one of the potential treatments currently being developed in the Alzheimer's space, called anti-Aβ therapy.'

A trial version of CapAIBL is available on MilxCloud, our web platform (milxcloud.csiro.au).

Pierrick Bourgeat
CSIRO Research Scientist

AI-driven technology to develop chatbots for therapy

AEHRC are leaders in the development of therapeutic chatbot technology.

Chatbots are increasingly used therapeutically to both capture and deliver information.

With a health system under pressure and clinicians struggling to keep up with administrative burdens, chatbots could help support clinicians ensure critical aspects of the therapeutic process, such as patient education, remote monitoring, and regular check-ins.

Through AI technology scientists at CSIRO's Australian e-Health Research Centre (AEHRC), are investigating whether clinician burden can be reduced without a deficit in the quality of patient care.

So far, results are promising.

The AI chatbot framework is currently used in three chatbots at CSIRO – Dolores (pain chatbot), Harlie (a chatbot developed for people living Parkinson's disease and other language disorders) and Quin (a smoking cessation chatbot) – was developed by CSIRO and partners.

The chatbots are currently being tested for therapeutic efficacy as well as design and user-experience.

Dolores is a chatbot who works in chronic pain.

The CSIRO research scientist who developed Dolores' framework, David Ireland, said pain management support is a vital component of the healthcare system but one that is often quite resource heavy due to the prevalence of chronic pain in the community.

He said, 'Pain-related diseases, low back pain and migraine, are the leading causes of disability and disease burden globally. Chronic pain is a disease state in its own right but studies have shown pain education is a useful tool in helping people manage their chronic pain which is often incurable.'

Dolores conducts a pain history interview, then provides an education session for people with chronic pain.

In a recent collaborative study Dolores was shown to be received positively in cohorts of patients across three different age groups.

Scientists looked at how receptive different age groups (adolescents, young adults, and adults) are to the use of Dolores. The study suggested high acceptability in quality, accuracy, usability in all groups.

Quin is a smoking-cessation chatbot, built by researchers at CSIRO, University of Queensland, and Prince Charles Hospital, that gives non-judgemental advice to quitters and sets up regular check-in appointments to improve consistency of therapeutic support.

Quin was developed using a combination of thematic analysis of Quitline transcripts, clinician consultation and research into user experiences of quit smoking apps.

Quin builds a profile of the user and then provides advice about specific issues & challenges the user is facing in their quit smoking journey. After a consultation, Quin builds a TODO list which might include tasks such as 'speak to your doctor about quitting'.

Although Quin provides AI support for a person looking to quit smoking, it is designed to encourage human interactions. For example, it will encourage users to phone the Quitline.

Quin is available 24/7 on the mobile phone to provide advice or encouragement. Early focus groups revealed some participants conveyed a sense of accountability to Quin.

Future research is focussed on a clinical trial of Quin's efficacy.

Harlie is an artificial conversational agent and is designed to converse with users with neurological conditions that may impair speech, such as Parkinson's and dementia, or even autism. The more people interact with Harlie, the more conversation topics she develops. Harlie could in the future be used as a virtual companion for social interaction therapy, coaching and remote monitoring.

David Ireland
CSIRO Research Scientist

Definitions

Activities of daily living: measure used in healthcare to refer to people's daily self-care activities.

Artificial intelligence: computer systems developed to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.

Biotechnology: technology based on biology, usually intended to improve human health and society.

Cloud: cloud computing is the delivery of computing services – including servers, storage, databases, networking, software, analytics, and intelligence – over the Internet ('the cloud') to offer faster innovation, flexible resources, and economies of scale.

FHIR: Fast Healthcare Interoperability Resources is the rapidly growing global standard for representing and sharing health information. When combined with SMART it can support EHR/EMR connected health applications for a variety of clinical purposes for both practitioners and their patients.

HL7: membership-based not-for-profit public company that facilitates the adoption of e-health in Australia by promoting effective use of standards and products developed by HL7 International and supporting their enhancement to meet local needs.

Internet of things: tools that allow electronic devices to 'speak' with one another.

Machine learning: field of study that gives computers the ability to learn without being explicitly programmed. There are two main ML tasks: classification and regression. Classification involves using a ML model to 'classify' some data according to a finite set of categories; for example, classifying the type of cancer found in a pathology report: breast, lung, etc. The simplest case being a binary classification – yes/no, true/false, cancer/not cancer, etc. Regression, in contrast, uses a ML model to predict a continuous value rather than a category. For example, predicting length of stay for a patient given their condition.

Magnetic resonance imaging: describes an imaging technique that uses the magnetic field and radio waves to take pictures inside the body.

Natural language processing: using computational techniques to analyse and synthesise natural language and speech.

Ontoserver: next-gen FHIR terminology server developed by the Australian e-Health Research Centre.

SMART: technology layer that rests on FHIR to provide identification of app users, permissions to access clinical data and facile access to that data when apps launch.

SNOMED CT: clinical terminology owned, maintained, and distributed by SNOMED International.

Software as a medical device: Software, which on its own is a medical device.

Acronyms

ABI	Acquired Brain Injury
ADL	Activities of Daily Living
AEHRC	Australian e-Health Research Centre
AI	Artificial Intelligence
AMT	Australian Medicines Terminology
ASD	Autism Spectrum Disorder
AWS	Amazon Web Services
CapAIBL	Computational Analysis of PET for the Australian Imaging, Biomarker & Lifestyle Study of Ageing
COVID-19	Corona Virus Disease identified in 2019
CP	Cerebral Palsy
CPU	Central Processing Unit
CSIRO	Commonwealth Scientific and Industrial Research Organisation
CT	Computed Tomography
CT-MR	Computed Tomography merged with Magnetic Resonance Imaging
ED	Emergency Department
EHR/EMR	Electronic Health/Medical Record
FHIR	Fast Healthcare Interoperability Resources
HL7	Health Level 7 – The International Standards Body for Pathology
ICD	International Classification for Disease
ICU	Intensive Care Unit
ML	Machine Learning
MRI	Magnetic Resonance Imaging
NLP	Natural Language Processing
PET	Positron Emission Tomography
QIMR	Queensland Institute for Medical Research
SaMD	Software as a medical device
SMART FHIR	Connected App Platform
SNOMED CT	Systematised Nomenclature of Medicine – Clinical Terms
SPECT	Single-photon emission computed tomography

As Australia's national science agency, CSIRO is solving the greatest challenges through innovative science and technology.

CSIRO. Creating a better future for everyone.

Contact us
1300 363 400
csiro.au/contact
csiro.au

For further information
Australian e-Health Research Centre
Janet Fox
+61 7 3253 3646
janet.fox@csiro.au
csiro.au/healthandbiosecurity

B&M | 23-00456