JOURNAL & PROCEEDINGS

of the

ROYAL SOCIETY OF NEW SOUTH WALES

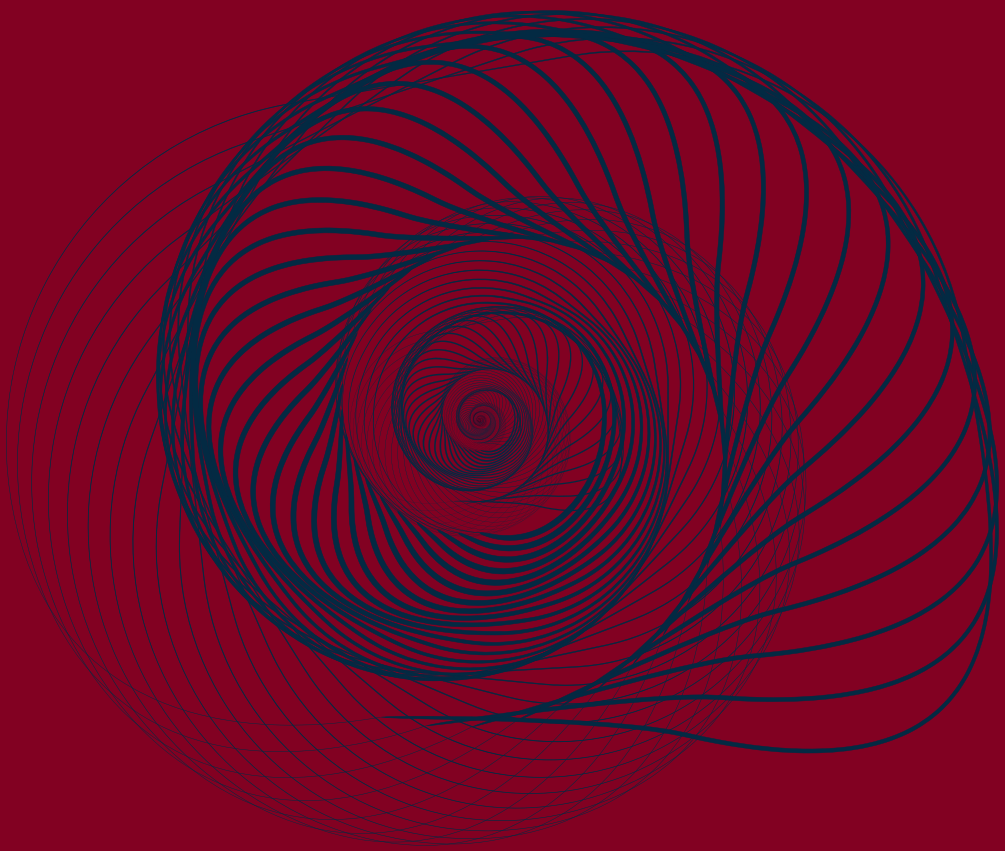

2026

Volume 159 Part 1

Numbers 501 & 502

ISSN (printed) 0035-9173
ISSN (online) 2653-1305

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, pp. 1–3. ISSN 0035-9173/26/0101-03

Editorial: Reporting the Forum

Robert Marks

Editor

Reporting the Forum

In May 2016 when I took over Editorship of the *Journal*, there had already been a Four Societies Forum at Government House, Sydney, on 14 September 2015, on the Future of Work. The eight papers presented at the Forum were published in the December 2015 issue of the *Journal*. Nine papers presented at the second Forum, on 29 November 2016, were published in the June 2017 issue of the *Journal*. Since then there have been 11 Forums, with presented papers published in the June issue of the *Journal* of the following year.

In 2018 one of the presenters, Sir Peter Glickman, did not supply us with a paper of his presentation, despite our requests. We made a transcript of his presentation, from the YouTube video, including images from his overheads. It was cleared with Sir Peter (then Chief Scientist with the NZ Government) and published in the June 2018 issue, with 9 other papers.

Unfortunately, since then fewer presenters have supplied us with papers of their presentations. At some inconvenience and cost, we transform the video presentations (on YouTube) into papers printed in the *Journal*.

But a departure in successive Forums has been a move away from individual presentations to group discussions of various topics. This happened especially in the 2024 Forum (published here in June 2025) and in the 2025 Forum (published below).

The problem for the Editor is that discussions are not papers: they are difficult to edit, and transcripts are unlikely to be read when the audio is available on YouTube.

I have reluctantly decided to recommend to my successor that discussions not be transcribed and published, thus avoiding the cost, time, and inconvenience of turning them into written versions for the *Journal*. (This is one reason for the delayed publication of this issue.) We can still publish individual papers (the Governor's introduction and any keynote addresses, etc), but leave the discussions to be viewed on YouTube.

Presentations and discussions in the 2025 Forum are included in this issue. The topic was "AI: the Hope and the Hype, "and the presenters were by and large optimistic about the consequence of the different forms of AI — generative, agentic, etc. — in the economy, despite a general belief that AI will adversely disrupt employment, as reflected in opinion polls. There is a growing literature about AI — its impacts, the prospect of more data centres in Australia, what the Pope has to say about AI, and so on, but this Editorial is not the place to pursue and ventilate this literature. There will be opportunity enough now and into the future.

A post-literate world?

My resolve to stop transcribing Forum discussions and instead to leave them to be viewed on the YouTube recordings (referred to above) was made for the convenience of the editorial staff. But a recent article in *The Atlantic* (Horowitch 2026) argues that an increasingly large proportion of the population would anyway prefer to watch

such events on the screen instead of reading transcripts. Indeed, Horowitch reports that Americans (and I would guess Australians) are reading much less than they used to, especially books, magazines, and newspapers. This is not just among the young: in America, this decline is across all age groups, gender, and education levels, even among college graduates. Moreover, the books that people do read are apparently written more simply than they used to be: shorter sentences, simpler syntax.

And yet people are apparently reading more words than previously: in emails, on social media, text messages, captions, etc. — but not longer written works that convey rich information, and might demand more thought and inferences from the reader. Horowitch dubs this a "post-literate" society.[1] (She characterises Donald Trump, from his writing and speaking habits, the first post-literate president. But not the last.) Joshua Meyrowitz (1985) predicted that the decline of print and rise of electronic media would ultimately push people toward populist leaders. These are worrying trends for democracies.

In the meantime, we push on publishing the *Journal & Proceedings* in a print version as well as on-line. One motivation is to capture and record in libraries across the globe activities of the Society and articles that might otherwise disappear.

Other papers in this issue

Other papers in this issue include Professor Emerita Maria Skyllas-Kazacos writing about the development of the vanadium batteries now being employed in solar electricity farms in China. The development of this type of battery has been her goal over many years at UNSW, and the designs are now being licensed for use in large-scale generation. A second paper is from a presentation (the 2011 Ben Haneman Memorial Lecture at the State Library of NSW) by the late Peter Tyler, erstwhile RSNSW historian, on European (i.e. non-British) scientists in colonial Australia. Whereas many of the British scientists were self-funded amateurs, the continentals were often government-funded. I am grateful to Peter Keeda for alerting me to this paper. Adrian Lee, at Westmead Hospital, received a 2025 RSNSW Early Career Early Career Research and Service Citation. This resulted in his writing a paper on the Sjögren's disease biobank and registry in NSW.

The issue contains three PhD abstracts, and, sadly, five obituaries (Heather Goodall, Stephen Hill, Emma Johnston, Sir Anthony Mason, and Keith Suter).

Housekeeping

After more than ten years, this issue is my second last as Editor. I am assured that efforts to find a replacement are well under way. As always, I wish to thank Jason Antony for his sterling efforts at type-setting the *Journal* and Lindsay Botten at maintaining the *Journal* web pages. The Editorial Board has been of great assistance, particularly Jessica Milner Davis and Brynn Hibbert. I would also like to thank Christie Slade and Roy MacLeod for their help. Dr. Toner Stevenson has reluctantly resigned from the Editorial Board; thank you, Dr. Stevenson.

Robert Marks, Editor
Balmain

1 "Post-literate" first appears in Marshall McLuhan's 1962 *The Gutenberg Galaxy*.

References

Horowitch R (2026) The age of reading is over. *The Atlantic*, August. https://www.theatlantic.com/magazine/2026/08/reading-crisis-postliterate-age/687618.

McLuhan M (1962) *The Gutenberg Galaxy: The Making of Typographic Man*. Toronto: University of Toronto Press.

Meyrowitz J (1985) *No Sense of Place: The Impact of Electronic Media on Social Behavior*, Oxford University Press.

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, pp. 4–24. ISSN 0035-9173/26/0104-21

A forty-year history of the development, commercialisation and implementation of the UNSW Vanadium Flow Battery

Maria Skyllas-Kazacos*, Ao Tang, and Yifeng Li*****

*School of Chemical Engineering, UNSW Sydney, NSW 2052, Australia
**Institute of Metal Research, Chinese Academy of Sciences, Shenyang, China
***VSUN Energy Pty Ltd, West Perth, WA 6005, Australia

E-mail: m.kazacos@unsw.edu.au

Abstract

Long-duration energy storage plays an important role in the energy transition as the deployment of renewable energy has been accelerating in recent years. Vanadium flow batteries invented 40 years ago at UNSW are regarded as one of the most favourable technologies for large-scale energy-storage applications because of their unique properties that allow energy and power to be decoupled. The development at UNSW from fundamental scientific studies of electrolytes, electrodes and membrane materials to stack design, sensors, modelling, simulation, system design and integration has formed an important foundation for the multi-GWh systems that have been installed and commissioned around the world particularly in China, Europe, North America and Australia to date. Many more GWh systems are in the pipeline to support the global energy transition to net zero, while on-going research efforts are continuing to further improve power density and to reduce cost.

Introduction

With the accelerating shift to renewable energy in recent years, the need for long-duration energy storage (LDES) is becoming increasingly apparent. While lithium-ion batteries are currently being deployed globally in large numbers, fires in multiple installations around the world are creating much concern about their safety (Chen Y et al. 2021). Despite the on-going reduction in lithium-ion battery costs in recent years, alternative battery technologies are being sought. Flow batteries are now emerging as a viable, cost-effective technology for LDES (Hou et al. 2024, Poli et al. 2024).

Flow batteries store energy in external electrolyte tanks, so energy storage capacity can be increased by simply adding more solution. As the electrolyte represents a fraction of the cost of the overall system, the cost per kWh of a flow battery decreases with increasing energy-storage capacity or energy-to-power ratios (E/P). Other attractive features are their very long cycle life and their inherent safety compared with lithium-ion batteries. In the case of Vanadium Flow Batteries (VFBs), a 20-year cycle life has already been demonstrated, while the indefinite life of the vanadium electrolyte gives it a significant residual value at the end of a project (Blume et al. 2024).

VFBs were invented more than 40 years ago by researchers at the University of New South Wales (UNSW) Sydney; while early commercial interest led to multiple field trials in the late 1990s and early 2000s, it is only in the last 5 years that the VFB technology has been experiencing a rapid growth in demand around the world. The use of vanadium in flow

batteries had been suggested earlier, but it was not until the first experiments conducted in 1984 by researchers at UNSW Sydney that vanadium redox couples were successfully employed in both halves of a working flow cell (Sum & Skyllas-Kazacos 1985, Sum, Rychcik & Skyllas-Kazacos et al. 1985). Earlier researchers had eliminated vanadium during their screening studies on the basis of cost, the relatively poor electrochemical reversibility of vanadium redox couples (NASA 1980), and the low solubility of pentavalent vanadium compounds in acidic solutions that would limit the practical energy density (Oei 1985). Despite these potential limitations, vanadium's multiple oxidation states made it a good candidate for a single element flow-cell that could overcome the problem of cross contamination observed with the NASA Fe/Cr battery (Thaller 1975; Hagedorn and Thaller 1982). It was therefore chosen as the starting point of the flow battery development program at UNSW.

In this review, we describe the early VFB research that was conducted at UNSW in the areas of electrochemical evaluation, electrolyte preparation, electrode screening and surface modification, membrane screening and modification, electrolyte thermal properties and stabilisation, modelling and simulation, sensors and battery management as well as battery design safety. The early licensing and field trials of the VFB are also described, together with the current commercialisation status around the world. Finally, we provide a future outlook on vanadium supply chain considerations, and important research and development areas for further performance improvements and cost reduction.

Early studies and first All-Vanadium Flow Cell

To evaluate the feasibility of an All-Vanadium Flow Battery, cyclic voltammetry was used to investigate the electrochemical reversibility of the V(II)/V(III) and V(IV)/V(V) couples in VCl_3 solutions in H_2SO_4 (Sum and Skyllas-Kazacos 1985; Sum, Rychcik and Skyllas-Kazacos 1985). Initial results showed poor reversibility for the V(II)/V(III) and V(IV)/V(V) couples, but, by roughening the glassy carbon electrode to produce surface defects, a dramatic improvement in the reversibility was observed. These results suggested that vanadium redox couples could demonstrate good reversibility for flow cell applications, but the low solubility of V(V) compounds still appeared to be a limitation.

With support from the Australian National Energy Research, Development and Demonstration Council (NERDDC), the UNSW research efforts were extended to explore the possibility of producing concentrated V(V) solutions by oxidising the much more soluble $VOSO_4$ compound in sulphuric acid. A 2 M vanadium electrolyte was successfully produced and tested, leading to the first "All-Vanadium Redox Flow Battery" patent being filed by UNSW in 1986 (Skyllas-Kazacos, Rychcik and Robins 1986).

Low-cost electrolyte process

Early work on the VFB involved the preparation of the vanadium sulphate electrolyte by dissolution of $VOSO_4$ in sulphuric acid[1], but the high cost of $VOSO_4$ ($300/kg at the time) was a major concern. While V_2O_5 (around $15/kg) could be sourced at a fraction of the price of $VOSO_4$, its very

1 $VOSO_4$ is vanadium exo sulphate, or vanadyl oxysulphate. [Ed.]

low solubility seemed prohibitive. A simple experiment led the way to the development of a number of vanadium compound dissolution processes that would provide the low-cost electrolyte production needed for commercial implementation of the VFB. Two graphite electrodes were immersed in a suspension of V_2O_5 powder in H_2SO_4 and a current was passed between them. After an hour or so, the powder was seen to be dissolving and a bright blue solution of V(IV) was being produced.

The simple beaker-based cell was then upgraded to a two-compartment cell that used a membrane to separate the two half-cells. The V_2O_5 suspension was placed in the negative half-cell while sulphuric acid was used in the positive half-cell. By adjusting the electrolysis time, the V_2O_5 powder could be electrochemically dissolved while further reducing the resultant solution to the 50:50 V(III):V(IV) mixture ($V^{3.5+}$) that forms the starting solution for the two half-cells of the all-vanadium redox flow cell. To avoid disintegration of the positive electrode during oxygen evolution in the positive half-cell, dimensionally stable anodes were considered, but were replaced by lead anodes to avoid the possibility of inert metal contamination of the vanadium electrolyte (Cao et al. 2018).

Several other chemical and electrochemical processes were subsequently developed, and a patent was filed in 1988 (Skyllas-Kazacos, McDermott and Kazacos 1988). This patent describes the process of dissolving and reducing an insoluble vanadium compound in an aqueous electrolyte using an electrochemical cell or by adding a chemical reductant such as oxalic acid. The use of the much cheaper V_2O_5 compound for electrolyte production was a major breakthrough that ensured the commercial viability of the VFB.

Later electrolyte process developments by the UNSW team included the reactive powder dissolution process where V_2O_3 and V_2O_5 powders are slowly mixed together in H_2SO_4 (or other supporting electrolyte) at around 80 °C in the appropriate molar ratio to produce a $V^{3.5+}$ electrolyte (Skyllas-Kazacos 2004). The commercial viability of this process depends on the relative cost of V_2O_3 compared to V_2O_5. Furthermore, the tendency for V_2O_3 to oxidise in air means that it must be stored in airtight containers before processing to avoid oxidation to V_2O_4 that would lead to a final vanadium electrolyte oxidation state higher than the $V^{3.5+}$ required for direct use in the VFB. An additional processing step to produce $V^{3.5+}$ would then be needed.

To overcome the higher cost of V_2O_3, an alternative process was developed where V^{3+} is electrochemically generated in the negative compartment of an electrolysis cell and then used as a reductant to dissolve the V_2O_5 powder into solution to produce the $V^{3.5+}$ electrolyte product. A portion of the $V^{3.5+}$ solution is returned to the electrolysis cell to produce more V^{3+} reductant that is fed into the mixing tank to dissolve more V_2O_5 powder in a continuous process (Kazacos and Skyllas-Kazacos 1994; Skyllas-Kazacos 1998).

Each of these methods was further extended to evaporate most of the water to produce a vanadium electrolyte concentrate or slurry that would reduce transportation costs (Skyllas-Kazacos and Kazacos 2002). The concentrate or slurry can then be reconstituted on delivery by simply adding water.

With improvements in processing methods for vanadium compounds in recent

years, the cost of V_2O_3 has become comparable to that of V_2O_5, making the V_2O_3/V_2O_5 reactive powder dissolution method simple and cost effective. This method and the suspended powder electrolysis method described previously are currently being used by most vanadium electrolyte producers to supply $V^{3.5+}$ electrolyte to the VFB manufacturers around the world.

Early electrode screening and modification

Various electrode materials were screened by the UNSW group in the 1980s and '90s, and graphite and carbon felts were selected as the most suitable electrode materials for both the positive and negative half-cells in the VFB (Rychcik and Skyllas-Kazacos 1987). Provided that the cell is protected from excessive overcharge to avoid corrosion of the positive electrode under oxygen evolution conditions, sufficient chemical and mechanical stability can be achieved when combined with a suitable low cost and chemically resistant bipolar substrate material (Zhong et al. 1997).

A wide range of carbon and graphite felts was tested, and certain materials were found to give much higher voltage and energy efficiencies than others during charge/discharge cycling in the VFB (Zhong et al. 1993). X-ray photoelectron spectroscopy (XPS) demonstrated the role of surface oxide functional groups on the electrochemical behaviour of the electrode material. Thermal, chemical and electrochemical treatments of carbon and graphite felts were therefore evaluated for the purpose of enhancing the electrochemical activity of these materials as electrodes in the VFB (Sun and Skyllas-Kazacos 1991; Sun and Skyllas-Kazacos 1992a; Sun and Skyllas-Kazacos 1992b). XPS measurements showed that the amount of chemisorbed oxygen on the activated graphite felt surface increased dramatically after treatment. The improved performance of the treated graphite felts for the vanadium redox reactions was thus attributed to the increased surface concentration of the C–OH and C=O functional groups that were formed during activation. In another UNSW study, Mn, Te and In modified graphite fibre electrodes exhibited an improved performance compared to the untreated one (Sun and Skyllas-Kazacos 1992a).

The above studies also demonstrated that with appropriate thermal, chemical and electrochemical treatments, both the wettability and electrochemical activity of low-performance carbon and graphite felts can be substantially improved (Sun and Skyllas-Kazacos 1992b). This was a critical milestone in achieving cost reduction for commercial application of the VFB technology, since lower-cost carbon felts could be used in stack production. Thermal treatment is currently used by most carbon felt manufacturers to produce high-performance felt electrodes for VFB applications.

Early membrane screening and modification

During the early stages of the UNSW vanadium battery R&D program, low-cost, chemically stable commercial membranes were not readily available. Although Nafion membranes showed excellent chemical stability in the highly oxidising V(V) solutions, they were too expensive for practical application at the time. Various off-the-shelf membranes and separators were therefore initially tested in a laboratory-scale flow cell to determine coulombic, voltage and energy efficiencies at different current densities, but suitable characterisation methods were

needed to assess the suitability of different membranes for use in the VFB. The UNSW group therefore developed a range of methods to characterise each membrane in terms of its ionic conductivity, vanadium-ion selectivity, water-transfer properties and chemical stability in the highly oxidising V(V) positive half-cell solution (Chieng, Kazacos and Skyllas-Kazacos 1992; Mohammadi and Skyllas-Kazacos 1995; Mohammadi, Chieng and Kazacos 1997; Mohammadi and Kazacos 1997; Sukkar and Skyllas-Kazacos 2003a; Sukkar and Skyllas-Kazacos 2003b; Sukkar and Skyllas-Kazacos 2004).

Of all the membranes tested, perfluorinated membranes such as Nafion showed the best long-term stability, with negligible weight loss observed even after five years of immersion in V(V). Despite its high chemical stability, however, Nafion was not initially considered for use in the VFB because of its high cost at the time, which would limit its competitiveness in commercial systems (Mohammadi and Kazacos 1997).

Water transfer across the membrane was also found to be a significant issue during charge–discharge cycling of the VFB. Early water-transfer studies showed a net water transfer into the negative half-cell when using an anion-exchange membrane (AEM), and into the positive half-cell in the case of cation-exchange membranes (CEMs) (Mohammadi, Chieng and Kazacos 1997). To address the shortcomings of both AEM and CEM membranes, the amphoteric ion-exchange membrane (AIEM) was first proposed for the VFB (Mohammadi and Kazacos 1996). Sulphonation of the anion-exchange membrane was used to introduce cation-exchange groups in order to reduce water transfer and improve performance in the VFB. A range of anionic and cationic polyelectrolytes were subsequently incorporated into a Gore-Select cation-exchange membrane to successfully reduce water transfer in the VFB (Sukkar and Skyllas-Kazacos 2003c).

Electrolyte thermal stability

The composition of the vanadium electrolyte plays a critical role in determining the practical operating-temperature range of the VFB. While V(II), V(III) and V(IV) sulphates show increased solubility at higher temperatures and lower sulphuric-acid concentrations (due to the common-ion effect), V(V) undergoes thermal precipitation to V_2O_5, and this increases at higher temperatures and lower proton concentrations (Kazacos, Cheng and Skyllas-Kazacos 1990). A typical operating-temperature range of 10–40 °C was thus proposed for 2 M vanadium in 5 M total sulphate electrolytes, but to extend this range and minimise precipitation risks in extreme climates, lower vanadium concentrations (1.5–1.8 M) are usually employed (Kazacos, Cheng and Skyllas-Kazacos 1990; Skyllas-Kazacos et al. 2016). The sulphuric-acid concentration is also adjusted to maintain an approximately 2:5 vanadium-to-total-sulphate ratio, but this can be further modified to suit climate conditions. For example, in very hot climates, a higher sulphuric-acid concentration can reduce the rate of V_2O_5 precipitation (Kazacos, Cheng and Skyllas-Kazacos 1990; Skyllas-Kazacos et al. 2016).

The operating-temperature range can also be extended without lowering the total vanadium concentration. State-of-charge (SOC) control avoids precipitation of V(V) at high temperatures and V(II)/V(III) at low temperatures (Corcuera and Skyllas-Kazacos 2012), while precipitation inhibitors

extend the induction time for precipitation in supersaturated electrolytes (Cao et al. 2018; Mousa and Skyllas-Kazacos 2015; Kausar, Mousa and Skyllas-Kazacos 2016). While an extensive list of additives was shown to dramatically increase induction time for precipitation of one or more of the vanadium ions in the electrolyte, a range of phosphate and ammonium compounds was found to be effective inhibitors for both half-cell electrolytes (Skyllas-Kazacos 1998; Kazacos and Skyllas-Kazacos 1994; Skyllas-Kazacos and Kazacos 2002; Mousa and Skyllas-Kazacos 2015; Kausar, Mousa and Skyllas-Kazacos 2016).

In the case of state-of-charge control, a battery management system was used to continuously monitor SOC and electrolyte temperature. The controller varies the upper and lower SOC limits to stay within a safe operating range that would minimise the risk of precipitation at low or elevated temperatures (Corcuera and Skyllas-Kazacos 2012). For simplicity, emergency discharge could involve turning on the pumps during standby to increase the rate of self-discharge. This avoids the capital and maintenance costs associated with active cooling systems and allows the VFB to be effectively used in hot climates (Skyllas-Kazacos 1986).

While restricting the SOC range can avoid precipitation of the vanadium ion at temperature extremes, this will inevitably reduce the practical energy density of the flow battery and limit available capacity. A preferred solution is to stabilise the vanadium electrolyte by using precipitation inhibitors. Hundreds of potential additives were screened and a large number of chemicals were found to be effective precipitation inhibitors for the V(II), V(III), V(IV) ions at low temperatures, while others were shown to reduce the rate of thermal precipitation of V(V) at elevated temperatures (Cao et al. 2018; Mousa and Skyllas-Kazacos 2015; Kausar, Mousa and Skyllas-Kazacos 2016).

Although the exact mechanism of precipitation inhibition is not fully understood, it is believed that these additives affect the precipitation rate by being adsorbed on the crystal surface to form a monomolecular layer that changes the crystal habit and growth rate. It is possible that the additive molecules are adsorbed on certain edges or corners that play a special part in the crystal growth. A small amount of the additive could thus be enough to block most of these active sites, causing a significant decrease in the crystal growth rate.

From the additive screening studies, ammonium phosphate and ammonium sulphate were selected for their ability to stabilise both half-cell solutions, and these were used to prepare a 3 M vanadium electrolyte that was successfully tested in a laboratory cell and in the UNSW electric golf cart in the late 1990s (Roe, Menictas and Skyllas-Kazacos 2015). While the 3 M solution was stable during charge-discharge at room temperature, its use at temperature extremes over a broad SOC range requires further optimisation.

Other UNSW Vanadium Redox Flow Cell research highlights

With grant funding from the NERDDC (Skyllas-Kazacos 1986; Skyllas-Kazacos 1989; Skyllas-Kazacos and Kazacos 1992) and the NSW Office of Minerals and Energy (Skyllas-Kazacos 1989; Skyllas-Kazacos 1991), as well as industry funding from Agnew Clough Ltd, Mitsubishi Chemicals and Kashima Kita Power Corporation, the VFB was taken from the initial concept stage at

UNSW in 1983 through the development and demonstration of several 1–5 kW prototypes over a 40-year period at UNSW in both stationary and mobile field trials. As part of this program, a wide range of research projects was undertaken. In addition to the electrolyte production, electrode, membrane and additive studies described above, other areas of research at UNSW have included electrolyte characterisation, measurement of electrolyte density, viscosity and conductivity as a function of vanadium and acid concentration, temperature and state-of-charge (Kazacos, Cheng and Skyllas-Kazacos 1990; Skyllas-Kazacos et al. 2016). The variation in these electrolyte properties as a function of state-of-charge can be used for the individual state-of-charge monitoring of each of the two half-cell solutions to identify any state-of-charge imbalance that will lead to capacity loss during long-term charge-discharge cycling. This allows the battery management system to initiate an electrolyte remixing operation to restore capacity. A state-of-charge monitoring method based on solution absorbance was also developed using the distinct colours of the different vanadium ions and the varying spectroscopic properties of each half-cell solution at different states-of-charge (Skyllas-Kazacos et al. 1989; Corcuera and Skyllas-Kazacos 2012).

While simple remixing of the two half-cell solutions can restore capacity loss associated with the different rates of vanadium ion transfer across the ion exchange membranes, capacity loss caused by side reactions such as hydrogen evolution during charging or air oxidation of V(II), requires chemical or electrochemical rebalancing. The use of oxalic acid was proposed as a chemical reductant which is added to the V(V) solution in the positive half-cell to partially reduce it to V(IV) in order to match the state-of-charge of the negative half-cell electrolyte. Oxalic acid was selected since it is oxidised to carbon dioxide and water during reaction, leaving no residual biproducts in the electrolyte (Skyllas-Kazacos et al. 1989).

Another major activity in the 1990s was the development of low-cost bipolar electrodes to replace the heavy and expensive graphite plates used as substrates for the graphite felts in stack assembly. Novel conducting plastic bipolar plate formulations led to manufacturing trials of bipolar electrode assemblies that were used in the fabrication of several 1–5 kW prototype stacks for the first field trials of the VFB in the 1990s. By combining a thermoplastic polymer such as high-density polyethylene, an elastomeric polymer, and a conductive filler material, a flexible, conducting plastic electrode with good physical strength, chemical stability and electrical conductivity could be produced (Zhong et al. 1993). This material was used to produce multiple bipolar electrodes by heat bonding graphite felt onto each side of the conducting plastic plate as illustrated in Figure 1. These bipolar plates were employed in the fabrication of several 5 kW prototype stacks at UNSW during the 1990s.

Fig 1: Bipolar electrode fabricated by heat bonding graphite felt sheets onto each side of the conducting plastic sheets developed by UNSW.

In recent years, mathematical modelling has been a major focus area at UNSW and has been used to assist in stack design optimisation and battery performance prediction. Mathematical models were therefore developed to predict shunt currents and pumping energy losses in order to assist with optimal flow-frame and stack design (Tang et al. 2013; Tang, Bao and Skyllas-Kazacos 2014; Li, Skyllas-Kazacos and Bao 2016). Capacity losses associated with the differential transfer of vanadium ions across the membrane were also modelled for a range of different membranes (Tang, Bao and Skyllas-Kazacos 2011), allowing the development of automated remixing and other procedures to restore capacity.

Thermal modelling of VFBs under a range of operating conditions and environmental temperatures has been used to predict temperature gradients in cell stacks during standby using different types of membranes. When the pumps are turned off during standby, the transfer of vanadium ions across the membrane gives rise to self-discharge and heat generation within the stack that would normally be removed by electrolyte flow. This leads to an increase in the stack temperature, the rate of which depends on the membrane, stack design and ambient temperature. The shapes of the stack temperature gradient reflect the rate of heat generation caused by the self-discharge reactions during standby and the rate of heat dissipation to the environment. Thermal modelling and simulation studies were conducted for various climate conditions in both open and in containerised VFB systems to optimise insulation and design of passive and active cooling systems (Tang, Bao and Skyllas-Kazacos 2012; Yan, Skyllas-Kazacos and Bao 2017; Shu et al. 2023a; Shu et al. 2023b).

An electrical safety analysis of different VFB configurations during electrolyte leakage has also been undertaken to identify measures needed to avoid potential safety hazards in the event of an electrolyte leak (Shu et al. 2025). Simulation studies showed that battery voltage plays a critical role in determining the magnitude and safety hazard caused by current that would flow through the body if a person were to step into a pool of electrolyte in ionic contact with the battery stacks. A wide range of conditions were simulated and results showed that safety risks associated with low-voltage battery configurations were low, but safety hazards increased with high-voltage configurations on wet concrete in close proximity to the leak if the negative terminal of the battery stack was grounded. While the body current under low voltage configuration is expected to be within a safe range, protective shoes are highly recommended under high-voltage configurations.

To ensure a long operating life of any flow battery system, a sophisticated battery management system is essential. Using a range of sensors and monitoring systems, a battery management system is currently being developed to detect the condition of the battery in real-time and remotely implement all the necessary actions to prevent precipitation at extreme temperatures, restore capacity and optimise electrolyte flow rate for maximum overall efficiency (McCloy et al. 2022).

The vanadium-halide redox flow cell was proposed in 2002 in an effort to increase the energy density of the VFB. The higher solubility of vanadium ions in mixed halide solutions was the initial motivation, but the high cost of bromine complexing agents to prevent bromine vapours during

charging has been an obstacle to its commercial development (Skyllas-Kazacos 2002; Skyllas-Kazacos, Mousa and Kazacos 2003; Kazacos and Kazacos 2006; Skyllas-Kazacos et al. 2010).

The vanadium-oxygen fuel cell (VOFC) continues to be an attractive technology that has the potential to more than double the energy density of the VFB by replacing the positive half-cell electrolyte tank with an air electrode (Menictas and Skyllas-Kazacos 2011; Risbud et al. 2019). Since the VOFC only requires half of the vanadium raw material per kWh of energy produced, its cost could also be potentially halved compared with the VFB. While the negative vanadium half-cell functions well, however, an efficient non-noble metal gas diffusion electrode is still needed to reduce costs and increase energy efficiency. This is an on-going area of research at UNSW.

Early licensing and field trials

In 1987, an exclusive international licence to the VFB technology was granted by UNSW to the Australian vanadium mining company Agnew Clough Ltd., leading to 3 years of industrial funding that enabled UNSW to fabricate and test the first 1 kW VFB laboratory prototype which produced an overall energy efficiency of 90% at a current density of 20 mA/cm^2 (Skyllas-Kazacos et al. 1991). Unfortunately, due to financial difficulties, the company was forced to relinquish its licence in 1991, allowing the University to seek other licensees (Skyllas-Kazacos 2023).

A Thai-based construction company was granted a licence to the technology in 1993 for the South-East Asian region (Skyllas-Kazacos 2023). This led to the installation of a 5 kW/15 kWh VFB in a vanadium-powered solar demonstration house just outside Bangkok (Figure 2 left).

Around the same time, Kashima-Kita Electric Power Corporation, a subsidiary of Mitsubishi Chemical Corporation, became interested in the VFB technology as a way of using the vanadium waste extracted from the soot produced by the Orimulsion fuel derived from the vanadium-rich Venezuelan pitch employed at the power station. A 5-year R&D collaboration program between the UNSW research group and the Japanese companies led to further advances in solution chemistry, stack design, improved materials and control systems (Skyllas-Kazacos 2023).

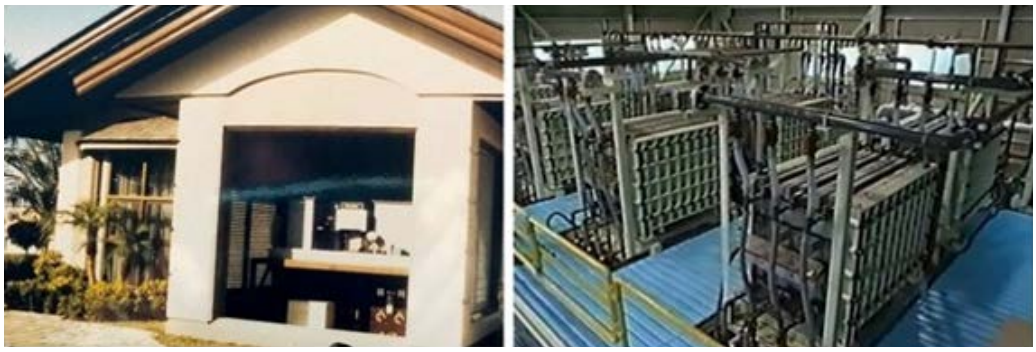

Fig 2: Early VFB Field Trials: (Left) VFB installation in Demonstration Solar House in Bangkok in mid-1990s; (Right) 200 kW/800 kWh VFB at Kashima-Kita Electric Power Station in Japan in 1996–97. Photos courtesy Rui Hong (left) and Chris Menictas (right).

By 1997, Kashima-Kita Electric Power Corporation had progressed to the design and installation of the first multi-kWh VFB in a load-levelling field trial at their power station in Japan (Figure 2 right). The system was a 200 kW/800 kWh battery that was tested over a wide range of conditions and an overall energy efficiency of 80% was achieved, demonstrating the excellent scalability of the VFB.

In the mid-1990s, the UNSW team installed a VFB in an electric golf cart (Figure 3). Using the UNSW ammonium phosphate-based proprietary stabilising agent, a 3 M vanadium electrolyte was produced by the UNSW team and installed in the electrolyte tanks that were located below the seats. In this and other prototype batteries and field trials at UNSW, the negative electrolyte was protected from air oxidation of the V^{2+} by the use of a thin (1 cm) layer of paraffin oil on the surface to act as a barrier to oxygen (Mousa and Skyllas-Kazacos 2015; Kausar, Mousa and Skyllas-Kazacos 2016). Paraffin oil is chemically stable in the negative half-cell solution, but its use requires greater care during the periodic remixing of the electrolytes to restore capacity. The golf cart was subjected to a large number of driving tests around the UNSW campus for several years before it was donated to the Powerhouse Museum in Sydney in 2005.

In the mid-1990s, the UNSW team installed a VFB in an electric golf cart (Figure 3). Using the UNSW ammonium phosphate-based proprietary stabilising agent, a 3 M vanadium electrolyte was produced by the UNSW team and installed in the electrolyte tanks that were located below the seats. In this and other prototype batteries and field trials at UNSW, the negative electrolyte was protected from air oxidation of the V^{2+} by the use of a thin (1 cm) layer of paraffin oil on the surface to act as a barrier to oxygen (Mousa and Skyllas-Kazacos 2015; Kausar, Mousa and Skyllas-Kazacos 2016). Paraffin oil is chemically stable in the negative half-cell solution, but its use requires greater care during the periodic remixing of the electrolytes to restore capacity. The golf cart was subjected to a large number of driving tests around the UNSW campus for several years before it was donated to the Powerhouse Museum in Sydney in 2005.

Fig 3: VFB powered electric golf cart at UNSW in 1997 (Left). Compact stack and tanks with hidden plumbing designed for the golf cart (Right). Photos courtesy Maria Skyllas-Kazacos UNSW Sydney.

The UNSW VFB technology was sold to the Australian ASX-listed company Pinnacle VRB in 1997, and several demonstration systems were installed by Pinnacle in South Africa, King Island in Australia, and in the USA, using battery stacks acquired from Sumitomo Electric Industries, electrolyte supplied by Highveld, and balance of plant sourced from a range of local suppliers. In 1999, Pinnacle VRB entered into a new licence agreement with Sumitomo Electric Industries that had been evaluating the UNSW VFB technology for several years and had built a 450 kW/1 MWh vanadium redox battery load leveling demonstration system at the Kansai Electric Power Plant in Japan, as well as several other VFB systems for wind energy storage and other stationary applications.

Pinnacle also licensed the VFB technology to the Canadian-listed company Vanteck VRB for the African continent and a 250 kW VFB system was commissioned by Vanteck in conjunction with the South African utility, ESCOM. A 250 kW/2 MWh VFB system was supplied by Pinnacle to US-based PacifiCorp in the early 2000s. The system was used by PacifiCorp to supply peak power capacity (charging in the off-peak hours) and provide end of line voltage support (supplying up to 250 kVAR of reactive power) in a remote area in southeastern Utah.

Despite the successful field testing of the VFB in a wide range of applications by SEI and the substantial early investment and capital raising, however, Pinnacle was unable to fight off take-over attempts and the company and all the VFB patents were eventually taken over by a Canadian company that successfully completed a number of projects in South Africa and the USA in the early 2000s. A short-lived peak in vanadium prices in 2007 followed by the 2008 Global Financial Crisis, however, eventually led to its insolvency and its assets being acquired in a fire sale by a Chinese-based energy storage developer.

After the sale of the UNSW patents in 1997, the UNSW team continued to explore alternative redox couple combinations and the Vanadium Halide Flow Battery (often referred to as the Generation 2 Vanadium Halide Battery) was patented in the early 2000s (Skyllas-Kazacos 2002; Skyllas-Kazacos, Mousa and Kazacos 2003; Skyllas-Kazacos and Kazacos 2006; Skyllas-Kazacos et al. 2010). In 2005, a new start-up company V-Fuel, was established

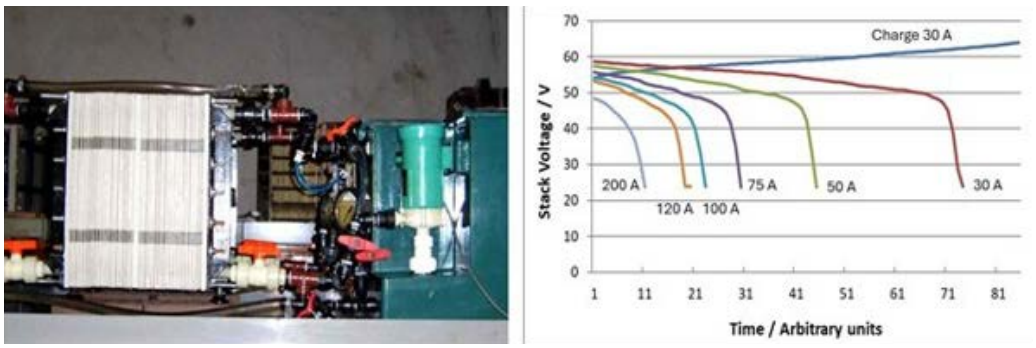

Fig 4: Left: L 5–10 kW V-Fuel stack using conducting plastic bipolar electrode with heat bonded graphite felt; Right: Discharge curves for V-Fuel stack at different currents after charging at 30 A. Photo courtesy Maria Skyllas-Kazacos.

with initial investment funding from the Victorian government-funded venture capital company, CEGT (Centre for Energy and Greenhouse Technologies) and a grant from the Australian government's Renewable Energy Development Initiative (REDI). The start-up successfully developed an improved perfluorinated membrane treatment process, identified suitable bromine complexing agents to bind the bromine produced at the positive electrode during charging (Kazacos and Kazacos 2012) and optimised a conducting plastic bipolar electrode assembly (Kazacos and Kazacos 2013). These electrodes were used in the fabrication of a 5–10 kW stack (Figure 4) that also included specially treated perfluorinated membranes. The company also designed a novel stack (Figure 5) that employed laser-welded bipolar electrode assemblies and could also be externally sealed by laser welding. However, the sudden jump in global vanadium price in 2007 combined with the effects of the Global Financial Crisis of 2008 prevented the company from attracting the necessary investment capital, so it was wound up and the patents transferred to UNSW's commercial arm New South Innovation.

When the basic UNSW All-Vanadium Flow Battery patent expired in 2006, groups around the world were free to develop and commercialise the VFB technology using the extensive research foundations provided by UNSW. One of the first companies to develop commercial systems using the UNSW technology outside of Japan was Cellstrom in Austria with its containerised 120 kWh CellCube system while multiple MW-scale field trials were conducted by Sumitomo Electric Industries in the early 2000s, verifying the efficiency and durability of the VFB in a wide range of applications.

Fig 5: V-Fuel integrated battery on a trolley, with pumps and piping incorporated inside tanks.

Current status

Since the expiration of the basic UNSW VFB patent in 2006, multiple companies have embarked on commercialisation of the VFB technology, with excellent performance demonstrated during long-term operation. More than 200,000 cycles were achieved in a 3-year VFB field trial by Sumitomo in the early 2000s, while a recent evaluation of a CellCube system has verified negligible capacity loss or performance deterioration after 10 years' operation (Li Y et al. 2024). These VFB companies continue to show a strong presence in the market, but others were forced to merge, restructure, or in some cases wind up, because of the slow growth of the energy storage market. Fortunately, this has changed dramatically over the last 5 years with the global push for net zero targets and increased renewable

energy uptake. While Li-ion batteries have attracted most of the attention to date due to their successful cost reduction in recent years, on-going safety issues and the growing need for longer duration energy storage are now starting to shift interest towards flow batteries.

At present, there are around 20 companies manufacturing VFBs around the world, with several of these being established in Australia to meet the rapidly growing demand for energy storage systems to achieve net zero targets. Since 2010, many new research groups around the world have been extending the original UNSW VFB research and have contributed significant improvements and cost reductions that are helping to enhance VFB performance and economic viability. The UNSW Vanadium Redox Battery technology is currently being successfully commercialised in a range of stationary applications such as wind, solar and load-leveling installations around the world and the demand for these systems is growing rapidly in China, North America, Japan, Europe and Australia.

Current R&D, commercialisation and implementation in China

Vanadium redox flow batteries have seen a fast development in China over the past decade, especially since 2020 when China announced its carbon neutrality strategy aiming to reach carbon peak in 2030 and carbon neutrality in 2060. With a skyrocketing deployment of renewable energy, China achieved a milestone in 2023 where clean energy generation (e.g., solar, wind, etc.) had for the first time broken through 50% of the newly installed power generation capacity. To sustain grid stability and security, therefore, long-duration energy storage technologies have been promoted to a high strategic level of national development, among which the VFB, with its high safety and flexibility in deployment, has been regarded as one of the most effective technologies for long-duration energy storage.

Fig 6: 200 MW/1000 MWh VFB in Xinjiang, China. Published with permission from Xingchen New Energy.

In 2024, over 440 MW/1.74 GWh of VFB systems were successfully installed and connected to the grid, where the capacity has been increased by 600% compared to 2023. Notably, this included the world's largest VFB installation of 250 MWh/1000 MWh system in Xinjiang Province. In addition, a 50 MW/200 MWh VFB system was recently installed in Gansu Province, while a 100 MW/400 MWh system has also been installed in Inner Mongolia. More recently, a 5-hour long-duration VFB project has been completed in Xinjiang, which achieves a total 200 MW/1000 MWh capacity as illustrated in Figure 6. This boom in VFB application is not only driven by the national strategies but is also attributed to the technical advances in VFB technology. Taking 4 hours storage capacity as an example, the average bid-winning price of a VFB system has significantly fallen to 2000 RMB/kWh in 2024. In the next few years, a further cost reduction to 1500 RMB/kWh is

expected, due to a well-developed material supply chain, enhanced battery efficiency and automated stack assembly (Figure 7), that has been dramatically reducing manufacturing costs.

Fig 7: Automated stack assembly at Rongke Power, China (reproduced with permission of Rongke Power).

Of all the VFB components, the vanadium electrolyte is still the most expensive part and may account for over 60% of capital expenditure. China owns ca. 37% of the total vanadium resources of the world. Currently, 80% of the vanadium electrolyte is produced by using vanadium from the by-product of the steel industry. Recent developments of short-processing vanadium electrolyte production technology (Zhang et al. 2022) have largely promoted the productivity, which has recently realised an electrolyte production capacity of over 60,000 m^3/year. With the shrinkage of the steel industry, however, other vanadium resources, such as stone coal, will be needed to meet the increased demand of vanadium electrolyte.

Apart from electrolyte, the power density of the VFB stack has been significantly increased in recent years. Large stacks such as 32 kW modules, can now deliver an energy efficiency over 80% at 200 mA/cm^2 constant charging-discharging mode, effectively reducing the material usage and lowering the power cost of the VFB.

Commercialisation and implementation in Europe and North America

The development activity of VFBs in Europe can be traced back to early 2000 (Mardilovich and Harrer 2023). In 2003, the first VFB system in Europe was installed, with 1 kW power and 50 kWh storage capacity. The system was operated in an off-grid configuration charged by solar and wind. The findings from the demonstration unit led to further research and development activities to improve the battery stack performance. A few years later, the concept of containerised VFBs started to emerge offering a standalone turnkey solution that fits all the necessary components, including tanks, power stacks, pumps and power electronics into a closed structure that has the same dimensions as a standard shipping container. Containerised units have gained interest because they simplify deployment on site, enable easy relocation and improve ingress protection. Early versions of containerised systems typically have power ratings ranging from 10 to 30 kW and capacity ranging from 40 to 130 kWh. One notable installation was the 10 kW/100 kWh system commissioned in Lichtenegg Energy Research Park, Austria, as part of a micro grid demonstration project in 2010. The Lichtenegg system has shown very little degradation after more than a decade of continuous operation according to a study in 2023 (Li et al. 2023).

Since 2010, research and development activities have continued to grow, with a strong focus on cost reduction while improving the system design and performance. More VFB manufacturers around the world have started to emerge, offering containerised VFB products based on a

modular concept. For example, some designs integrate multiple small VFB modules into a standard shipping container unit. The containerised units can be stacked in multiple levels to minimise the total footprint of large-scale installations. Studies have shown that the normalised footprint of a VFB installation can be less than a lithium battery installation of the same capacity due to the non-flammable nature of a VFB and the ability to be stacked (Reber, Jarvis and Marshak 2023). Some designs also are certified for road and sea transport with electrolyte filled in the tanks, which eliminates filling on site and simplifies the commissioning process (Cols 2025). Many MW scale installations in Europe and North America include:

- the 5 MWh Energy Superhub Oxford project in the UK (Cols 2024)
- the 1 MW/8 MWh project in Spain supported by the government energy research institute CIUDEN (Huh and Han 2025)
- a 2 MW/20 MWh system installation in Fraunhofer ICT as part of a European-funded project (Seiler et al. 2025)
- a 2 MW / 8 MWh system for G&W Electric in Illinois to use in a resilience micro-grid application (Whitehead et al. 2025)
- an 8.4 MWh Chappice Lake Solar Storage project in Canada for utility application (Cols 2024)
- a 2 MW/8 MWh system in San Diego California for grid stability application (Figure 8).

Fig 8: VFB projects in Europe and North America (top left: system in Lietenegg, top right: system in Spain, bottom left: system in Oxford Energy Superhub, bottom right: system in San Diego).

Besides offering DC VFB systems, several suppliers focus on the development and manufacture of the battery cell stack, the most important component in a VFB system. One important aspect includes optimising cell stack design to improve the stack performance, such as reducing the internal resistance by using a zero-gap reactor design with interdigitated flow field targeting high-power applications (Austing 2025). The stack manufacturing process has also been improved by utilising a fully automated stack assembling process to lower the cost (Seipp 2025). The overall stack performance has dramatically improved over the past decade. According to a study conducted in 2022, the average stack round trip efficiency is above 80% and some can reach as high as 88%, comparable with lithium-ion batteries (Li, Kienbaum and Lüth 2023).

While the development activities on the VFB stacks and containerised modules are expanding, with the increase in renewable energy penetration, the use of VFB is shifting towards longer duration (8–12 hours) due to better economics compared to lithium batteries. Several companies focus on system integration, including integrating containerised systems using a high voltage concept to ensure a seamless integration with existing proven power conversion systems (PCSs). The high-voltage concept offers better system efficiency by eliminating the need for DC-DC converters while still minimising shunt current losses in the common manifolds. Non-containerised solutions offering better flexibility to fit different applications and the ability for future expansion have gain interest, such as the 200 MW/800 MWh VFB system constructed in a building in Dalian China and the 15 MW/60 MWh system for the Hokkaido electric power network. In 2025, an AI data centre began construction in Laufenburg (Switzerland), which will include an 800 MW/1.6 GWh VFB in the basement of the building to support the grid. This is one of the largest flow battery projects in Europe and the world (Wyss 2025).

Activities in Australia

Australia's energy landscape is currently undergoing a rapid transition aiming to reach 82% renewable energy by 2030, according to The Department of Climate Change, Energy, the Environment and Water Australia (Edis 2025). Scalable and flexible long-duration energy storage solutions are critical to reach this ambitious goal. In 2024, the Australian Government released the National Battery Strategy as a key part of the government's Future Made in Australia agenda. The strategy outlines how the Australian Government will support domestic battery industry to:

- improve Australia's energy security
- ensure its place in global battery supply chains
- drive economic growth
- create positive social, economic and resilience outcomes for all Australians.

As of 2024, 35% of total electricity across the country is generated from renewable sources (DCCEEW 2025). VFBs are recognised as a favourable technology for large-scale energy storage systems to manage the intermittent supply and demand fluctuations of renewables. Vanadium, as an important critical mineral, plays a vital role in the VFB value chain and is a key enabler of VFB deployment. Australia has the world's best vanadium reserves, including: deposits in Meekatharra, West Australia;

Fig 9: Vanadium electrolyte manufacturing facilities in Western Australia (left) and Queensland (right).

in Julia Creek, Queensland; and Speewah, Western Australia. While the vanadium mining projects are being developed, several plants have been setup to produce vanadium electrolyte, including a 35 MWh annual production capacity plant in Western Australia commissioned in 2024 using the mixed oxide method (Skyllas-Kazacos 2004) and a 35 MWh annual production capacity plant in Queensland commissioned in 2023 using the V_2O_5 powder electrolysis method (Skyllas-Kazacos and Kazacos 2002). (Figure 9).

Besides electrolyte production, several important VFB pilot systems have been installed in Australia, including the 2 MW/8 MWh Yadlamalka Energy project in South Australia, the 250 kW/750 kWh Berrinba project in Queensland, and the 78 kW/230 kWh Horizon Power project in Kununurra, Western Australia (Noack et al. 2025). In early 2025, the WA Labor government announced the plan to build a 50 MW/500 MWh VFB in mining town Kalgoorlie by 2029 to support the electricity grid and is committing AUD$150 million funding towards the project.

Conclusion and outlook

Renewable energy generation has grown rapidly around the world and is expected to continue growing, as countries aim to reduce fossil fuel consumption and the associated greenhouse gas emissions. LDES plays an important role in balancing the power grid demand and supply associated with the inherent intermittent nature of renewable sources. VFBs, a 40-year-old technology invented at UNSW, are regarded as one of the most effective technologies for large-scale energy storage and long-duration applications. The early development of the VFB at UNSW focused on low-cost electrolyte manufacturing processes, electrolyte stabilisation, screening and modification of electrodes and membranes, novel bipolar stack design, fabrication and testing, modelling and simulation studies, as well as system monitoring and control. The fundamental research conducted at UNSW has been the key enabler for the commercialisation and deployment of the technology around the world, particularly in Japan, China, Europe, North America and Australia. To date, multi-GWh capacity VFBs have been installed and connected to grids around the world and many GWh projects are in the pipeline. The scale of VFB deployment is expected to continue its growth and form an important part of the global energy transition.

As MW- and GW-scale VFBs continue to be implemented around the world, the supply of vanadium will become a critical factor that is already being addressed, with

new vanadium mines expected in Australia and elsewhere in the coming years. The logistics of transporting thousands of cubic metres of electrolyte to remote locations will however require considerable research to reduce transportation costs. The production of an electrolyte concentrate or slurry, first proposed by Skyllas-Kazacos and co-workers in the 1990s (Kazacos and Skyllas-Kazacos 1994, Skyllas-Kazacos and Kazacos 2002), is currently being considered as a possible solution. By producing an electrolyte concentrate or slurry, the electrolyte weight and/or volume could be reduced by up to 50%, saving on costs of shipping and transportation over long distances. The concentrate or slurry can then be reconstituted by simply adding water prior to use. Further research on the storage, transportation, and reconstitution of the electrolyte concentrate is currently being undertaken, as is the effect of electrolyte impurities on battery performance and gassing side reactions. Electrolyte impurities are a critical issue which requires further research to determine whether lower purity vanadium raw materials could be used in the VFB to allow significant cost reduction.

Another important area of future research is the development of non-fluorinated membranes that can help to address increasing global concerns relating to "forever" chemicals. Many research groups and companies are developing such membrane materials (Huynh et al. 2025) but they still need to undergo long-term testing to verify their durability in the VFB for 20 years or more. The operation of VFBs in extreme hot or cold climates is also an area of further research and development. While such systems are already operating in many hot or cold regions of the world, the need to use active heating or cooling has been a drawback, since it reduces overall energy efficiency. New electrolyte formulations, including the mixed sulphuric acid/hydrochloric supporting electrolyte developed by Pacific Northwest National Laboratories (Li et al. 2011), are being considered, but the possible evolution of chlorine during charging requires special chlorine scrubbing equipment to minimise corrosion and prevent human exposure.

References

Austing JG (2025) Business cases for flow batteries. *IFBF Vienna Austria*.

Blume N, Zielinski O, Becker M and Minke C (2024) Prospective life cycle assessment of chemical electrolyte recycling for vanadium flow batteries: A comprehensive study. *Energy Technology* 12, p. 2300750. DOI: 10.1002/ente.202300750.

Cao L, Skyllas-Kazacos M, Menictas C and Noack J (2018) A review of electrolyte additives and impurities in vanadium redox flow batteries. *Journal of Energy Chemistry* 27: 1269–1291.

Chen Y et al. (2021) A review of lithium-ion battery safety concerns: The issues, strategies, and testing standards. *Journal of Energy Chemistry* 59: 83–99. DOI: 10.1016/j.jechem.2020.10.017.

Chieng SC, Kazacos M and Skyllas-Kazacos M (1992) Preparation and evaluation of composite membrane for vanadium redox battery applications. *Journal of Power Sources* 39: 11–19.

Cols JL (2025) Effective partnerships for utility scale VFB deployments. *IFBF Vienna Austria*.

Corcuera S and Skyllas-Kazacos M (2012) State-of-charge monitoring and electrolyte rebalancing methods. *European Chemical Bulletin* 1: 511–519.

DCCEEW (2025) Australian energy statistics — update report 2025. *Department of Climate Change, Energy, the Environment and Water (DCCEEW)*.

Edis T (2025) Bridging the gap to 82% renewable electricity generation by 2030. *Clean Energy Council*.

Hagedorn NH and Thaller LH (1982) *Design flexibility of redox flow systems*. US Department of Energy.

Hou Z et al. (2024) Towards a high efficiency and low-cost aqueous redox flow battery: A short review. *Journal of Power Sources* 601: 234242.

Huh J and Han S (2025) 1 MW/8 MWh vanadium flow battery project installed in Spain. *IFBF Vienna Austria*.

Huynh TTK et al. (2025) Construction of high-performance membranes for vanadium redox flow batteries: challenges, development, and perspectives. *Nano-Micro Lett*. 17, 260. DOI: 10.1007/s40820-025-01736-x.

Kausar N, Mousa A and Skyllas-Kazacos M (2016) The effect of additives on the high-temperature stability of the Vanadium Redox Flow Battery positive electrolytes. *ChemElectroChem* 3: 276–282.

Kazacos N and Kazacos G (2013) *Conducting plastic bipolar electrode assembly*. Patent Application.

Kazacos M, Cheng M and Skyllas-Kazacos M (1990) Vanadium redox cell electrolyte optimisation studies. *Journal of Applied Electrochemistry* 20: 463–467.

Kazacos M and Skyllas-Kazacos M (1994) *High energy density electrolyte for vanadium redox battery*. Patent Application PCT/AU94/00711.

Li L et al. (2011) A stable vanadium redox-flow battery with high energy density for large-scale energy storage. *Adv. Energy Mater.* 1: 394–400. DOI: 10.1002/aenm.201100008.

Li Y, Skyllas-Kazacos M and Bao J (2016) A dynamic plug flow reactor model for a vanadium redox flow battery cell. *Journal of Power Sources* 311: 57–67.

Li Y et al. (2024) Long term performance evaluation of a commercial vanadium flow battery system. *Journal of Energy Storage* 98: 111790. DOI: 10.1016/j.est.2024.111790.

Li Y, Kienbaum D and Lüth T (2023) Comparison study of different commercial vanadium flow battery stacks. *IFBF Prague Czech Repub*.

Li Z et al. (2024) Prospective life-cycle assessment of vanadium electrolyte recycling. *Energy Technology* 12: 2300750.

Mardilovich P and Harrer M (2023) Industrial applications of flow batteries — CellCube: history, progress, and outlook of commercial VRFB deployment. In: Roth C, Noack J and Skyllas-Kazacos M (eds) *Flow Batteries: From Fundamentals to Applications*, vol. 3, pp. 1079–1098.

Mardilovich I and Harrer B (2023) *Commercial deployment of vanadium flow batteries*. IFBF Proceedings.

Menictas C and Skyllas-Kazacos M (2011) Vanadium–oxygen fuel cell development. *Journal of Power Sources* 196: 10737–10747.

Mohammadi T and Kazacos MS (1996) Modification of anion-exchange membranes for vanadium redox flow battery applications. *Journal of Power Sources* 63: 179–186.

Mohammadi T and Skyllas-Kazacos M (1995) Preparation of sulfonated composite membrane for vanadium redox flow battery applications. *Journal of Membrane Science* 107: 35–45.

Mohammadi T, Chieng SC and Kazacos MS (1997) Water transport study across commercial ion-exchange membranes. *Journal of Membrane Science* 133: 151–159.

Mohammadi T and Kazacos MS (1997) Evaluation of the chemical stability of some membranes in vanadium solution. *Journal of Applied Electrochemistry* 27: 153–160.

Mousa A and Skyllas-Kazacos M (2015) Effect of additives on the low-temperature stability of vanadium redox flow battery negative half-cell electrolyte. *ChemElectroChem* 2: 1742–1751.

NASA (1980) Redox flow cell development and demonstration project. *Journal of Power Sources* 5: 382–384. DOI: 10.1016/0378-7753(80)80060-7.

Noack J, Stott C, Scarinci S and Luo B (2025) Flow batteries in Australia with a focus on Queensland. *IFBF Vienna Austria*.

Oei D-G (1985) Permeation of vanadium cations through anionic and cationic membranes. *Journal of Applied Electrochemistry* 15: 231–235. DOI: 10.1007/BF00620938.

Poli N, Bonaldo C, Moretto M and Guarnieri M (2024) Techno-economic assessment of future vanadium flow batteries based on real device/market parameters. *Applied Energy* 362: 122954. DOI: 10.1016/j.apenergy.2024.122954.

Rahman F and Skyllas-Kazacos M (2017) Evaluation of additive formulations to inhibit precipitation of positive electrolyte in vanadium battery. *Journal of Power Sources* 340: 139–149.

Reber D, Jarvis SR and Marshak MP (2023) Beyond energy density: flow battery design driven by safety and location. *Energy Advances* 2: 1357–1365. DOI: 10.1039/d3ya00208j.

Risbud A et al. (2019) Vanadium–oxygen fuel cell performance. *Electrochimica Acta* 312: 1–10.

Roe S, Menictas C and Skyllas-Kazacos M (2015) A high-energy-density vanadium redox flow battery with 3 M vanadium electrolyte. *Journal of The Electrochemical Society* 163: A5023.

Rychcik M and Skyllas-Kazacos M (1987) Evaluation of electrode materials for vanadium redox cell. *Journal of Power Sources* 19: 45–54.

Seiler Y, Schäffer M, Tübke J and Noack J (2025) Past, present and future of the vanadium redox flow battery (VRFB) system at Fraunhofer ICT embedded into EU-Horizon project SMHYLES. *IFBF Vienna Austria*.

Seipp T (2025) What's the perfect stack size? *IFBF Vienna Austria*.

Shu B et al. (2023a) Thermal modelling and simulation studies of containerised vanadium flow battery systems. *Batteries* 9: 196.

Shu B et al. (2023b) Hybrid cooling-based thermal management of containerised vanadium flow battery systems in photovoltaic applications. *Processes* 11: 1431.

Shu B et al. (2025) *Electrical safety analysis of VFB configurations*. IFBF Proceedings.

Skyllas-Kazacos M (1986) *Development of redox cell for remote area energy storage*. NERDDC Final Report 788.

Skyllas-Kazacos M (1989) *Development of vanadium redox cell for remote area energy storage*. NERDDC Final Report 1009.

Skyllas-Kazacos M (1991) *Design and development of vanadium prototype battery*. NSW Dept. of Minerals and Energy Final Report.

Skyllas-Kazacos M (1998) *Stabilised electrolyte solutions, methods of preparation thereof and redox cell*. Australian Patent 696452.

Skyllas-Kazacos M and Kazacos M (2002) *High-energy-density vanadium electrolyte solutions*. US Patent 6,468,688.

Skyllas-Kazacos M (2004) *Vanadium redox battery electrolyte*. US Patent 2004/0241552.

Skyllas-Kazacos M (2023) *History of UNSW vanadium battery licensing*. UNSW Internal Report.

Skyllas-Kazacos M et al. (1989) *State of charge of redox cell*. Patent Application PCT/AU89/00252.

Skyllas-Kazacos M et al. (2016) Vanadium electrolyte studies for the vanadium redox battery — a review. *ChemSusChem* 9: 1521–1543.

Skyllas-Kazacos M, McDermott R and Kazacos M (1988) *Vanadium salt dissolution process*. Patent Application PCT/AU88/00471.

Skyllas-Kazacos M, Rychcik M and Robins RG (1986) *All-vanadium redox battery*. US Patent 4,786,567.

Skyllas-Kazacos M and Kazacos M (2006) *High energy density vanadium electrolyte solutions*. US Patent 6,468,688.

Skyllas-Kazacos M and Kazacos M (2012) *Bromine complexing agents for halide flow batteries*. Patent Application.

Skyllas-Kazacos M, Mousa A and Kazacos M (2003) *Vanadium halide flow battery*. Patent Application.

Skyllas-Kazacos M, Rychcik M and Robins RG (1986) *All-vanadium redox battery*. US Patent 4,786,567.

Sukkar T and Skyllas-Kazacos M (2003a) Water transfer behaviour across cation-exchange membranes. *Journal of Membrane Science* 222: 235–247.

Sukkar T and Skyllas-Kazacos M (2003b) Modification of membranes using polyelectrolytes. *Journal of Membrane Science* 222: 249–264.

Sukkar T and Skyllas-Kazacos M (2004) Membrane stability studies for vanadium

redox cell applications. *Journal of Applied Electrochemistry* 34: 137–145.

Sum E and Skyllas-Kazacos M (1985) A study of the V(II)/V(III) redox couple for redox flow cell applications. *Journal of Power Sources* 15:. 179–190. DOI: 10.1016/0378-7753(85)80071-9.

Sum E, Rychcik M and Skyllas-Kazacos M (1985) Investigation of the V(V)/V(IV) system for use in the positive half-cell of a redox battery. *Journal of Power Sources* 16(2): 85–95. DOI: 10.1016/0378-7753(85)80082-3.

Sun B and Skyllas-Kazakos M (1991) Chemical modification and electrochemical behaviour of graphite fibre in acidic vanadium solution. *Electrochimica Acta* 36: 513–517. DOI: 10.1016/0013-4686(91)85135-T.

Sun B and Skyllas-Kazacos M (1992a) Modification of graphite electrode materials for vanadium redox flow battery application — I. Thermal treatment, *Electrochimica Acta* 37: 1253–1260. DOI: 10.1016/0013-4686(92)85064-R.

Sun, B. and Skyllas-Kazacos, M. (1992b) Chemical modification of graphite electrode materials for vanadium redox flow battery application — part II. Acid treatments. *Electrochimica Acta* 37: 2459–2465. DOI: 10.1016/0013-4686(92)87084-D.

Tang A et al. (2013) Investigation of the effect of shunt current on battery efficiency and stack temperature in vanadium redox flow battery. *Journal of Power Sources* 242: 349–356.

Tang A, Bao J and Skyllas-Kazacos M (2011) Dynamic modelling of the effects of ion diffusion and side reactions on the capacity loss for vanadium redox flow battery. *Journal of Power Sources* 196(24): 10737–10747.

Tang A, Bao J and Skyllas-Kazacos M (2012) Thermal modelling of battery configuration and self-discharge reactions in vanadium redox flow battery. *Journal of Power Sources* 216: 489–501.

Tang A, Bao J and Skyllas-Kazacos M (2014) Studies on pressure losses and flow rate optimization in vanadium redox flow battery. *Journal of Power Sources* 248: 154–162. DOI: 10.1016/j.jpowsour.2013.09.071.

Thaller LH (1975) *Electrically rechargeable redox flow cells*. US Patent 3,996,064.

Whitehead AH, Harrer M, Sun J and Pokorny P (2025) The economics of state of health management for the VFB. *IFBF Vienna Austria*.

Wyss P (2025) Flexbase: Developing a multi 100 MW FB in Switzerland. *IFBF Vienna Austria*.

Yan Y, Skyllas-Kazacos M and Bao J (2017) Effects of battery design, environmental temperature and electrolyte flowrate on thermal behaviour of a vanadium redox flow battery in different applications. *Journal of Energy Storage* 11: 104–118.

Zhang X et al. (2022) A novel process to prepare high-purity vanadyl sulfate electrolyte from leach liquor of sodium-roasted vanadium slag. *Hydrometallurgy* 208: 105805. DOI: 10.1016/j.hydromet.2021.105805.

Zhong S, Kazacos M, Skyllas-Kazacos M and Haddadi-Asl V (1997) *Flexible conducting plastic electrode and process for its preparation*. US Patent 5,665,212.

Zhong S, Padeste C, Kazacos M and Skyllas-Kazacos M (1993) Comparison of the physical, chemical and electrochemical properties of rayon- and polyacrylonitrile-based graphite felt electrodes. *Journal of Power Sources* 45: 29–41. DOI: 10.1016/0378-7753(93)80006-B.

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, pp. 25–45. ISSN 0035-9173/26/01025-21

European scientists in colonial Australia

The late Dr Peter J Tyler

Erstwhile Historian, Royal Society of NSW
Merewether Scholar, Mitchell Library, 2008

Presented at the Mitchell Library, Sydney, 24 November 2011
The 2011 Ben Haneman Memorial Lecture[1]

Ben Haneman

First, let me say how privileged I am to have been invited to present the fifth biennial Ben Haneman Memorial Lecture, following in the footsteps of some distinguished international scholars. I thank the Australian and New Zealand Society of the History of Medicine and the Trustees of the State Library of New South Wales for giving me this opportunity.

Next, I must say how apposite that this lecture should be given in the Friends Room of the Mitchell Library, surrounded by this unrivalled collection of the works of Miguel Cervantes donated to the Library by Ben Haneman. This includes 1,100 editions of *Don Quixote* in fifty different languages, the oldest published in 1620. Ben was a passionate bibliophile, like David Scott Mitchell, the benefactor whose bequest paid for this building. If a day passed in which he had not purchased another book, new or old, Ben felt empty and deprived. His literary interests were much wider than Spanish classics, as he showed in a lecture to the medical staff at St. George Hospital the year before he died: "One of my publicly declared vices is book-collecting. I think I have collected nearly a hundred different hospital histories" (Haneman 2000).

Unlike David Scott Mitchell, Ben Haneman did not spend his later years at home with his books. At the time of his death, he was still practising as a consultant physician, and he was active in numerous organisations that fostered his interests — medical history, science, libraries, Spanish culture, and the Jewish Board of Deputies. He was President of the Australian Society of the History of Medicine when he died. Most of all, he was devoted to this Library, where he served as a volunteer, and became the first Honorary Life Member of the Library Society. It was in this building that he collapsed suddenly almost exactly ten years ago, succumbing the following day without regaining consciousness. He would not have wished for a better end to a full life.

Another of Ben's passions was support for the Australian indigenous community, and he spent ten years of his spare time in the Mitchell Library preparing a complete index of now-defunct Aboriginal newspapers and journals. In his memory, therefore, I acknowledge with respect the traditional owners of the land on which we are meeting — the Gadigal people of the Eora nation.

For those of you who did not know Ben, I should say a few words about his life. Born in Italy of Jewish parents in 1923, Benedetto Haneman migrated to Australia

1 This address is available at the State Library of NSW. (Search for "2011 Ben Haneman Memorial Lecture".) I have sought a copyright holder unsuccessfully. Please contact me if you have information — Editor.

with his family as a young boy. As the only foreigner at Carlton Public School in the 1920s, he was known as "the Dago kid" and had stones thrown at him by boys who later became his patients. Leaving Canterbury Boys' High School at the early age of 15, he began studying medicine at the University of Sydney, graduating in 1944 just after his 21st birthday. After residencies in Perth, Ben began practice as a general practitioner at Carlton in 1947, as well as becoming a Clinical Medical Assistant at St. George Hospital. He went on to become a Fellow of the Royal Australasian College of Physicians, as a specialist in gastroenterology. He also served on the Library Committee of that College. Throughout his career, Ben nurtured young medical students, first as Warden of Clinical Studies at St. George Hospital, and later as Warden of Warrane College at the University of NSW. An inveterate traveller, he became an honorary professor at the University of Navarre, and in 1984 King Juan Carlos of Spain decorated him with the Cross of a Knight of the Order of Civil Merit. He was created a Member of the Order of Australia in 1988.

As one of his obituary writers said, Ben Haneman was a man of "extraordinary energy, enthusiasm, erudition, generosity, compassion and good humour" (Wilson 2002).

Lastly, I should emphasise how necessary it is for us to commemorate such a person, and to remind ourselves at regular intervals of his contribution to our society. In these days of instant gratification and a loss of corporate memory within so many of our institutions, it is vitally important that the Mitchell Library can still provide us with centuries of knowledge and understanding that otherwise would be lost.

Colonial arts and sciences

A few historians have drawn attention to the relationship of the arts and sciences to the colonial enterprise in white settler societies such as Australia, and have concluded that they formed part of the larger thrust of imperialism that so engaged western powers from the sixteenth century to the end of the nineteenth century. These writers suggest that as provincial outposts, the colonies simply provided information and resources to the home country in furtherance of imperial ambitions. This interpretation does not allow much scope for local originality or creativity, and neglects the independence imposed by the "tyranny of distance." Fortunately, some historians have taken a more nuanced view, for example the late Ann Moyal (1993), who pointed out deficiencies in both the isolationism and intellectual imperialism theories, and Roy MacLeod (2004) in his magisterial biography of Sydney University Professor Archibald Liversidge.

Recent trends in imperial history have tended to emphasise a networked concept of empire, with complex relationships between colonies, and with the metropole, rather than seeing them as subservient dominions (see Lester 2006). I do not wish to engage in that discourse this evening, but instead draw attention to some colonial experiences that require an alternative explanation. In doing so, I hope to redress some of the Anglocentric bias we find in much Australian history, as well as the relative lack of interest that has been shown in the development of cultural and intellectual life of this nation (see Fairburn 2004).

In particular, I want to connect with some areas of Ben Haneman's experience by describing — in roughly chronological

order — the sometimes overlooked role of scientists from Continental Europe in nineteenth century New South Wales and to depict in some detail a notable example from the field of medical science. The people I will speak about were born and educated in either France, Germany, Austria, Sweden, Poland, or Russia. Some were only transient visitors, while others remained here. It is by no means an exhaustive list; numerous other scientists from Continental Europe also visited. You must bear in mind that the nineteenth century was a period of intense nationalism as well as imperialism in Europe; small princely states were being amalgamated into larger entities, and national boundaries were shifting as a result of recurrent conflict. So the country of origin of these people will not always coincide with modern understandings of their birthplace.

French scientists in Australia

We know that European navigators had been exploring this region of the world for at least three centuries before a settlement was established in 1788. Partly they were driven by innate curiosity, partly to search for the fabled *Terra Australis Incognita* which they believed was necessary to balance the land masses of the northern hemisphere, but mainly by avarice. Uncharted territories promised unimagined riches. Many Dutch, British, and possibly Spanish, seamen had bumped into these shores but had not been impressed by what they had seen. They drew some charts but found little water or provisions for their crew, and few even bothered to take specimens of unfamiliar objects back home. Very few were like Dirk Hartog, who left behind a commemorative plate as tangible evidence of his visit to the western coast of New Holland in 1616.[2]

This all changed when Lieutenant James Cook sailed up the east coast of New Holland in His Majesty's Bark *Endeavour* in 1770. His expedition was overtly scientific, even if covertly political. We must remember that this was at the peak of the European Enlightenment, the Age of Reason when centuries of religious dogma were being challenged by rational explanations of the phenomena of nature. Funded by the Royal Society of London to observe the transit of Venus in Tahiti, Cook's vessel carried a small scientific party led by Sir Joseph Banks, the wealthy amateur botanist. Amongst his party was the Swedish naturalist **Daniel Solander** (1733–1782), a student of Carl Linnaeus, who when he stepped ashore in Botany Bay became the first university-educated scientist to set foot on Australian soil.

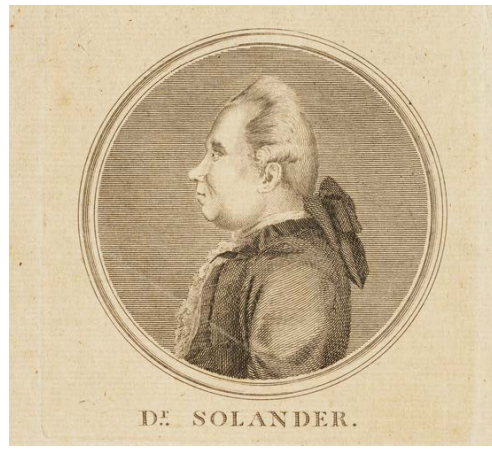

Botanist Daniel Solander, who accompanied Sir Joseph Banks and James Cook on their visit to Botany Bay in 1770. [SLNSW: PXA1028, item 9 (Album ID 824241)]

2 See https://www.migrationheritage.nsw.gov.au/exhibition/objectsthroughtime/hartog/index.html — [Ed.]

And here we immediately note an important distinction. In class-ridden England of that period, science, or natural philosophy as they called it, was the province of leisured gentlemen, who tended towards the observational rather than experimental modes of study. The professional scientist in this case was a foreigner. To the English, the fact that a person was paid to do research rather than for love of the subject implied that their conclusions would be compromised. To be a professional by definition excluded one from respectable society for much of the nineteenth century. Even the term "scientist" did not enter the English language until 1833.[3] It was very similar to the class-based distinction later made on the cricket field, when "players" and "gentlemen" in the same team occupied separate dressing rooms (Knight 1991: 42).

The situation was different in Continental Europe, where French governments both before and after the Revolution patronised science, and in the German universities, which provided systematic training in various scientific disciplines and a culture of experimental research. Because most educated Australians in the early nineteenth century brought with them British traditions, they came with a "pattern of scientific activity in which social distinctions were extremely important and social distances preserved." As Atkinson remarked, "The First Fleet sailed from England two years before the outbreak of revolution in Europe, and for many years in this remote corner of the globe the eighteenth century stood still." (Atkinson 1997: 7)

So it is hardly surprising that the next scientific expedition to touch these shores — and like Cook with a hidden imperialist agenda — was that of the French Count **Jean-François de Galaup de La Pérouse** (1741–1788), who arrived in Botany Bay just as Arthur Phillip and the First Fleet were preparing to move to Sydney Cove. You will recall that Arthur Phillip also had a European heritage: his father was German. La Pérouse and his two ships, *Buoussole* and *Astrolabe*, were well-equipped with scientific instruments and qualified personnel to use them, including a respected astronomer, **Joseph Lepaute Dagelet** (1751–1788), who immediately set up an observatory at Frenchman's Beach.

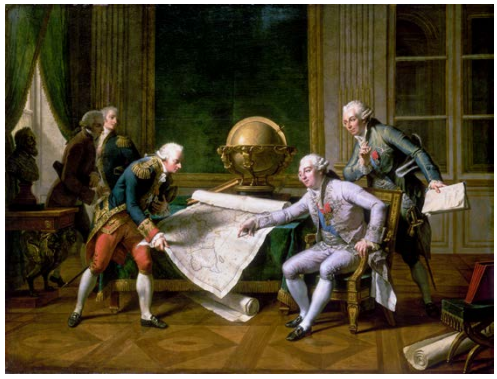

Louis XVI giving final instructions to the Comte de la Perouse, 1785, before his expedition to New Holland. Painted by Edouard Nuel. [SLNSW: ML39 (a128478)]

Dagelet was an Associate Member of the French Academy of Sciences, and was able to provide marine Lieutenant William Dawes with advice on setting up an observatory in Sydney Cove and establishing its precise location. Meanwhile, the other scientists in the French party were busy collecting and illustrating the unique flora and fauna they found here, as well as making

3 William Whewhell, of Trinity College, Cambridge, first used the term at a meeting of the British Association in 1833. (Hooker 2004: 10)

ethnographic studies of the indigenous population. Amongst them was Father **Claude Francois-Joseph Receveur** (1757–1788), who had been selected by the Académie des Sciences to serve as a naturalist and chaplain aboard the *Astrolabe*. A scientist, soldier and Franciscan priest, he was known as an "*insectologiste*," but he also made geological collections during the expedition (Duyker 2011a: 13). Unfortunately he was recovering from injuries received in Samoa and died soon after they arrived in Botany Bay, thus becoming the first scientist buried in Australia. His tomb can still be seen in the suburb of La Perouse, where it is currently undergoing restoration with French funding (Duyker 2011b: 24).

While the convict settlement at Sydney was still struggling for survival, another Frenchman, the **Chevalier Bruny D'Entrecasteaux** (1737–1793) led another well-equipped scientific party which explored and charted coastal van Diemen's Land in 1791–93, as well as making a number of landfalls on the mainland to collect specimens. In the words of Ann Moyal, when his botanist **Jacques-Julien Houtou de la Billardière** (1755–1834) published his findings in Paris, his work immediately "linked Australia directly to the European scientific enterprise, and provided a major early tool of Australian botany" (Moyal 1993: 24).

The pattern of under-resourced British and state-of-the art French expeditions was exemplified in 1802 when Matthew Flinders and **Nicolas Thomas Baudin** (1754–1803) unexpectedly encountered each other while charting the continental coastline from opposite directions. Nevertheless, despite the fact that their two nations were at war with each other, the commanders exchanged information civilly, true to the spirit of the notional "Republic of Letters" to they both subscribed. Baudin sailed on to Sydney for his crew to recuperate from illness, but although they were entertained cordially, Governor Philip Gidley King was suspicious of their motives in charting Van Diemen's Land, and immediately despatched Lieutenant Bowen to establish a settlement on the Derwent estuary (Blainey 2001: 78).

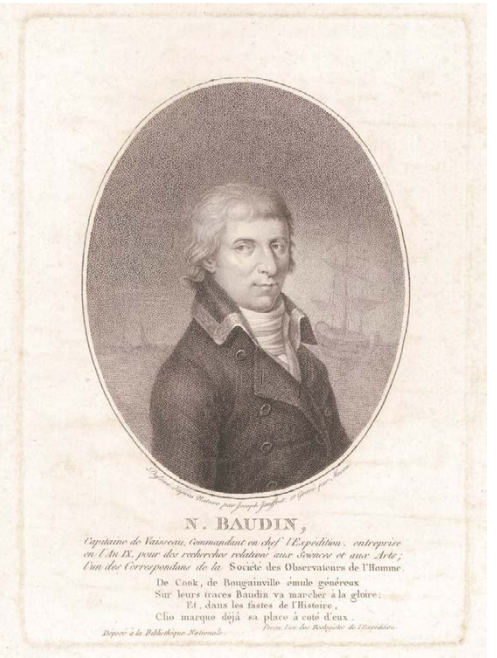

Nicolas Thomas Baudin, engraving by AJ Mecou, c.1804. [SLNSW: P2/467 (a128093)]

Baudin was conducting a purely scientific expedition in his appropriately-named ships *Le Geographe* and *Le Naturaliste*, with 23 distinguished scientists aboard, including the anthropologist **François Auguste Péron** (1775–1810) (Jones 1998: 2572). Some beautiful charts drawn by Baudin's naval colleague, the cartographer/naturalist **Louis-Claude Desaulses de Freycinet** (1779–1842) were recently displayed in the Dalgety Walkway

in Mitchell Library. This was a period of intense rivalry between England and France in the aftermath of the French Revolution and the expansionist dreams of Napoleon, for which Flinders afterwards suffered years of incarceration on Mauritius. Flinders had but one scientist aboard the decrepit *HMS Investigator*, the English botanist Robert Brown, plus the outstanding Austrian botanical artist **Ferdinand Lukas Bauer** (1760–1826) (Gilbert 1966). The difference in equipment is partly explained by the fact that the French received government sponsorship, while the British relied largely on private patronage (Liebersohn 2006: 97).

De Freycinet returned to Australia and the western Pacific as leader of his own expedition in 1818–19, and gathered a great deal of additional scientific material, much of which was lost when his ship *L'Uranie* was wrecked on the Falkland Islands on the return voyage. Although this procession of French scientists over a period of thirty years made many notable discoveries which were widely publicised throughout Europe, they did not attempt to usurp the British settlers who arrived here before them. Their visits are remembered now only by the taxonomy of some unique floral and faunal species, and the names of a few geographical features.

Macquarie and the Russians

After the turbulent early years of the Colony, Governor Lachlan Macquarie began to restore some semblance of civilisation to the nascent settlement at New South Wales. Macquarie was a visionary builder, but not an intellectual. He didn't share the current enthusiasm for collecting and classifying natural history specimens, but he did set up the beginnings of a botanical gardens in the Governor's Domain, and appointed Charles Fraser (or Frazer) as its superintendent around 1816. But Fraser was a gardener, not a botanist. At least Macquarie — or perhaps his wife Elizabeth — had the good sense to take back home to Scotland the two beautiful chests of natural history specimens that had been given to him, and which are now treasured artefacts of the Mitchell Library.

Residents of Sydney Town were surprised one morning in March 1820 to discover a strange Russian ship in their harbour, the *Blagonmerenny*. It turned out to be part of Admiral Bellingshausen's scientific exploratory expedition sponsored by Tsar Alexander I, and was soon joined by a naval sloop, the *Otkrytie* (Branagan 2010). These intruders fuelled Australian fears about invasions from the north that have persisted for more than two centuries. Several scientists were amongst the party, including astronomer **Pavel Tarkhanov** and surgeon naturalist **Fedor Ivanovich Stein** (d.1845). Governor Macquarie gave them a courteous if somewhat reserved welcome, allowing them to set up a temporary observatory on Kirribilli Point, and permitting them to make a geological excursion into the Blue Mountains, accompanied by botanist Allan Cunningham. Their findings were largely forgotten both here and in Russia, partly because they were not of great significance, or in some cases inaccurate. A later commentator remarked that Stein had "made much theory … out of little fact" (Aurouseau 1972). Even the names they gave to notable features in the landscape were later replaced by good English names by Sir Thomas Mitchell and others (Branagan 2010: 25).

The Philosophical Society of Australasia

Towards the end of Macquarie's term as Governor of New South Wales, a small group of educated officials and settlers established the Philosophical Society of Australasia. This became the first scientific body in the Colony, and some people regard it as a distant predecessor of the Royal Society of NSW, even though it was moribund for nearly thirty years after. When Major-General Sir Thomas Brisbane arrived in Sydney to become the new Governor in December 1821, his retinue included two astronomers to assist him with observations from his planned observatory at Parramatta. One of these was the German-born and educated **Christian Carl Ludwig Rümker** (1788–1862).

Governor Brisbane readily accepted an invitation to become the President of the Philosophical Society and actively participated in their gatherings, even hosting meetings at Government House.[4] His personal physician Dr Macleod and his private astronomer Carl Rümker also joined the Society, educated gentlemen both. Rümker, Brisbane and their technician colleague James Dunlop[5] began a rigorous program of celestial observations from the new observatory in the grounds of what we now know as Old Government House. They quickly made a number of significant discoveries, including the predicted return of Enke's comet. Sir Thomas Brisbane was awarded the Gold Medal of the Royal Society in London, and in turn he granted Carl Rümker 1,000 acres of land at Picton as a reward. To the intense annoyance of Brisbane, Rümker left his job at the observatory after twelve months to become a farmer, so the Governor tried to rescind the land grant (Bergman 1960).

When Brisbane was recalled to England in 1825, the government acquired his observatory. Rümker then returned to Parramatta, and discovered another new comet. Governor Ralph Darling appointed him to the position of Government Astronomer, the first person to hold this title, but Sir Thomas Brisbane still nursed a grudge, and used his influence to have Rümker dismissed from government service. Rümker returned to his native Hamburg and became director of the observatory there, while continuing to work on his Australian star catalogue. He received many international awards for his achievements, but is little recognised in this country (Bergman 1967).

But before this drama occurred, the members of the Philosophical Society subscribed to form a small collection of objects as the nucleus of a natural history museum which was housed in the office of the Colonial Secretary, who at this time was Frederick Goulburn, another member of the Philosophical Society. His successor as Colonial Secretary was Alexander Macleay, a noted amateur entomologist, who expanded the collection, much of which found its way into the Australian Museum that was established in 1836, although a lot also ended up in the Macleay family's private collection at Elizabeth Bay which eventually became the Macleay Museum at the University of Sydney.

The Poles

The Polish naturalist **Jan (John) Lhotsky** (1795–1866) received a grant from the

4 Minutes, Philosophical Society of Australasia, 21 November 1821.

5 See Rutledge (2009). [Ed.]

King of Bavaria to carry out botanical and zoological research in South America and Australia. He arrived in Sydney in 1832, but his application for appointment as zoologist at the museum was rejected. However, he did receive some government support for geological exploration in the Australian Alps and Snowy River region, discovering gold and other minerals. His specimens were exhibited and a report of the expedition was published in 1834–35. Nevertheless, he failed to obtain other employment despite representations by the Prussian ambassador in London, and was reduced to selling vegetables for a living. In disgust, he moved to Hobart in 1836 where again he found only intermittent scientific employment before returning to London after eighteen months (Whiteley 1967).

Another Polish scientist who faced difficulties in Australia was (Sir) **Paul Edmund de Strzelecki** (1797–1873). After a colourful early life and extensive travels, he arrived in Sydney in 1839 with letters of introduction to Governor Gipps. An inveterate explorer, he began a geological survey of the Colony, later extending this work to Victoria and Tasmania. He did not get on well with Gipps, and was deeply disappointed when Rev. WB Clarke was commissioned instead of him to survey for minerals in 1851. Nevertheless, he is still remembered in the Hunter Valley for his research of coal deposits and chemical analysis in that region. Strzelecki's *Physical Description of New South Wales and Van Diemen's Land* laid the basis of Australian palaeontology, for which he was awarded the founder's medal of the Royal Geographical Society in London, where he returned to live. He remained active in public life, and retained a special interest in fostering emigration to Australia. His contributions were recognised through a knighthood, Fellowship of the Royal Society of London, and an honorary Doctorate from Oxford University (Heney 1967).

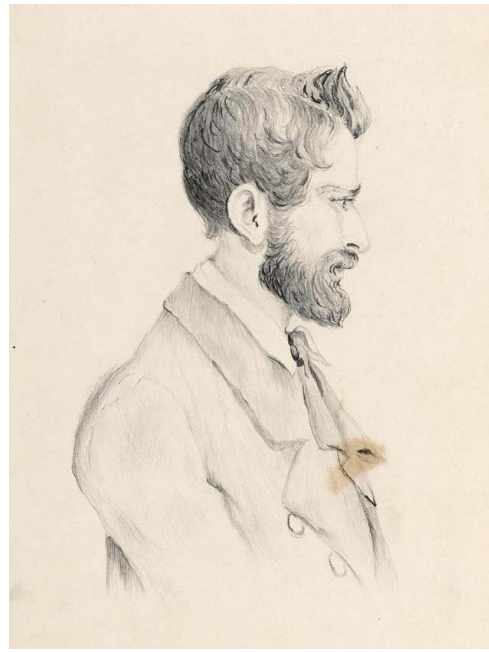

Explorer naturalist Ludwig Leichhardt. Portrait by Isobel Fox, 1846. [SLNSW: DG P1/4 (a128456)]

The Germans

Friedrich Wilhelm Ludwig Leichhardt (1813–1848), like La Pérouse, is perhaps best known to us as an explorer whose final expedition disappeared, or as the name of a suburb in Sydney. We might even think of Leichhardt as the inspiration for Patrick White's 1957 novel *Voss*, but Ludwig Leichhardt would have preferred to be remembered as a naturalist. After studying philosophy and the natural sciences at the Universities of Berlin and Göttingen, but without graduating, he came to Sydney in 1842, where he gave lectures on the geology and botany of the region while lobbying

Governor Gipps to be appointed curator of a national museum or of the botanic gardens. Although the position at the botanic gardens was vacant, the selection committee wanted an Englishman rather than someone who was likely to send interesting specimens to Paris or Berlin (Bailey 2011: 87). Having failed to secure a scientific appointment, Leichhardt began exploring more widely, his privately-financed party reaching the site of present-day Darwin after an arduous overland journey that took 14 months and opened up much unfamiliar territory (Erdos 1967).

Sadly, we do not know what became of his third expedition, and presumably he perished somewhere in the central desert at the age of only thirty-five before he was able to contribute more of his scientific knowledge to this country. Leichhardt was probably the first person to express the notion that type specimens and other reference or unique materials ought to remain in or at least return after description to Australia, rather than being deposited in overseas collections (Home 1988). Before his disappearance, Leichhardt's work was valued highly for the accuracy and restraint of his reports, gaining the Patron's Medal of the Royal Geographical Society in London, as well as recognition by the King of Prussia and the Geographical Society of Paris. Yet in Australia, twenty years after his death he was vilified as being "careless, slovenly, wholly unfitted for leadership, and very lax in his religious opinion" (Erdos 1967: 103). Later writers continued to build on this negative assessment, culminating with the ornithologist AH Chisholm's opinion that Leichhardt was an incompetent psychopath. It is notable that adverse comments reached their zenith during both the First and Second World Wars, when anti-German sentiment was at its peak.

Only in recent years has there been a more balanced appraisal of his work. His most recent biographer remarked that "If he had lived, there is no doubt this intellectual powerhouse would have made enormous contributions to our knowledge of the natural world" (Bailey 2011: 348).

Leichhardt's rejection for a museum position was an early example of the prevailing British, or at least Scottish, xenophobia that existed in New South Wales. Leichhardt himself alleged that Alexander and William Macleay opposed his appointment (Gilbert 2000: 27). Governors Macquarie and Brisbane were both Scots, as were the Macleays, while the first four superintendents of the Botanic Gardens — Charles Fraser, Richard and Allan Cunningham, Charles Moore — all boasted of Scottish heritage. Malcolm Prentis (2008) has written about the influence of the Scots in Australia, and this should be a topic for another lecture. During the nineteenth-century, Australians with a Scottish legacy excelled in agriculture, commerce, medicine, education and the sciences — partly because of their legendary Presbyterian thriftiness, but more significantly because of the superior educational opportunities in Scotland at that time, where merit counted for more than pedigree, as in England. I am told that Robert Louis Stevenson wrote that "the greatest good fortune that any man can have in life is to be born a Scot" (Haneman 2000).

Nevertheless, the most notable of all Colonial botanists was **Baron Ferdinand von Müller** (1825–1896), who was appointed Colonial Botanist in Victoria in 1853, and

later Director of the Melbourne Botanical Gardens.[6] After training as a pharmacist, he completed a PhD in botany from the University of Kiel, becoming the only academically-qualified botanist in Australia during the nineteenth century. Most other people in the colonies termed botanists were more accurately described as botanical collectors. Von Müller travelled extensively outside Victoria as well, and — serving Ben Haneman's interest in the history of medicine with a combination of his pharmaceutical and botanical skills — von Müller reported on the possible medicinal value of certain plants in the treatment of consumption (TB), rheumatism and scurvy. Like Ludwig Leichhardt almost two decades before him, von Müller argued strongly for the description and classification of natural history specimens to be carried out in Australia, and not in England or Europe (Morris 1974). Unlike most of their local contemporaries who deferred to British authority, these men were true nationalistic Australians.

From the 1850s until the end of the nineteenth century Germans were the largest group of non-British migrants to Australia (Doust 2004: 159). A German Association (*Deutscher Kranken Verein*) met weekly in Sydney from 1853, and was soon followed by a German Glee Club which organised musical events.[7] These are forerunners of the present day Concordia Club, which started in 1883, and still has a choir and a folk-dancing group.

Indeed, despite the misgivings of the British-born colonial establishment, another German did become curator of the Australian Museum. **Johann Ludwig (Louis) Gerard Krefft** (1830–1881) was born in Prussia. At the age of twenty he migrated to the United States, but two years later moved to Australia, attracted by the Victorian gold rush.

Gerard Krefft came to Sydney in 1860 where he was appointed assistant curator of the Australian Museum, becoming curator in 1864 despite some dissent amongst the trustees. Krefft systematically enlarged the Museum's collections, while his broad knowledge of zoology and geology and extensive research earned him international respect in the scientific community. He corresponded with Charles Darwin and other notable scientists overseas. Gerard Krefft was particularly interested in snakes, and his book *The Snakes of Australia*, published in 1869 became the standard reference work for many years. In addition to his books, he published some 200 scientific articles as well as many more popular pieces for the Sydney newspapers.[8] Because he was a capable watercolour artist, he was able to illustrate these articles with appropriate images, although some of his works were also illustrated by the sisters Harriett Scott and Helena Forde, who were cousins of David Scott Mitchell (Chanin 2011: 107).

Gerard Krefft was anxious to make the Museum a popular attraction for disseminating scientific knowledge. He wrote to a friend that a collection "of all kinds of bugs, beetles, butterflies and cockroaches without explanation is about as good a vehicle to education [as] ... the window of any large grocery establishment." He told

6 See Müller (1876, 1879, 1880, 1881, 1883, 1886, 1886,1888, 1889). [Ed.]

7 *Sands' Sydney Directory for 1870.*

8 See Krefft (1863, 1865, 1867, 1868, 1873). Krefft was one of the first Councillors of the Royal Society of NSW, having been a member of its forerunner, the Philosophical Society of NSW (Whitley 1958: 24) [Ed.]

Premier Henry Parkes that he wanted displays that were better than "all the rotten fishes crammed into bottles and covered by a brown fluid."

Zoologist Gerard Krefft, curator at the Australian Museum 1864–74. Photograph by H Newman, 1871. SLNSW: P1/912 (a4219012)]

Ever since his appointment, Krefft's relationship with the Museum trustees had been tense. Some of the trustees, such as (Sir) William Macleay, were intent on building up their private natural history collections; often they were competing with the Museum in acquiring exhibits. In June 1874 the Museum trustees held a special meeting at which twelve charges were levied against Krefft, ranging from drunkenness to disobedience. He was accused of wilfully smashing a fossil jawbone, while in a dramatic scene just after the meeting started, one of the trustees rushed in with two policemen brandishing a number of indecent photographs which they had found on the premises. Krefft appealed to Henry Parkes to intervene. Parkes replied: "I have great respect for your undoubted ability & am truly sorry that you should be involved in such a disagreeable difficulty. ... But you must learn to keep a cool temper & a respectful bearing even to gentlemen who may be opposed to you."

Nevertheless, the trustees found Krefft unfit to be curator and dismissed him. He maintained that they did not have the power to do this, because he had been appointed by the Governor and Executive Council. Krefft locked himself inside his residential apartments at the Museum. He was then forcibly ejected by the trustees, although some of his property, including his medals, manuscripts and personal papers, was retained by them. Reports vary as to whether a single bailiff or two prize-fighters were involved, but historians agree that somebody was employed to break down the door into Krefft's quarters, and carry him, still seated in his armchair, into College Street where he was unceremoniously dumped. His small son and a visitor were ejected at the same time. Mrs Krefft was absent, but returned to find her husband and son in the street and their apartment barricaded. Krefft still refused to concede that he had lost his job. He rented premises around the corner in Stanley Street, and had a rubber stamp made with the wording: "Gerard Krefft. Curator of the Australian Museum. 73 Stanley St., Sydney." This he applied to all his extensive correspondence subsequently.

Protracted litigation then commenced between Krefft and the Museum trustees. He won two cases and was awarded some financial compensation.

To remove any doubt about his position, Krefft's dismissal was confirmed by the Governor and Executive Council in July 1878. Parliament voted him £1,000 for arrears of salary, but the government refused to pay this sum unless Krefft renounced all claims against the trustees or the government. He asked the Supreme Court to compel the Treasurer to pay the amount voted by Parliament, but this appeal was rejected, which is not surprising, because the Colonial Treasurer, the Chief Justice, and the Attorney-General were all trustees of the Australian Museum, his (former) employers. Krefft complained to the Crown Solicitor: "I think that after twenty years of faithful service in various Museums under Her Majesty's Government I would have been entitled to as much consideration as the notorious bushranger ... Gardiner received from some of the Trustees ... " The affair became a *cause célèbre* in Parliament and the Press, which lampooned the Trustees as "the finest extant specimens of the *Dilettanti* genus — *humbug* species" (Chanin 2011: 111).

Because the Museum trustees still held his medals and testimonials, Krefft could not get other employment despite a strong recommendation from (Sir) Henry Parkes. His private collections were damaged and many of his research papers could not be published. This left him in a precarious financial position that ended in bankruptcy in 1880. The newly-formed Zoological Society of New South Wales offered him a cottage and a small salary of £100 a year to be its secretary and curator, but these arrangements were never finalised because the disheartened, depressed and impoverished Krefft died early the following year from "congestion of the lungs" — probably pneumonia — aged fifty-one.

Australia lost a distinguished scientist who could have contributed much more to our understanding of the natural environment. Although his reputation was tarnished by his detractors in Australia, his achievements were recognised overseas, with a mountain named in his honour, and the award of a Knighthood by King Victor Emmanuel of Italy. In retrospect, he was badly treated by the Sydney Establishment, yet his own personality was unhelpful in resolving the conflict. His uncompromising, belligerent attitudes — perhaps a characteristic of his Teutonic upbringing — caused him to be ostracised in a close-knit, conservative community.

The age of the collector

The nineteenth century was the great age of the collector. *Magpies, Squirrels & Thieves* is the apt title of Jacqueline Yallop's 2011 book about this phenomenon. Wealthy enthusiasts paid high prices for unusual plants, birds, animals or geological curiosities from places like New South Wales, as well as Greek antiquities, Chinese ceramics or Italian renaissance art. This market was exploited by officials and even convicts, who could sell unique specimens to visiting ships' captains who in turn sold them for a handsome profit when they returned to their home ports in Britain or Europe. The activities of casual collectors were supplemented by more systematic efforts by people engaged by European patrons. The Mitchell Librarian, Richard Neville, has

pointed out that "European interest in natural history specimens created the earliest export market of Australian commodities to Europe" (Neville 2010).

It seems a little ironic that these hunters and gatherers were intruders in a society where the original inhabitants had survived by hunting and gathering for at least 40,000 years!

Although women could not take part in the discussions of the learned societies, they were able to participate in scientific activity as collectors and illustrators. Painting of botanical or faunal specimens was considered an appropriately lady-like activity at which numerous colonial women excelled. **Amalie Dietrich** (1821–1891) from Saxony was from the same social class, but very different inclination. She developed a passionate interest in natural history from her chemist husband. After their marriage broke up, in 1863 Mme. Dietrich was employed by the Hamburg merchant JC Godeffroy to collect material in Australia for his private museum. After spending eight years here (mainly in Queensland), she returned to work in Godeffroy's museum, and later was given a position in the Hamburg Botanical Museum.

Amalie Dietrich wrote that "I speedily forget the discomforts of heat and mosquitoes in the unbounded feeling of joy that animates me when, at every step, I light upon treasures that no one has secured before me" (Gilbert 1972). She even caught and dissected crocodiles nearly seven metres in length — a sort of early, female "Crocodile Dundee"![9] Historian Ann Moyal remarked that "In a colonial environment, her very foreignness and outsider status provided her strength" (Moyal 1993: 95). She did not feel obliged, nor did society expect her, to behave like a sheltered British lady.

Now let us return to the French influence in Australian science, this time with a particular orientation to one of Ben Haneman's abiding interests — medical history.

Rabbits and Pasteur

In an attempt to provide gentlemanly sport for his guests and neighbours, pastoralist Thomas Austin imported twelve pairs of field rabbits from England in 1859 to stock his property, "Barwon Park," near Geelong. With this act, he unwittingly almost destroyed the profitability of the very industry that had brought him wealth. By the end of the nineteenth century the "furry hordes" had spread throughout the temperate parts of mainland Australia, consuming valuable pastures and causing massive land degradation. The European wild rabbits (*Oryctolagus cunciculus*) colonised Australia very quickly, travelling up to 100 km a year in open country. The rabbits adapted quickly to the hotter, drier climate because their burrowing habit helps to insulate them from temperature extremes, and most Australian soils are particularly suited to digging. Unlike Europe, there were no predators that could follow them underground.

In desperation, the NSW government in August 1887 offered a prize of £25,000 to anyone who could devise a "method or process not previously known in the Colony for the effectual extermination of rabbits ... [but] not be injurious ... to horses, cattle, sheep, camels, goats, swine or dogs" (Colman 1999: 65). No mention was made of possible harmful effects on humans or

9 *Crocodile Dundee* is an iconic 1986 Australian film comedy starring Paul Hogan as a knockabout crocodile hunter from northern Australia.

wildlife. A Royal Commission (known as the Intercolonial Rabbit Commission) was formed to examine all claims for the prize. All the colonies were represented, with the chairman being Dr HN (later Sir Normand) MacLaurin, who would later become Chancellor of the University of Sydney.

The lavish prize — representing something over ten million dollars in today's money — attracted over 1,500 responses from all over the world, ranging from the plausible to the downright silly. One hundred and fifteen (including forty-two from Australia) proposed a biological attack on the pests, using diseases such as smallpox, syphilis or hydrophobia.

Louis Pasteur (1822–1895) in Paris was one of those attracted by the lucrative reward. He argued that an epidemic should be propagated amongst the rabbit population through the inoculation of chicken cholera (*Pasteurella septica*), which he had used successfully at the Pommery champagne vineyards in France. Although the technique seemed to promise an elegant solution to the problem, there was strong opposition from some quarters. Public health and quarantine officials were dubious. What if pastoral livestock became infected, or the native Australian animals? Even more seriously, was there a risk that the disease could spread to the human population, with unpredictable consequences?

Nevertheless, Pasteur's international reputation as the father of modern bacteriology ensured that his proposal was treated seriously. Pasteur had already accepted Honorary Membership of the Royal Society of NSW in 1883 for his contributions to scientific knowledge. After an exchange of correspondence with the Minister for Mines, his proposal was accepted and Pasteur sent a team to Sydney led by his nephew and research assistant, Dr **Adrien Loir** (1862–1942). Thus French science returned to Australia after a lapse of sixty-five years. Even before the Pasteur Institute team arrived, the NSW Board of Health, which — like the Rabbit Commission led by the autocratic and sceptical Scot Dr MacLaurin — was warning against the introduction of such a "loathsome and pathologically undefined" disease (Rolls 1984: 198). At his instigation, the Board of Health initially rejected Pasteur's request to introduce chicken cholera to Australia.[10] However, after testing the virulence of the organism in its own laboratory, the Board agreed to allow field trials on condition that no objections were received.

Faced with such opposition, the Pasteur team had to find an acceptable location to conduct their experiments. Eventually, Rodd Island in Iron Cove was selected as a safe site where the virus could be confined without spreading to the mainland. Domestic accommodation, a laboratory and artificial burrows were hastily erected on the island. The Rabbit Commission appointed the Viennese-born Dr **Oscar (Otto) Katz**, a protégé of Sir William Macleay, as Chief Expert Officer to supervise the experiments. Katz is a shadowy figure about whom little is known. He arrived in Sydney in 1886, claiming to have earned his PhD as a student of the great German microbiologist Robert Koch. There is no evidence that this was true (Dando-Collins 2008: 129–132). His intervention was guaranteed to antagonise the French team, because there was intense professional rivalry and personal dislike

10 NSW Board of Health, Minutes of proceedings, 18 April 1888.

between Robert Koch and Louis Pasteur, as well as their ingrained nationalistic hostility.

The results of the first trials were not encouraging. Although the rabbits inoculated with the disease soon died, it was not contagious to the other rabbits in the warren. To make matters worse, native birds became infected with the virus. In its final report in December 1889, the Rabbit Commission concluded "that no scheme has been propounded for the destruction of rabbits which complies with the terms of the proclamation." So the Government kept their £25,000 — and Australia kept the rabbits. Acrimonious discussions took place before Dr Loir and his men returned to France, although he did come back to Rodd Island the following year to prepare vaccines against anthrax and pleuro-pneumonia in cattle (Rountree 1987: 13). With increasing public opposition to the alienation of the island for a laboratory, Rodd Island was proclaimed a public recreation reserve in 1894, becoming a popular picnic venue (Davies 1984: 56). The Pasteur Institute then moved its operations to Double Bay.

The French historian of science Anne Marie Moulin (1992: 308) has pointed out that Louis Pasteur was engaged in his own program of scientific imperialism, setting up branches of his Institute throughout the world, and not only in the French colonies in Africa and Asia. She suggests that his acolytes functioned like a monastic order, with their own "saint" as founder, "spreading a new Gospel, the so-called 'Pasteurian methods'" of bacteriology.

Despite the negative report on the efficacy of chicken cholera, the Queensland Government bacteriologist, CJ Pound, was still distributing the microbes ten years later to anybody who requested them (Rountree 1988). He argued that the Rabbit Commission decision was erroneous, with the implication that professional rivalries or prejudice had prevented the Pasteur team from completing its experiments. That was certainly their view and that of the French Consul in New South Wales. Although derided in Australia, Loir's reputation in Europe was consolidated following the publication in 1892 of the earliest scientific treatise on microbiology in Australia[11] (Rountree 1987: 13).

I am now going to trespass a little beyond my nineteenth-century brief. In 1905, the Polish-born microbiologist Dr **Jean Danysz** (1860–1928) from the Pasteur Institute in Paris wrote to the NSW Government advising that a new organism was available for rabbit destruction. Danysz had already conducted a successful campaign against cattle plague in South Africa. However, after its earlier experience, the NSW government was not keen to become involved, until a group of pastoralists formed the Rabbit Destruction Fund Committee to bring Danysz to Sydney in 1906. The government reward offer had lapsed, but the Committee soon raised £10,000 from subscriptions. Danysz kept the new pathogen secret, but it was probably a bacterium of the *Pasteurella* species, not a virus (Coman 1999: 77). At the request of the Premier, the Board of Health reluctantly allowed its microbiologist, Dr Frank Tidswell to oversee these experiments, which took place on Broughton Island, off Port Stephens. He reported that the disease was potent among rabbits held in close cap-

11 See Loir (1891, 1891). [Ed.]

tivity, but ineffective when the animals had a free range (Rountree 1988: 36).

Because of the secrecy surrounding the precise nature of the disease, there was a lot of resistance to the Danysz experiments with his "filthy … treacherous … hell-broth" as the Melbourne *Age* described it.[12] The trade union movement urged the government to prevent dissemination of the organism, "in the interests of health and human life" (Coman 1999: 79). Their motives were probably mixed — rabbit was cheap meat for the working classes ("underground mutton"), and they were reluctant to see the squatters become even richer at their expense. Coming after their previous experience in Sydney, the French scientists accused the Australians of sabotage (Rountree 1988: 41). Danysz returned to Paris in 1907, leaving his assistant **A. Latapie** to continue the experiments in Australia. Danysz was later responsible for significant discoveries on the role of non-specific antigens in the treatment of chronic and anaphylactric disease.[13] So we lost another original thinker.

Despite the opposition, the pastoralists demanded a larger-scale trial in typical rabbit-infested country on the mainland. After much discussion, the States' public health and veterinary officers concluded there was no danger, but still refused permission (Rolls 1984: 206). Over-ruling this recommendation, in 1908 the NSW Premier GC Wade informed the Board of Health that he proposed to allow the "Danysz microbe" to be used. By this time, the President of the Board of Health was the formidable Dr John Ashburton Thompson, the first trained epidemiologist to work in Australia. The Board tartly replied that it "infers from the communication that as a body they are not called upon to express an opinion or to take any responsibility."[14] The Premier was not prepared to proceed in the face of this threat from the medical hierarchy, and so the rabbits remained until the introduction of other viral agents such as myxomatosis and the calicivirus half-a-century later.

The Russian

Nicholai Nicholaievich Mikluho-Maklai (1846–1888) was a scientific polymath: medico, anthropologist, linguist and marine biologist. Educated in Russia and Germany, he is usually known simply as Nicholas Maclay in Australian records. During the early 1870s he lived in New Guinea villages studying the local inhabitants, and part of that coast is named after him. After further travels, he arrived in Sydney in 1878 where he soon struck up a friendship with W.J. Macleay, who shared his interest in sharks, and provided some laboratory space at his home, Elizabeth Bay House.[15] The Russian joined the Linnean Society, which published 34 of his research papers, while he lobbied the government on the need for a proper laboratory of marine sciences in Sydney. He raised money through donations and obtained a grant of £300 from the colonial government. The new building at Camp Cove, near Watson's Bay was finished in 1882. The *Sydney Mail* remarked approvingly:

12 *The Age*, 22 May 1907, p.6 https://trove.nla.gov.au/newspaper/article/198614691 [Ed.]

13 Archives of the Institut Pasteur, Paris, Ref.FR IP DAN. Website: www.pasteur.fr/infosci/archives/dan1.html.

14 NSW Board of Health, Minutes, 21 January 1908.

15 See Jones (1986). [Ed.]

> The position of the station at Watson's Bay, upon the shallow basin of Port Jackson, and close to the deep water of the Pacific, with large fresh-water swamps and lagoons in the immediate neighbourhood, and a vast tract of forest country to the north … offer … a combination of opportunities for study as can hardly be met with elsewhere.[16]

Maclay then went to Russia, Berlin, Paris and London, lecturing about his collections in an effort to raise additional funds. This effort was not successful, although the Russian Geographical Society awarded him a gold medal and the Czar presented him with a certificate of honour. To add to his disappointment, when he returned to Sydney in 1883 he found that many of his collections and records had been lost in the Garden Palace fire.

Nevertheless he began work at the new marine laboratory and prepared several papers for publication, still managing time to marry the widowed daughter of Sir John Robertson, five times Premier of New South Wales. They had two sons[17], but on a trip to Europe with the family in 1888 Nicholas Maclay died unexpectedly (De-M Maclay 1974). He was a distinguished scientist who did not live long enough to fulfil his ambition to document the marine life of Sydney Harbour. He was fortunate to have the friendship of WJ Macleay, with a coincidentally similar surname. Although the research station closed, its activities have recently been revived as the Sydney Institute of Marine Science at Chowder Bay on the opposite side of the Harbour.

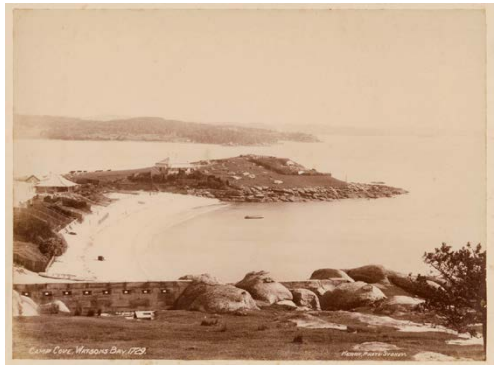

Marine Biological Station established at Camp Cove by Nicholas Maclay, 1881. Photograph by Kerry & Co., between 1885–95. [SLNSW: SPF/810 (a089810)]

A common thread?

So far we have looked at professional scientists from Sweden, France, Russia, Germany, Austria and Poland. Many of these European visitors were given a hard time in Australia, so it is not surprising that while some persevered, others gave up. Some had established international reputations when they arrived; usually they had specialised training in scientific disciplines which their local counterparts lacked. They tended to be experimenters rather than collectors.

So we are left with the question: is there a common thread linking all these experiences? They were working in a pioneering settler society where intellectual pursuits were not highly valued. They were coping rather well in a language that was not their own, and where few people could speak to them in their native tongue. Because they were usually better-educated than the people around them, they were regarded as outsiders, an elite whose difference made them stand out from the crowd. Prejudice

16 *Sydney Mail*, 14 May 1881, quoted in Sydney Harbour Federation Trust leaflet.

17 And a grandson, Paul Maclay (pron. Macl-eye), the ABC announcer and actor who died in 2002. [Ed.]

and jealousy surrounded them. Louis Pasteur accused the Australian authorities of "ignorance, blindness, prejudice, despicable tricks," and wrote to a friend that "my assistants will have seen the worst aspects of humanity and of men without honour."[18] It is notable than none of them worked in the academy. The University of Sydney opened in 1852 and science subjects were taught as part of the Arts degree from its inception. Academic staff were mostly Englishmen or Scots, who naturally brought their practices and "memorised culture"[19] of Britain with them. Who were the principal villains? The names of Sir William Macleay and Sir Normand MacLaurin recur in several of these episodes. In a biography of MacLaurin written by his granddaughter, he was described as "despotic … austere and formidable" (Mackerras 1991: 40–41). Both were Scots; both were Vice Chancellors of the University of Sydney. They were not unanimously against foreign scientists, but they certainly played favourites. And at least they didn't return "Home" to die like many of their compatriots from the respectable classes of Sydney.

Yet was it simply a matter of xenophobia, a distrust of the "other" — the foreigner? After all, those "transplanted Britons" were themselves strangers to the country. Or was it perhaps something in the personality of the Europeans, in their assertive way of doing business in a strange environment, that jarred with traditional British reserve? As one obituary notice remarked of Gerard Krefft: "If he had been as much as home with men as with animals, or could have charmed his trustees as cleverly as he did his snakes, his fate would have been a much fairer one."[20]

Were the men I have discussed simply paranoid? After all, we must not forget that a number of other European-born and educated scientific workers settled in Sydney, leading successful and rewarding lives. Amongst many others, one can cite the Italian Dr **Tomasso Fiaschi** (1853–1927), who pioneered new techniques of medicine and surgery; the German chemist Dr **Adolph (Carl) Leibius** (1833–1893),[21] who was senior assayer at the Royal Mint in Macquarie Street; and a French geographer living in Sydney, **Edmond Marin la Meslée** (1852–1893), who according to Russel Ward "was a propagandist for Australia in France and became more interested than many native or British-born colonists in the growth of Australian nationality" (Ward 1974).

Conclusion

Whatever their idiosyncrasies, most of the Europeans who remained here became members of the Royal Society of NSW, some even accepted imperial honours, but they were unlikely to have been attracted to the ideals of "Imperial Science under the Southern Cross," exemplified by such English scientific expatriates as Archibald Liversidge and Rev WB Clarke.[22] Was it just a career move for them, or were they, like

18 Letter to Mrs Priestley, 29 June 1888, held by Archives du Ministère des Relations Extérieures, Paris, quoted in Chaussivert & Blackman (1987), p.21.

19 The phrase is attributed to Donald Horne, cited in Macleod (2009: 11).

20 See Whitley (1958). [Ed.]

21 See Leibius (1864, 1869, 1872, 1872, 1884, 1891). Leibius was elected President of the Royal Society of NSW in 1890–91. [Ed.]

22 This is the title of a biography of Liversidge (MacLeod 2009).

Ben Haneman, true lovers of science, and of Australia?

Let us conclude with a few words by Ben. He was a proud Fabian socialist who believed that civilisation is gradually progressing, albeit with many setbacks along the way. He wrote:

> I believe that at present the Goths and Vandals are having a free kick at the fabric of our culture and civilization. It is my belief that it is the duty of every man and woman to be an activist ... to heal the world or make it a slightly better place.

But, like many of the other people we have looked at this evening, perhaps Ben Haneman was really just another Don Quixote, tilting at windmills.

References

Atkinson A (1997) *The Europeans in Australia*, Vol.1, Oxford University Press.

Aurousseau M (1972) The Russian visitors of 1820, *Newsletter of the Royal Australian Historical Society*. No.113, April, p.5.

Bailey J (2011) *Into the Unknown. The tormented life and expeditions of Ludwig Leichhardt*. Pan Macmillan.

Bergman C (1960) C.C.L. Rümker, Australia's first government astronomer. *Journal of the Royal Australian Historical Society* 46(5): 247–289, November.

Bergman GFJ (1967) Christian Carl Ludwig Rümker (1788–1862). *Australian Dictionary of Biography (ADB)*, Vol. 2.

Blainey G (2001) *The Tyranny of Distance*. Pan Macmillan, revised edition.

Branagan DF (2010) The Russian Expedition's Sydney visit in 1820 and some forgotten Blue Mountains names. *JProcRSNSW* 143: 1–28. https://doi.org/10.5962/p.361632.

Chanin E (2011) *Book Life. The Life and Times of David Scott Mitchell*. Australian Scholarly Publishing.

Chaussivert J and M Blackman (eds) (1987) *Louis Pasteur and the Pasteur Institute in Australia*. The French-Australian Research Centre, Sydney.

Coman B (1999) *Tooth and Nail: The Story of the Rabbit in Australia*. Text Publishing.

Dando-Collins S (2008) *Pasteur's Gambit. The Australasian Rabbit Plague and a Ten Million Dollar Prize*. Vintage Books.

Davies S (1984) *The Islands of Sydney Harbour*. Hale & Iremonger, Sydney.

de M-Maclay RW (1974) Nicholai Nicholaievich Mikluho-Maklai. *ADB*, Vol.5, pp.248–250.

Doust J (2004) Setting up boundaries in colonial Eastern Australia: Race and Empire. *Australian Historical Studies* 123, April.

Duyker E (2011a) *Père Receveur. Franciscan, Scientist and Voyager with Lapérouse*. Dharawal Publications.

Duyker E (2011b) A grave on the shore of Botany Bay. *Trust News Australia*, November.

Erdos R (1967) Friedrich Wilhelm Ludwig Leichhardt. *ADB* 2: 102.

Fairburn M (2004) The state of Victorian studies. in Australia and New Zealand. In: M Taylor & M Wolff (eds) *The Victorians Since 1901: Histories, Representations, Revisions*. Manchester UP.

Gilbert L (2000) *Mr McLeay's Elizabeth Bay Garden*. Mulini Press.

Gilbert LA (1966) Bauer, Ferdinand Lukas (1760–1826). *ADB* 1: 73.

Gilbert LA (1972) Amalie Dietrich. *ADB* 4: 73.

Haneman B (2000) Some history about St. George, lecture to medical staff, St. George Hospital, Kogarah, May. Transcript in the author's possession.

Heney H (1967) Strzelecki, Sir Paul Edmund de (1797–1873). *ADB* 2: 494.

Home RW (1988) History of Science. *The Australian Encyclopaedia* (5th edn), p.2577.

Hooker (2004) *Irresistible Forces*. Melbourne UP.

Jones BO (1988) Science. *The Australian Encyclopaedia* (5th edn), p.2572.

Jones S (1986) Gentlemen scientists: the Macleays of Elizabeth Bay House. *JProcRSNSW* 118: 211. https://doi.org/10.5962/p.361208.

Knight D (1991) Tyrannies of distance in British science. In: RW Home and

SG Kohlstedt (eds) *International Science and National Scientific Identity*. ASHPS, Kluwer.
Krefft JLG (1862) On the vertebrated animals of the lower Murray and Darling, their habits, economy, and geographical distribution. *TransPhilSNSW* 1: 1–33. https://doi.org/10.5962/p.345620.
Krefft JLG (1863) On snakes observed in the neighbourhood of Sydney. *TransPhilSNSW* 1: 34–60. https://doi.org/10.5962/p.345642.
Krefft JLG (1863) Description of a new fish from the Hawkesbury (*Therapon pittii*). *The Sydney Morning Herald*, Nov. 12, p. 8. https://trove.nla.gov.au/newspaper/article/13093253.
Krefft JLG (1865) On the manners and customs of the Aborigines of the lower Murray and Darling. *TransPhilSNSW* 1: 357–374. https://doi.org/10.5962/p.345642.
Krefft JLG (1866) On the dentation of *Thylacoleo carnifex*. *The Empire*, Aug. 2, p. 4. https://trove.nla.gov.au/newspaper/article/60593068.
Krefft JLG (1867) The vertebrata of Tasmania, recent and fossil. *TransPhilSNSW* 1: 30–41. https://doi.org/10.5962/p.346101.
Krefft JLG (1868) On the improvements effected in modern museums in Europe and Australia. *TransRSNSW* 2: 15–25. https://doi.org/10.5962/p.346108.
Krefft JLG (1869) *The Snakes of Australia*. Sydney: Government Printer.
Krefft JLG (1873) Australian natural history: mammals of Australia and their classification. Part I — Ornithodelphia and Didelphia. *TransRSNSW* 7: 135–147. https://doi.org/10.5962/p.346156.
Leibius CA (1864) On osmium and iridium, obtained from New South Wales gold. *TransPhilSNSW* 1: 210–215. https://doi.org/10.5962/p.345632.
Leibius CA (1869) On a new apparatus for reducing chloride of silver. *TransRSNSW* 3: 168–170. https://doi.org/10.5962/p.346172.
Leibius CA (1872) On an improved method of separating gold from argentic chloride, as obtained in gold-refining by chlorine gas. *TransRSNSW* 6: 67–70. https://doi.org/10.5962/p.346140.
Leibius CA (1872) Remarks on the fallacy of a certain method of assaying Antimony Ores given by some Manuals of Assaying. *TransRSNSW* 6: 71–72. https://doi.org/10.5962/p.346141.
Leibius CA (1872) Remarks on in ore, and what may appear like it. *TransRSNSW* 6: 73–74. https://doi.org/10.5962/p.346142.
Leibius CA (1884) Notes on gold. *JProcRSNSW* 18: 37–41. https://doi.org/10.5962/p.358966.
Leibius CA (1891) President's address *JProcRSNSW* 25: 1–45. https://doi.org/10.5962/p.359096.
Lester A (2006) Imperial circuits and networks: Geographies of the British Empire. *History Compass* 4/1.
Liebersohn H (2006) *The Travelers' World*. Harvard UP.
Loir A (1891) Notes on the large death rate among Australian sheep, in country infected with Cumberland Disease, or splenic fever. *JProcRSNSW* 25: 46–52. https://doi.org/10.5962/p.359097.
Loir A (1891) Notes on a spontaneous disease among Australian rabbits. *JProcRSNSW* 25: 89–93. https://doi.org/10.5962/p.359103.
Mackerras C (1991) *Divided Heart*. Sydney: Little Hills Press.
MacLeod R (2009) *Archibald Liversidge, FRS. Imperial Science under the Southern Cross*. Royal Society NSW/Sydney UP.
Morris D (1974) Mueller, Sir Ferdinand Jakob Heinrich von (1825–1896). *ADB* 5: 306.
Moulin AM (1992) Patriarchal science: The network of the overseas Pasteur institutes. In: P Petitjean et al (eds) *Science and Empires*. Kluwer Academic.
Moyal A (1993) *"A Bright and Savage Land." Scientists in colonial Australia*. Penguin.
Müller JW von Baron (1876) Description of fossil fruits in siliceaous deposit, Richmond River. *JProcRSNSW* 10: 239–239. https://doi.org/10.5962/p.358785.
Müller JW von Baron (1879) *Ottelia præterita*. *JProcRSNSW* 13: 95–96. https://doi.org/10.5962/p.358859.
Müller JW von Baron (1880) A catalogue of plants collected during Mr. Alexander Forrest's geographical exploration of north-west Australia in 1879. *JProcRSNSW* 14: 81–96. https://doi.org/10.5962/p.358886.

Müller JW von Baron (1881) Census of the genera of plants hitherto known as indigenous to Australia. *JProcRSNSW* 15: 185–300. https://doi.org/10.5962/p.358917.
Müller JW von Baron (1883) Additions to the census of the genera of plants hitherto known as indigenous to Australia. *JProcRSNSW* 17: 187–189. https://doi.org/10.5962/p.358957.
Müller JW von Baron (1886) Description of an unrecorded *Ardisia* of New Guinea. *JProcRSNSW* 20: 43. https://doi.org/10.5962/p.358995.
Müller JW von Baron (1886) Further additions to the census of the genera of plants hitherto known as indigenous to Australia. *JProcRSNSW* 20: 75–77. https://doi.org/10.5962/p.358999.
Müller JW von Baron (1888) Considerations of phytographic expressions and arrangements. *JProcRSNSW* 22: 187–204. https://doi.org/10.5962/p.359047.
Müller JW von Baron (1890) Record of hitherto undescribed plants from Arnheim's Land. *JProcRSNSW* 24: 73–79. https://doi.org/10.5962/p.359084.
Neville R (2010) At first glance. *SL Magazine*, State Library of NSW, Spring, p.14.
Prentis M (2008) *The Scots in Australia*. UNSW Press.
Rolls E (1984) *They All Ran Wild: The Animals and Plants that Plague Australia*. Angus & Robertson.
Rountree PM (1987) The Pasteur Institute and Australia. In: J Chaussivert and M Blackman (eds) Louis Pasteur and the Pasteur Institute in Australia. Kensington, NSW: French-Australian Research Centre, UNSW.
Rountree PM (1988) Jean Danysz — A forgotten visitor to Australia. *Australian Microbiologist* 9(1): 35.
Rutledge S (2009) For further consideration: Mr. James Dunlop, Esq. *JProcRSNSW* 142: 17–28. https://doi.org/10.5962/p.361621.
Ward R (1974) Marin la Meslée, Edmond Marie. *ADB* 5: 210.
White P (1957) *Voss*. Eyre & Spottiswood.
Whitley GP (1958) The life and work of Gerard Krefft. *Proc. of the Royal Zoological Society of NSW*: 21–34. https://www.biodiversitylibrary.org/item/122965#page/23/.
Whitley GP (1967) Lhotsky, John (1795–1866). *ADB* 2: 114.
Wilson G (2002) Obituary, Benedetto ("Ben") Haneman. *Medical J. of Australia* 176, 18 March: 292.
Yallop J (2011) *Magpies, Squirrels and Thieves: How the Victorians Collected the World*. Atlantic.

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, pp. 46–47. ISSN 0035-9173/26/01046-02

The Sjögren's disease biobank and registry in NSW

Adrian YS Lee

Department of Clinical Immunology & Immunopathology, Westmead Hospital & NSW Health Pathology, Westmead, NSW; Westmead Institute for Medical Research, The University of Sydney, Westmead, NSW

Email: adrian.lee1@sydney.edu.au

Abstract

Sjögren's disease (SjD) is a common systemic autoimmune disease that causes dryness symptoms, fatigue and pain. There are currently no targeted therapies available to treat this disorder. As a valuable resource to address this unmet need, the Western Sydney Sjögren's Syndrome Biobank (WS3B) was established in 2021 to promote the understanding, research and health outcomes of SjD in New South Wales and Australia. Five years on, the WS3B has contributed substantially to local, national and international research efforts.

Letter to Editor

Dear Editor,

Autoimmune diseases result when the body's immune system fails to properly distinguish self from non-self, creating a misguided attack on otherwise healthy tissues. Examples of these include type I diabetes mellitus, lupus, and Sjögren's disease (SjD). SjD is one of the most common systemic autoimmune diseases, classically resulting in profound dryness symptoms, pain, fatigue and organ damage (Lee et al. 2023b). The disease was named after a Swedish ophthalmologist, Henrik SC Sjögren, who was one of the first to describe the condition in the 1930s (Jonsson 2021). SjD patients are typically middle-aged females and experience a significant impact on their emotional, functional and social health (Lee et al. 2026; Lin et al. 2025). Immunologically, B cell dysregulation and autoantibodies are hallmark features and patients are at high risk of developing B cell cancers (Lee et al. 2023b). Unfortunately, there are currently no cures or targeted therapies for SjD and patients globally, although several are in trial (Pelkas et al. 2025).

To address these unmet needs in SjD care and research in Australia, I established the Western Sydney SjD Biobank (WS3B) at Westmead Hospital in 2021. This is the first of its kind in Sydney and one of the very few in Australia. The WS3B is a biobank and registry of carefully curated SjD patients and controls. I acknowledge my colleagues, the support of Westmead Hospital and Westmead Institute for Medical Research, and the generous financial support from the Leece Family and Western Sydney Local Health District for facilitating the development and growth of this valuable resource. I am also incredibly grateful to my own patients for

helping advance the knowledge and understanding of SjD in Australia and globally.

The key aims of the WS3B include:

- To improve the understanding and diagnosis of SjD in Sydney and, in particular, Western Sydney;
- To promote research, teaching and education in SjD in New South Wales and Australia; and
- To improve health outcomes for people living with SjD in New South Wales and across Australia, by advancing precision medicine approaches and developing better therapies.

Now at its 5-year anniversary, the WS3B has over 100 patients and has contributed to local, national and international studies. These have ranged from basic science to clinical and epidemiology-based projects. Some of the key contributions of the WS3B include improving the phenotype of SjD patients to facilitate predictive medicine (Lee et al. 2023a), and characterising the impact that one's heritage has on disease manifestations (Lin et al. 2025).

Plans are underway to form a coordinated national SjD research movement. It is hoped that the WS3B will be an indispensable resource for future and current clinicians and researchers in New South Wales and beyond.

Acknowledgements

I am grateful to the numerous colleagues who have supported the biobank and registry. I also am grateful to the patients and their families for their generous biospecimen donations and contributing to research. I am also grateful to the Royal Society of NSW for the honour of receiving a RSNSW Early Career Early Career Research and Service Citation 2025.

Contact details

Dr Adrian Lee
Department of Clinical Immunology
Westmead Hospital
Hawkesbury Road
Westmead NSW 2145
AUSTRALIA

Ph: +61 2 8890 5555

References

Jonsson R (2021) Henrik Sjögren (1899–1986): the syndrome and his legacy. *Annals of the Rheumatic Diseases* 80(9): 1108–1109.

Lee AYS, Putty T, Lin M-W, et al. (2023a) Isolated anti-Ro52 identifies a severe subset of Sjögren's syndrome patients. *Frontiers in Immunology* 14: 1115548.

Lee AYS, Wang JJ, Gordon TP, et al. (2023b) Phases and natural history of Sjögren's Disease: a new model for an old disease? *Arthritis Care & Research* 75(7): 1580–1587.

Lee AYS, Zembrzuska H, Franke KB, et al. (2026) Patients' perspectives of living with Sjögren disease: a systematic review of qualitative studies from the OMERACT Sjögren disease working group. *Seminars in Arthritis and Rheumatism* 77: 152929.

Lin MW, Volk S and Lee AYS (2025) Primary Sjögren disease patients of Asian and Middle Eastern ancestry have early-onset but similar disease activity. *Internal Medicine Journal* 55(5): 856–859.

Pelkas C, Franke KB, Vincent FB, et al. (2025) Novel therapies in Sjögren's disease: A systematic review of the literature. *Best Practice & Research Clinical Rheumatology* 39(4): 102084.

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, pp. 48–67. ISSN 0035-9173/26/01048-20

2025 Royal Society of New South Wales and Learned Academies Forum: "AI: The Hope and the Hype"

Thursday, 6 November 2025
Government House, Sydney

Her Excellency the Honourable Margaret Beazley AC KC, Governor of New South Wales

Emeritus Professor Christina Slade FRSA FRSN, President, Royal Society of NSW

The Honourable Victor Dominello, Director of the UNSW–UTS Trustworthy Digital Society Hub, Future Government Institute

Professor Sally Cripps, UTS, Forum Program Chair

Professor Pascal Van Hentenryck, A. Russell Chandler III Chair, Georgia Institute of Technology

Opening Addresses

Her Excellency the Honourable Margaret Beazley AC KC

Bujari gamarruwa
Diyn Babana Gamarada Gadigal Ngura

Good morning, distinguished guests, Fellows of the Royal Society of New South Wales, representatives of Australia's Learned Academies, leaders from government, industry and research, and all joining us in person and by livestream.

I pay my respects to the Traditional Owners of the lands on which we meet — the Gadigal of the Eora Nation — and to Elders past, present, and emerging. I extend that respect to all Aboriginal and Torres Strait Islander peoples with us and acknowledge the particular relevance of the topic of today's discussions to First Nations peoples, their culture and intellectual property.[1]

It is my great pleasure to welcome you to the 2025 Royal Society of New South Wales and the Learned Academies Forum and its topical subject: "AI: The Hope and the Hype."

May I begin by acknowledging:

- Emeritus Professor Christina Slade, President of the Royal Society of New South Wales and Chair of this Forum.
- The Honourable Victor Dominello, Director of the UNSW–UTS Trustworthy Digital Society Hub, and former NSW Minister for Customer Service and Digital Government.
- Professor Pascal Van Hentenryck, A. Russell Chandler III Chair and Professor at the Georgia Institute of Technology; and
- Professor Hugh Durrant-Whyte, Chief Scientist and Engineer of New South Wales.

Each of these distinguished leaders — representing the fields of philosophy, civic leadership, science and engineering — presents a vital dimension to today's discussion. I also acknowledge the generous support of the many bodies that have contributed to this Forum.

1 https://www.abc.net.au/news/2025-08-23/calls-to-protect-indigenous-intellectual-property-from-ai-cultur/105680182

How does one define AI?

As AI evolves at an unprecedented pace, the definition itself continues to evolve. Indeed, in recent years, the Organisation for Economic Co-operation and Development (OECD) has struggled to reach consensus on a definition, acknowledging that "There is no clear red line but a continuum of features characterising what we think of artificial intelligence and where the 'magic' happens."[2]

Its current (2023) definition, a definition which the Australian Government's Digital Transformation Agency also employs,[3] is that it is "a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. (It notes that) different AI systems vary in their levels of autonomy and adaptiveness after deployment."[4]

AI's applicability and transformative potential in a seemingly endless range of environments — from governance to science and engineering, medical research, law, education, health and the public and community sectors — is matched by extraordinary commercial investment. McKinsey forecasts an astonishing $7 trillion US in global AI investment over the next 5 years.[5] It has fuelled expectations of governments, with the Australian Government's Productivity Commission forecasting growth of $116 billion in gains over the next decade and a 4.3 percentage-point boost in labour productivity.[6] It has also fuelled an industry of new AI scholarship with an incredible 403,897 books, journals and papers published worldwide in 2023.[7]

A quick old-school Google search of the range of recent media articles about AI — noting that AI relies on patterns, and thus historical rather than current data,[8] "threw up" some arresting headlines:

- Australia's culture of AI-powered innovation[9]
- How AI falsely named an innocent journalist as a notorious child murderer[10]
- The Sydney hospital using AI to improve diagnoses under pressure[11]

2 https://oecd.ai/en/wonk/definition

3 https://architecture.digital.gov.au/domain/ai

4 https://oecd.ai/en/wonk/definition

5 Forum summary: https://www.royalsoc.org.au/events/rsnsw-and-learned-academies-forum-2025/

6 ibid

7 https://ourworldindata.org/grapher/annual-scholarly-publications-on-artificial-intelligence?tab=discrete-bar&time=latest

8 "Artificial intelligence is often hailed as the technology of the future. Yet it primarily relies on historical data, reproducing old patterns instead of fostering progress" — Alexander von Humboldt: Institut für Internet und Gesellschaft: https://www.hiig.de/en/why-ai-is-currently-mainly-predicting-the-past/

9 https://www.afr.com/technology/australia-s-culture-of-ai-powered-innovation-20251015-p5n20l

10 https://www.smh.com.au/national/how-google-ai-falsely-named-an-innocent-journalist-as-a-notorious-child-murderer-20251024-p5n52d.html

11 https://www.smh.com.au/national/nsw/the-sydney-hospital-using-ai-to-improve-diagnoses-under-pressure-20251010-p5n1jp.html

- Big Tech deploys Orwellian doublespeak to mask its democratic corrosion[12]
- AI can help the environment, even though it uses tremendous energy[13]
- Tech billionaires seem to be doom prepping. Should we all be worried?[14]

And this last:

- "Over 50 Percent of the Internet Is Now AI Slop, New Data Finds" ... with the intriguing subline: "Humans aren't finished ... (pause) yet."[15]

All of which led another editorial to lead with a lament:

- Can we just have one day when no one mentions AI?[16]

It then went on to quote a recent Gallup poll in the US which showed that 49% of people think AI is just another tech advance that will improve our lives and 49% think it will harm humans and society.[17]

Australians are less inclined to sit on the digital fence: a snap Roy Morgan SMS survey of 18000 Australians in early October revealed that 65% of us believe that "artificial intelligence (AI) creates more problems than it solves" compared to 35% who believe it "solves more problems than it creates."[18]

On any measure, Artificial Intelligence can be a polarising topic. It has entered our public discourse not only as a technological innovation, but as a social force — one that promises to reshape the way we learn, work, govern, create and even imagine what it means to be human. Indeed, we should remember Alan Turing who, way back in 1950, pressed the question: "Can machines think?"[19]

On one hand, we hear of AI's capacity to help solve some of our greatest challenges — climate change, diseases and medical conditions, resource scarcity. On the other, we are reminded that technology has often outpaced regulation, and that many promises may remain — at least in the near term — unfulfilled.

This duality — the hope and the hype — underscores the tension that societies around the world now grapple with how to harness AI's enormous potential while guarding against its risks. The question arises: who is leading whom in intelligence? Who imagines the future?

12 https://www.theguardian.com/commentisfree/2025/oct/15/big-tech-deploys-orwellian-doublespeak-ai-deployment

13 https://abcnews.go.com/Technology/wireStory/ai-environment-tremendous-energy-5-ways-126751869

14 https://www.bbc.com/news/articles/cly178345240

15 https://futurism.com/artificial-intelligence/over-50-percent-internet-ai-slop

16 https://www.afr.com/technology/can-we-just-have-one-day-when-no-one-mentions-ai-20250811-p5mlvs

17 ibid; Drilling down into that Gallup poll result: https://news.gallup.com/poll/692435/major-threat-next-tech-thing.aspx

18 https://www.roymorgan.com/findings/10048-ai-press-release-october-2025

19 CSIRO: *Artificial Intelligence for Science, Adoption Trends and future development pathways*: November 2022: https://www.csiro.au/en/research/technology-space/ai/Artificial-Intelligence-for-Science-report
"One of the pivotal papers in AI was written by Alan Turing and published in October 1950. This paper opens with the words 'I propose to consider the question, can machines think?'. It lays out the future challenges for AI to solve. Turing's question is still being asked but remains unanswered." — Turing A (1950) Computing machinery and intelligence in Mind — A quarterly review of psychology and philosophy 59(236): 433–460.

Even more importantly: who has the power and where is the equity? noting that: "AI algorithms need big datasets to learn from, but several cultures and groups of the human population are absent or misrepresented in existing datasets."

As a paper published by the now insecurely-funded US National Library of Medicine has pointed out: "AI is thus prone to reinforcing bias which can lead to poor or fatal outcomes"[20] — or entrenching disadvantage.[21] All of this has led the European Union to legislate the EU AI Act — the first comprehensive regulation on AI by a major regulator anywhere.[22]

Over the course of today's sessions — AI and the Law; in Communities; in Health; AI in Practice and AI Research & Future Directions — we will probe precisely these tensions. Let us hold in mind a few guiding principles:

- Use evidence over assertion. Let claims about AI's power be backed by data, by peer review, by real-world trials and critical evaluation.
- Prioritise ethics, justice, and human rights. Technological possibility must be tethered to human values: fairness, transparency, accountability, inclusivity.
- Factor in governance and oversight. Institutions — governmental, academic, civil society — must evolve to keep pace, balancing innovation with precaution.
- Reflect on interdisciplinarity and humility. No discipline has all the answers alone. We need collaboration across law, technology, the social sciences, medicine, and beyond.
- The need for continuous assessment and adaptation of our legal, regulatory, educational frameworks.

AI is a shared responsibility and Forums such as this are essential. They bring together the broad spectrum of expertise needed to ensure that our collective intelligence keeps pace with our artificial one. As more than one academic has pointed out: "We are still in the early stages of this history, and much of what will become possible is yet to come."[23]

The Royal Society of New South Wales — the oldest learned society in the Southern Hemisphere — has always stood for rigorous, open, cross-disciplinary dialogue. That tradition is alive today in this gathering of scientists, policymakers, ethicists, and industry leaders.

It is, therefore, with great pleasure that I declare open the 2025 Forum of the Royal Society of New South Wales and the Learned Academies' "AI: The Hope and the Hype." Thank you.

20 Natalia Norori et al: 'Addressing bias in big data and AI for health care: A call for open science: National Library of Medicine (USA)'; https://pmc.ncbi.nlm.nih.gov/articles/PMC8515002/

21 "Artificial intelligence is often hailed as the technology of the future. Yet it primarily relies on historical data, reproducing old patterns instead of fostering progress": Alexander von Humboldt: Institut für Internet und Gesellschaft: https://www.hiig.de/en/why-ai-is-currently-mainly-predicting-the-past/

22 (2023) https://artificialintelligenceact.eu/

23 https://ourworldindata.org/brief-history-of-ai. Also, the view of the University of Sydney's Dr Rob Nicholls: "*Dr Nicholls still believes the industry is 20 or 30 years off truly mimicking humans, and that AI as it currently stands is immense data systems that are very good at predicting trends and patterns.*" https://www.abc.net.au/news/2025-09-22/is-ai-the-fourth-industrial-revolution/105790912

Christina Slade, President, Royal Society of NSW

Thank you, Your Excellency. A characteristically broad ranging and wise talk. I should say that I have now heard Her Excellency talk on many, many different issues, and she brings scholarship and wiseness and wisdom to us. And I think those principles are ones that are really worth our while thinking about more seriously.

It's my great honour to be the president of the Royal Society of New South Wales. I acknowledge the traditional owners of the land on which we meet: the Gadigal people of the Eora nation. I welcome you all to the Royal Society of New South Wales and Learned Academies' Forum, which is yet again being held most generously at the residence of our patron, Her Excellency. I thank the office of the Chief Scientist of New South Wales for supporting the Forum this year, not only financially, but actually personally.

So it's a great joy for me. This is the eleventh in the series which draws together the Royal Society, the Learned Academies, and members of civil society to debate issues of pressing importance. I should tell you that the last seven of these have been recorded and are available on our YouTube channel.[24]

The forums are designed to cross disciplinary boundaries to allow a broadly based, informed and civil discussion on a theme. Our fellows and members come from across the academy, but from beyond the academy as well: government, business, professional startups. Last year, our topic was "Threats to Democracy," which was held in the very week of Trump's election to the presidency. And no doubt, many of us have noted an election in the United States this week or yesterday.

This year we're equally topical: AI, the hope and the hype. It has been convened by the redoubtable Professor Sally Cripps, who not only brings her own enormous expertise, but her extended network of colleagues who are going to be speaking with us today. In fact, when we had this idea of artificial intelligence, we put together a committee to organise the event. And it was about two months in and a great deal of debate. And Sally said, "This is my area of expertise. Just leave it to me."

Now, I suspect she has regretted that very much, but it has been amazing. So I want to welcome you, the speakers, the Fellows and Members of the Royal Society of New South Wales and of the Learned Academies, Fellows of the Royal Society of Arts of London, and also some newcomers. We've worked very hard to have a new group here. I do hope that those of you who are new to the Society will consider becoming a Fellow or a Member.

I'm now turning to Sally Cripps, who's the convenor of today's forum. You'll see in your booklet the list of enormously long and impressive list of Sally's achievements. I wanted to say that one of the things that I think is most exciting about Sally's work is that she's working in a very highly technical area, enormously specialised, and yet she has managed to work across a whole range of issues looking for global impact. Her concerns and her interests are about global impact. So I'm delighted to introduce her as the chair of the forum.

24 See https://www.youtube.com/watch?v=HutO4CpRKaY?t=26 [Ed.]

Sally Cripps, Forum Program Chair

Thank you for those very kind words. It's really a pleasure to be here with you. As we've heard today from the Governor and from the President, we really are at a crossroads of grand hopes and rampant hype. We heard the Governor (who has also stolen some of my statistics, I was going to say) talk about the extraordinary investments that we have going on in AI. McKinsey, as was mentioned, forecasting a seven trillion in global investment. And in the last two months, NVIDIA has announced a hundred billion dollar investment in OpenAI. And also Google announced a $15 billion investment in India in setting up an AI data centre. There is a huge amount of optimism. Australia has moved in line with this optimism. We have gone from being quite cautious about AI, to looking at it as a game changer for productivity.

And the UK has also shifted in this narrative, but I would like to offer a really cautionary note about this productivity gains: a recent article looked at 59 countries, and where those gains in productivity go.[25] And the more developed the country, the more likely those gains go to profits and not to wages. What that means is that income inequity — or just in general wealth inequity — is exacerbated or can be exacerbated by new technology such as AI. This is a real challenge for politicians to recognise and stop talking about it as if everybody is naturally going to share in it. Everybody could naturally share in it, but only with deliberate action on our part will that happen. Economists have sounded a further warning: this may be less transformation than hype, with one arguing that's only about 1% of increase to global GP over the next decade.

Markets are wobbling and what we seek to do today is to explore this tension between the boundless excitement of hype and hope, which is, to my way of thinking, grounded optimism shaped by governance, ethics, and fairness. So as you see, I'm in front of you, but the AI panic is here. I don't think it's a panic, and I'm hoping today we'll draw that out by having a look at various aspects of AI. The first one is the Law, the second is Communities, the third is Medicine, then in Practice, and finally wrapping it up will be what the Future of Research is in AI. So these are the five things that we're going to be looking at today. And Christina said, I said, "Leave it to me." And actually what I did was ask some really remarkable people and left it to them.

The people you're going to hear talking today — the people who are on the panels — have all come together to do this. And I'm so forever grateful for all they're doing. I actually just picked up the phone and said, "Hey, what do you think?" And they ran with it. So I'm really grateful.

I did want to say just a couple of minutes about AI, because there's so much confusion. Today, we are going to be hearing about all sorts of AI. I'm not going to define AI, but I will give a taxonomy. So: first, Generative AI, I think of as that's what you'll all be familiar

25 Xiao A et al. (2024) Bridging the digital divide: the impact of technological innovation on income inequality and human interactions. *Humanit Soc Sci Commun* 11: 809. https://doi.org/10.1057/s41599-024-03307-8

with ChatGPT.[26] Fantastic at prediction so much so that it can write essays.

Generative AI — Prediction & Automation:

- Produces text, images, summaries;
- Useful for drafting, triage, translation;
- Risk: errors, hallucinations, misinformation;
- Policy relevance: needs transparency, human oversight.[27]

And then, second, Probabilistic or Causal or inferential AI, which is really about understanding the world.[28]

Probabilistic & Causal AI — Understanding Systems:

- Models uncertainty;
- Causal AI identifies true drivers of outcomes;
- Supports targeted, equitable education, health, justice policy;
- Policy relevance: strengthens evidence-based decisions.

And the third type I would call: Adaptive AI, which is that which engages in systems and adaptively learns and changes as it goes along.[29]

Adaptive AI — Continuous Learning:

- Continuously updates as new data arrives;
- Supports adaptive trials and real-time policy learning;
- Helps governments refine interventions safely;
- Policy relevance: enables responsive, iterative governance.

So they're just my framing. There are many others, but that's what we're going to hear about today.

I'm going to introduce our introductory speaker, Victor Dominello. Victor is a globally recognised leader in public transformation and digitalisation. He's listed as in the top 25 innovative government innovative leaders. Most of you here will remember Victor as NSW Minister for Customer Service. And most of you, I presume — at least the statistics tell me — 85% of us have done it. We have all got our driver's licence on our phones. This was all led by the wonderful Victor Dominello. So he's a superstar. And so without any further ado, I'll ask Victor to give the Introductory Talk.

26 Examples of Generative AI apps include chatbots such as ChatGPT, Claude, CoPilot, DeepSeek, Google Gemini and Grok; text-to-image models such as Stable Diffusion, Flux, Midjourney, and DALL-E; and text-to-video models such as Veo, LTX and Sora. [Ed.]

27 How dangerous is Gen AI? See Chayka K (2026) A.I. has a message problem of its own making. *New Yorker*, April 15, https://www.newyorker.com/culture/infinite-scroll/ai-has-a-message-problem-of-its-own-making [Ed.]

28 GenAI is deterministic but Causal AI is probabilistic. GenAI looks probabilistic because it samples from probability distributions. But those distributions are hard-coded correlations in past data. Given the same prompt and weights, GenAI's behaviour is fully deterministic. Causal AI, on the other hand, starts with the admission that the world itself is uncertain. In a nutshell, GenAI predicts words; Causal AI predicts worlds. [Ed.]

29 Adaptive AI is a form of AI that learns, adapts, and improves as it encounters changes, both in data and the environment. It lets organisations become more flexible, resilient, and effective when faced with complex challenges in daily operation. [Ed.]

The Honourable Victor Dominello, Director, UNSW–UTS Trustworthy Digital Society Hub; Future Government Institute

Thank you very much, Sally. It's actually 87% with a 93.6% satisfaction score. Lost a couple of points between friends. Thank you, Excellency for that amazing introduction. Today, I'm glad you put the "and" — not the "or" — between the hope and the hype, but everyone talks about AI. I've been fortunate enough since leaving politics and government two and a half years ago to be all around the world, from the White House to Mongolia — even in places like Mongolia, they're talking about some of the things that we've done here in New South Wales, particularly through the Service NSW app, because what we did here in New South Wales through that service transformation was really Government 3.0. Now, I'll just walk you back. Everyone talks in every boardroom I go to, in every government department I talk to, they're all talking about AI. Everyone's talking about AI and the evolution of AI.

Where is it going? You got Grok 2, Chat 3, Claude 4, Gemini 5. And as Sally's identified, that's just one stream of AI. Everyone's talking about the various stages, as you do. We talk about the evolution of connectivity. Last decade was 2G and then 3G, now we're in 4 and 5G. Next decade it'll be 6G. There are stages, there are gates. No one is talking about the evolution of government because if you don't talk about the evolution of government, then government stays this rock type figure that becomes reactive rather than proactive. And in the age of AI, that's moving so fast, that is like oil and water. It's not going to mix. Governments have to evolve. They have to become more agile and have to move a bit faster, just a bit faster would help. If I could walk you through the stages of government evolution and then maybe address where I think the hope and the hype is, and then sit down.

The stages of government: think about if you were to distil government, ultimately what's its interface with people, with the citizens, the customers it serves. Now, I say customers because it also provides services to refugees who are not citizens, to tourists who are not citizens, to businesses who are not citizens, to visa holders. So when it's dealing with or providing services to organisations, people, communities, how does it interface? It doesn't get a Ouija board out and say, "Hmm, I wonder what our client's wanting to do today." It normally interfaces through a form, right? If you want to receive a grant, if you want to receive a grant, you're going to have to fill out a form. Even if you have to pay money to tax, you've got to fill out a form. If you want a driver's licence, fill out a form. That's the interface. Up until 1990, we're Government 0.0. That's the paper form. That was the interface for many, many years before that.

In 1990, a big meteor hit the planet. It was called the Internet and it hit every government around the world. We went from paper form, Government 0.0, to web form, Government 1.0. So we're still communicating via forms, but now we also had a web form that we could do 24/7 when we're at the beach, but still in very siloed structures and still pretty much one-way traffic — the government basically requesting information. There wasn't much engagement. 15 years later, about 2005, another big meteor hit the planet. It was called the smartphone, social media. We went to the platform era, Government 2.0, and that means we've got

mobile apps and we're starting to have a bit more engagement with the customer. And now we're starting to hear a bit more noise around social media.

Well, they like these services, or they don't like those services, but still very siloed. Health does its piece, education does its piece, transport, police — they all do their thing. 15 years later, another big meteor hit the planet. It was called COVID, and then you bring in AI. COVID forced governments all around the world to go to "uniform." Think about it: in New South Wales alone, if we said to you, "We'll get you through COVID, but you're going to have to download a transport app for the border pass, you're going to have to download a police app for QR check-ins, you're going to have to download a health app for the vaccine certificate, you're going to have to download a finance app for your Dine and Discover vouchers, and then you're going to have to call each department if there's a problem on different numbers."

There is no way we would have got through it. "Uniform" doesn't mean it's one, but it means it's a unified experience. One major dominant app — the Service NSW app, one major website … we started looking through that lens. That's Government 3.0, and that's what service is. It's the beginning of what that could look like. In another 15 years' time, 2035, another big meteor was going to hit the planet, and it's called AGI — however that's framed. It will be 6G. That's definitely going to come next decade. We are going to have a lot more robotics, humanoid robotics. It's going to be an augmented world. As Her Excellency said now, I can't even tell now who's real and who's not. The level of augmentation in the next decade is going to be almost impossible to discern, but it's going to be a fundamentally different world.

We might even have the beginnings of quantum in some enterprise level. That is, that's going to be intelligent government. There's no such thing as intelligent government now. We might have smart departments, but there's no such thing as intelligent government because the left hand doesn't know what the right's doing. What we are finding is that with the private sector, they are moving rapidly into that G3, G4 environment. Governments have gone from G3, which was forced upon them by COVID, and most governments have now gone back to G2, which is siloed, because the Empire Strikes Back.

Ultimately, the source of power in government is around authority — the premier or the prime minister — and money, the treasurer. And things revolve around those two suns. There's no real third sun that represents the human, to make sure that we get that horizontal seamless experience rather than the silos. What the challenge for us is, particularly as we are talking about the hope and the hype of AI, is how do we get governments to evolve? Because I can guarantee you this: government is the elephant in the room. In every discussion you have, government is the elephant in that room. It could either be a good elephant and remove roadblocks, get rid of regulation and red tape, and let us move faster. Or it could be a slow elephant as most governments are, and just sit there and just slow things down. Or it could be a rampaging bad elephant such

as the Robodebt scheme[30] and like that it tramples all over us.

So how do we get government to move from bad to slow to good so that we can capture the hope of AI? I'll end on this note. There are so many things that we can unpack, and I'm sure a lot of the great speakers after me will go through that. But this is the reality. This is the hype. I had my brain MRI'd recently, my beloved mum's got dementia. As a result, I've now taken on a role of trying to promote brain health. So I thought I'd go and get an MRI. Spoke to the neurosurgeon, and he's using AI now to slice my brain into a thousand pieces and find dots that probably wouldn't have been seen before. Good bill of health so far, so that's great. But it's just extraordinary listening to these doctors and researchers about how powerful this will be.

This will literally be a tool to profoundly reduce suffering. We will have modern tech miracles. The blind will see. The deaf will hear. The lame will walk with brain-computer interfaces and everything else that's coming on board. I went to a neurotech conference the other day and they were talking about neural dust. Forget brain computer interfaces, neural dust. The world is going to change so fast. But this is the hype. What happens if all that power, all that control becomes centralised, then all of a sudden we're in a bad, bad place? Her Excellency was referencing a book about autocracies. We could have corporate autocracies, because autocracies are just about power. And power in the 21st century, as it has always been really, is about data who controls your data. So that's another lens that we should probably think about in the Hope and Hype era.

There is so much hope and it will be realised hope. But equally, there's a lot of hype in terms of: it may not go where we'd like it to be, but hype in a negative sense. If we don't get it right, it will be very, very bad.

Thank you. Your Excellency: it's been so wonderful to see you again. And thank you for the invitation and to another superstar. Sally, you're amazing. Thank you very much. Cheers.

Sally Cripps, Forum Program Chair

Thanks, Victor. We need to work that superstar bit out so that we can sort of tag. I was going to call another superstar, but before I call him a superstar, I just need to tell you a bit about him. Our next speaker is Professor Pascal Van Hentenryck. Pascal is the real deal when it comes to AI. He runs the NSF-funded AI Research Institute for Advances in Optimization. And they're looking at such a range of problems, from energy, the environment, and health, to supply chains. He really is across it all. And I know he's going to give us a lovely talk about what he believes AI is and what he has done in that area. I don't want to take up any more time and welcome another superstar.

30 See https://en.wikipedia.org/wiki/Robodebt_scheme [Ed.]

Keynote

AI: The Hype and the Hope

Pascal Van Hentenryck

Director, NSF AI Institute for Advances in Optimization
Georgia Institute of Technology

Thank you, Sally. And good morning, everyone. I still feel like a little kid from Belgium. That's where I come from. It's going to be very difficult to follow Her Excellency and Victor and Sally, but I'll try to do it as best I can. I'm the director of one of the national AI institutes in the United States.[31] We do use Inspire research. And a lot of the things that I will tell you in the second part of the presentation is going to be about what we do in the institute for actually trying to make a difference using AI. First, I'd like to thank you for the invitation. It's very humbling and I'm deeply honoured to be here today and trying to talk to you.

When I talk about AI, I like to give some context. In 1966, Joseph Weizenbaum, a professor at MIT, wrote the first chatbot, Eliza,[32] a chatbot mimicking a therapist.[33] It was so compelling that people using the chatbot actually revealed deep secrets about their private lives. He wrote an entire book that I was forced to read when I was an undergraduate in part of a philosophy class.[34] But it actually got me interested in AI. But this is telling you that this is not a new area. We actually have been looking at AI for more than 60 years. I was part of the first AI summer when expert systems were really the hot topic, and that was a big bubble that burst. And then we entered a very, very long winter. But typically in the winter, you do a lot of good works.

So that's where we invented most of the technologies which are the foundation of AI these days. And that's where people at AT&T Labs — Bell Labs at the time — invented convolutional neural networks, which are the core of what is basically used in computer vision.[35] Statistical natural-language processing — you can think of it as the core of ChatGPT these days — was invented at the time.[36] This was a very big move from grammar-based natural-language processing to statistical natural-language processing. The core of robotics was invented at that

31 The AI Institute for Advances in Optimization (AI4OPT) https://www.ai4opt.org/about/about-ai4opt [Ed.]

32 Named after Eliza Doolittle in Shaw's "Pygmalion." [Ed.]

33 Weizenbaum J (1966) ELIZA — a computer program for the study of natural language communication between man and machine. *Com. of the ACM* 9: 36–45 https://doi.org/10.1145/365153.365168.

34 Weizenbaum J (1976) *Computer Power and Human Reason: From Judgment To Calculation*. San Francisco: WH Freeman.

35 See Hubel DH and Wiesel TN (1968) Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology* 195(1): 215–243. DOI: 10.1113/jphysiol.1968.sp008455 [Ed.]

36 See Charniak E (1993) *Statistical Language Learning*. Camb.: MIT Press.

time here in Australia. And other technologies like constraint programming[37] as well, which are very trendy now in manufacturing, were also invented at that time. So we are in the middle of a very hot summer for AI right now. And I want to tell you what actually changed because there are three things that were really important in actually bringing AI where it is today.

The first one is data. It's the web, the point that we have access to a very, very large amount, a massive amount of the data that have been produced by humans for many, many years. This was not the case before.

The second thing that I'm sure you are all aware about is our GPUs. Now, I heard about GPUs in 1995 because they were used in computer graphics for actually doing video games and so on. But in a sense, the mathematics behind computer graphics and the mathematics behind AI is actually the same. They are the same concept. And this is why GPUs now are so useful for AI technologies.

The third thing that people typically forget is the democratisation of the tools. Big tech — Meta, Google, and so on — were actually very, very good at releasing software that made the technology accessible to a very wide range of the public, a large number of scientists. They were able to build the new machine-learning models and the new tools that AI is relying upon these days. This has made a huge difference for the field.

This is an example of what happened in 2010. People at Stanford started building this massive amount of data, concerning images. It's called ImageNet.[38] It was at the time a very, very large data set of images. What happened is the computer vision system got better and better. In 2015, they started beating humans in actually classifying images.[39] So this progress in having this data set, and then the GPUs, made the big difference in computer vision, and they started to be used for actually diagnosing cancer,[40] for recognising the faces of people when you board a plane and things like this.

At the same time, researchers on recurrent neural networks made a lot of progress as well. This is what was at the core of Apple's Siri and Amazon's Alexa. It's all about recognising speech, recognising language. Bloomberg said that this was the main commercial outcome of AI technology up to that point. This is natural-language processing. This is the technology at the core of ChatGPT these days. It relies on predicting the next word when you see a sentence, and so the next set of words. Google made a big, big contribution to that in around 2018 when they introduced the transformers technologies. This is what is used in ChatGPT. It's all about predicting the next word. It's a statistical machine, and I'll come back to that when I talk about hallucinations and so on. And, of course, we got Waymo. If you go to San Francisco, you can take a trip on Waymo. If you go to Atlanta or Phoenix, you can get trips to move from one point to another using this technology, which once

37 See Rossi F, van Beek P and Walsh T (2006) *Handbook of Constraint Programming*. Elsevier.

38 See https://www.image-net.org/index.php [Ed.]

39 See Markoff J (2015) A learning advance in artificial intelligence rivals human abilities. *New York Times*, 10 December. [Ed.]

40 O'Connor M (2019) AI may help radiologists reduce missed breast cancer cases, June 18, *Radiology Business*. https://radiologybusiness.com/topics/artificial-intelligence/ai-radiologists-reduce-missed-breast-cancer-cases

again was invented in the late '80s and beginning of the '90s.

Obviously, when people talk about AI these days, they talk about generative AI, the large language models (LLMs).[41] I won't talk too much about them, but I have already mentioned some other things. I'll tell you what they are good at in a moment. The big thing happened when OpenAI actually released ChatGPT, which changed the perception of AI and made AI come to public awareness. Obviously, we knew about that technology. I don't think any academic would have released that technology. I didn't think big corporations would have released that technology, but somehow it got released. One of the things that people have been building now — and this is going to be the future in a sense — is chatbots based on large language models. And one of the things that they try to do is understand the intention of the speaker.

That's the difference between using a search engine and using a chatbot these days. They try to understand what you are trying to do. And they act — this is not a single tool, this is something they're trying to see, "Oh, you are trying to do this. I'm going to use that tool." They can do content generation, as I mentioned before. They can do summarisation, they can do query answering. They also try to maintain conversations such that they understand what you are trying to achieve. And so this is kind of what the state of the art is in some of these technologies.

2025 is supposed to be the year of agentic AI.[42] Victor actually talked about AI for government. This is one of the promises of agentic AI. What agentic AI is, if I summarise it, is automated workflows. When we look at big corporations, including universities, we are really working at a very low level. Where is the data? How do I access this database? The promise of agentic AI is that we are actually looking at what people are trying to achieve, and are trying to build tools so that that process is automated. People are supported who are actually implementing these workflows. So this is one of the big, big promises. It's not yet realised. A lot of companies are trying to make it happen, but it's not as easy as people believed because we humans know a lot about the processes that we have automated and we need to bring that knowledge inside the AI system.

So this is some of the hype that you see. We won't need any programmers in the next 10 years. I did something really bold this year. I told my students they have projects that they have to implement using programming languages and so on. I told them, "You can use any tool that you want." This was forbidden the previous year. And you know what? The grades didn't change. They are the same because the fundamentals are not going to change. You still need to understand what you are doing when you are actually using AI.

What about doctors? Are we going to need doctors? My daughter is a pediatrician in intensive care. She came to me about six months ago saying she never wanted to know about computer science. She thought I was working too hard. And six months ago, she called me and said, "Daddy, I want to learn about AI because I need to diagnose these kids and the doctors don't have

41 See https://www.esafety.gov.au/industry/tech-trends-and-challenges/generative-ai.

42 See https://www.ibm.com/think/topics/agentic-ai [Ed.]

enough information." I'll talk about some of the tools later on, but we still need doctors, but we need better tools for them. It was predicted 10 years ago that they would not be a radiologist and that we would have autonomous diagnoses everywhere. As far as I know, all the radiologists are still commuting to work and looking at X-rays and MRIs at this point. They will replace lawyers. We'll talk about the law. I can tell you that I don't think they're going to disappear, but they are going to be much more productive.

And then teachers will replace teachers. Again, I think education is going to change fundamentally, but we'll still need teachers. So the role of generative AI is going to be a productivity tool. It's going to improve productivities.

"AI won't replace humans, but humans using AI will." This is a quote I like, and I'll come back at the end of this talk to tell you what I think the future is going to be, but this is part of the future. The people using AI are going to be much better than the people not using AI. This is going to be a tool that essentially everybody should use. When you talk about generative AI, it is a hallucination machine. Hallucination is not necessarily bad. As a scientist, I have to hallucinate. That's part of my job. But the public has to understand about the concept of hallucination, what these technologies are actually doing such that they can use them properly.

So AI at the edge is going to make a big difference. It's going to change the way we educate. It's going to change education. It's going to change the way we learn. It's going to change the way we access knowledge. Some people expressed doubts earlier about knowledge, and I think this is a big concern, but our relationship with knowledge is going to change. It's going to change the way we communicate. This is also something that we need to take care of. And it's also going to change the way every profession is going to operate.

It's also going to change the way we think about us being humans compared to machines. We can talk about that a lot. This is where I think there are a lot of dangers. A lot of people anthropomorphise AI based what the technology can do now. But I can tell you that the technology is very, very far from actually being human. I think these conversations sometimes are somewhat dangerous because we are much more than an electrical system. We are also biological machines, chemical machines and so on. It's very interesting when you look at popular books like science fiction books.[43] They all talk about consciousness and things like this. So if you read this book, this is actually pretty interesting.

I'm also doing the national AI strategy for the country of Panama. And I want to touch upon some of the issues that you mentioned because of all these dangers: how the technology is going to affect knowledge, understanding how we communicate. I think you need regulation. What we do for the country of Panama is telling them how they can actually boost their economy, where they have to invest in AI. But we also want to give them a governance framework:how to make the government much more responsive, how to make sure that you have the sandboxes that you can experiment in, and so on. So governance and equity are super important. In Panama,

43 Such as Dan Brown's *The Secret of Secrets*.

you have a very modern city and very rural areas. We have to make sure that whatever we propose to them will boost equity. We are going to boost every segment of the population.

Let me talk about the hope now. This is going to be very different from what you hear about AI. This is what we are doing in the AI Institute that I'm leading. In a sense, this is the road less travelled. This is what people don't talk about in the press, but this is where I believe AI is going to make a huge difference.

Science is going to be completely revolutionised. You probably have heard about what Google did with AlphaFold.[44] It's actually completely changing the way you can do a structure prediction for proteins. I worked on that topic. I had papers in that topic 20 years ago. What they did is remarkable: it shortened the time to discovery by about six months. They got a Nobel Prize in chemistry for this.[45] They deserve it tremendously. So science is going to change completely.

Engineering is going to be completely revolutionised. I'll touch upon some of these topics later on again, but I want to talk about engineering as well. What I'm going to do now is tell you a couple of stories, stories that we have in my institute. There are some stories that exist, but there are many others, but these are stories that I know very well.

What I want to convince you is that AI is actually ready to be deployed in critical infrastructure. Not only is it ready, it's going to be the critical component. If you don't have it, you won't be able to manage this infrastructure correctly. And it all comes from the fact that the infrastructures that we are building are getting bigger and bigger, and humans have a lot more difficulty in controlling them. We have to control them in real time, and I'm going to show you examples of that. But the reality in the field is that this big infrastructure doesn't change very fast. When you think about a power grid, you don't build a transmission line every day. You don't put a new key gantry in Port Botany every day. You don't change your assembly lines very quickly.

Those things are very stable. So you end up running into the same problems over and over and over again with a lot of data, because through digitisation you have accumulated a lot of data. And we are also very good these days at forecasting the future. Therefore, we have a lot of forecast data that we can use. As I told you, data is the new electricity and I'm going to show you how we can leverage it. And I'm going to talk about power systems. This is one of the areas which are personalised. And this is, I think, a fundamental area for the future. This is how you run a power system in the United States. It's very much the same in Australia or in Europe. You try to decide who is going to generate power for every hour of the next day. That's the commitment of the generators.

Then in real time, you have to choose how much they are going to produce. Because, to keep the lights on, you have to make sure that at any point in time, the generation and the consumption are equal. And so what you're going to see in the next slide is from the French system, but this is like if the US would take over France. I'm not saying this

44 Jumper J et al. (2021) Highly accurate protein structure prediction with AlphaFold. *Nature* 596: 583–589. https://doi.org/10.1038/s41586-021-03819-2.

45 See https://www.nobelprize.org/prizes/chemistry/2024/press-release/.

is a good idea, but this is how they would run the system. On the left, you see the energy that is being produced. The big circles are the nuclear generators. They're not at the right place intentionally. You see also renewable energy there in the Southwest. On the right, you see the reserves. The reserves are there if something happens, if you lose a generator, if the lines are down.

And so these things work. Why am I talking about this? I'm talking about this because there is increased volatility and increased complexity these days, because of renewable energy, because of batteries, because of the data centres that were mentioned before, because everything is being electrified, and also because you have a lot of things that the system doesn't see, with solar panels on the roof. There are cars that are being charged and you don't want a blackout to happen.

This happened in Spain about three months ago, they had a blackout for about 24 hours. What people would like to have now, what the operators of this system want, is a real-time assessment platform. They want to run a lot of scenarios very quickly to actually see if there is a risk in this system. But one of these scenarios is 15 minutes. So you have to run your system in like two minutes.

What are you going to do? We need two things. We need a crystal ball for predicting the future, that's predicting what the load is going to be, what the renewables are going to be. We have made a lot of progress in actually forecasting technologies in the last couple of years. And then you need speed. This is where we take these optimisation problems that we have to solve. We use the massive amount of data and we use machine learning, because machine learning can approximate any function. Then in real time, you have tuned your machine-learning model and you can run it on a CPU and take milliseconds. Essentially what I'm telling you is that what we are doing is we are taking what we used to do in real time and moving it offline. We are training the machine learning model.

The analogy that I can give you is when you do sports, you train your body such that you build this muscle memory. That's what we are doing here. We are basically building the muscle memory of the critical system, such that in real time it goes really, really fast. But we have to keep the lights on. So machine learning on its own is not sufficient. You have to have this integration between machine learning and symbolic reasoning, between optimisation and mathematics, essentially. And so this combination is what you're going to see in the future, these combinations of generative AI of prediction methods and the symbolic part of AI that people have been working on for many years as well. When you do that, you solve the problems, near-optimal solution in milliseconds, and you can train this algorithm in about 30 minutes. What happens to your system when you do that?

Instead of taking 15 minutes, you replace every one of the optimisation models by these machine-learning models, this AI system, and now you are going to take five seconds. It's a game changer. We have a tool that could not be done five years ago, and it can be used for operating your system. 15 minutes, five seconds. Roughly speaking, this was the same thing for actually more complex models. We can do that for hydro, and batteries, which is going to be very important. You improve the efficiency of the system by 7%, these are billions of dollars,

and you can do that again in milliseconds, very, very fast.[46]

I'm going to talk a little bit about logistics and supply chains now. This is a project that we had with one of the largest packet delivery services in the United States.[47] They run thousands of terminals. They send millions of packages every year. One of the things that they have to do is load consolidation. They have to choose the path that these packages are going to follow, such that they minimise the number of trailers, the number of trains, the number of planes. There are many routes that the packages can take from Atlanta to Chicago. The key point there, again, we use this new AI system that I talked about to move from what took about an hour into three seconds. Now we can start automating this and the planners can actually look at this and say, yes, this is a good plan. No, I need to change this because I know something that the system doesn't know.

There are always planners in the loops. This is probably the most compelling application that we have built. So this is based on a very large supply chain you're going to see in the moment. But they start from the bill of materials, the suppliers, the way you build the products. It can be chips, for instance, chip design, the way that the customer demands, and you do that from a very large number of SKUs for an horizon which is about three years. The problems are massive. They're massively scaled. I have never seen problems that large. They have very large numbers, which is very difficult for machine learning in general. This is a problem in the pharmaceutical industry supply chain.[48] Again, we move from 40 minutes to three seconds. For a planner, it makes a huge difference because now they can actually look at different scenarios in real time.

Consider the semiconductor industry. It's one of the largest semiconductor factories in the world. This is two million variables. Half of them are integers. And at this particular point, again, we move from 40 minutes to about a couple of seconds.

Fashion is being affected by AI. This is with a company that we work with in Colombia. Typically, what is happening in the fashion industry is every six months you decide what you're going to produce for the next six months, what kind of things we want. I mean, this is boring, but what young people want to wear in the next six months. And so typically they send people all over the world to actually look at what is happening in Barcelona, in Milan, in Paris, Tokyo, in Sydney, and so on. And then they decide what to manufacture.

But currently they don't do that anymore. They scrape the web. They actually look at the fashion shows. They are looking at what

46 Zhang H et al. (2025) Weather-informed probabilistic forecasting and scenario generation in power systems. *Applied Energy* 384, 15 April 2025, 125369. https://doi.org/10.1016/j.apenergy.2025.125369; Tuo M et al. (2025) F. Inertia-constrained generation scheduling: sample selection, learning-embedded optimization modeling, and computational enhancement. *IEEE Transactions on Power Systems*. https://doi.org/10.1109/TPWRS.2025.3639407; Tuo M, Li X and van Hentenryck P (2026) Machine learning-assisted dynamics-constrained day-ahead energy scheduling. *Electric Power Systems Research* 252: 112449. https://doi.org/10.1016/j.epsr.2025.112449.

47 Ojha R et al. (2025) Outbound load planning in parcel delivery service networks using machine learning and optimization. *Transportation Science* 59. https://doi.org/10.1287/trsc.2024.0672.

48 Akhlaghi VE, Zandehshahvar R and Van Hentenryck P (2025) PROPEL: Supervised and reinforcement learning for large-scale supply chain planning. arXiv. https://arxiv.org/abs/2504.07383.

the images of people in Paris, in Barcelona, and so on. And then they do price optimisation, but not only their own data, but they can look on the web what the competitors are doing. It's a completely different game. And this requires multimodal learning using images and text and so on.

I want to talk about healthcare because I think, as Victor said, healthcare is going to change dramatically. Let me share a quote from a TV documentary.

> Right now, the pathway through a healthcare system for somebody with dementia is a journey. It's a nightmare just like it is for anybody with any complicated chronic disease. It's really tragic because most people still get delayed diagnosis, no diagnosis, rural America, underrepresented minorities, there is such poor access to healthcare that that's number one.
>
> — PBS Video, Allan Levey (head of the Emory Goizueta Brain Health Institute[49])

This was from a PBS documentary in the US. I'm proud of that at the end of the day, but I'm going to talk about what is happening. Victor talked about brain health before. In the United States, if you want to have an appointment with a neurologist in the state of Georgia, it takes about 10 months. It's predicted to actually get to two years in the next five years. Delayed care in brain health is brain loss, so you really need to change the healthcare system. We are trying to look at the workflows and changing them completely. You could do the same for government, for business scope, the way you run business systems as well. It's thinking differently about how you organise care or government. What we are doing with Dr. Allan Levey, who's a very famous neurologist for Alzheimer's disease and so on, is building a brain health centre.

What that means is that we are using AI technology to do better triage, to do better diagnosis, to do better treatments. AI is used in every one of the steps, but there are obviously doctors in the loop all the time. This is how you do diagnosis triage. This is a chatbot for triage, but it's very different from what you would be using with an LLM. You are actually being driven by what is called the test, a CDR test, which is actually looking at various symptoms and trying to diagnose them. But this is still a conversation. We still use natural language processing for doing this, but we are driven by the medical knowledge. And what we sent — and this is very interesting — is something like this. What you see is essentially the summary of the conversation that the chatbot has with the doctor. What is very important is when we make a recommendation, we actually tell the doctor, "This is why we are making this recommendation." So the provenance of the recommendation is there.

I was very lucky to go to Tasmania a few years back and I was very stunned because in Tasmania, you know where a particular animal is being grown, what grass is being fed. And when you see that particular piece of meat in the US, we know exactly where that particular animal came from. This is the same here. We have the provenance of the data. We know exactly why we are making this decision, and this is going to be super important.

I told you stories. There are many more stories. I think one of the things that's going to happen is that we are going to use AI for

49 See https://goizuetabrainhealth.emory.edu/about/index.html [Ed.]

creativity, designing new products, designing new molecules. This is a project that we have inside the Institute: how do you build new enzymes? It's called combination. It's very interesting, combination of generative AI and Bayesian optimisation. We design better aircraft. This is again, using design now and AI for actually generating new ways of thinking about this that were not possible before.

We are building new plants by not using the soil, but directly giving nutrients to the roots of the plant, monitoring how they grow, monitoring using computer vision, and then using reinforcement learning to actually control the way they look, the way they taste, and the way they grow.

Finally, I have one more. Manufacturing is going to change completely as well, detecting defects much more quickly and integrating with supply chains.

A last thing: I love Sydney. Sydney has a very big financial industry. The financial industry is changing very quickly, but there is also massive risk. Fraud detection has to be taken care of. When I work with Panama, this is one of the big recommendations that we are telling them. You need to be on top of this. There is massive data. The transactions are going so fast that you need to contact the malicious agents using AI as well.

The future of X is the future of AI plus X. I think the future of X, where X is any discipline, is going to be the future of AI plus X. You do not ask an engineer if they are using mathematics these days. In five years' time, you won't ask them if they are using AI. It's going to be obvious and it's going to be the same for every profession. I believe the future of X is the future of AI for X. Thank you very much.

Questions and Closing

Sally Cripps: What a fantastic talk. We've got about five minutes for questions, so I'm going to throw to the audience and see if there are any questions.

Q1: My question is related to the usage of water in generating electricity. Jeff Bezos said that to solve it, you have to go into space and use the very low temperatures there and 24 hours of sunlight. What do you think?

Pascal Van Hentenryck: I agree that, and this was mentioned by Her Excellency before, that AI is consuming a massive amount of electricity, a massive amount of water, and water is getting really a scarce resource at this point for the Earth. I think that AI can help in actually managing these processes much better. One of the things that you need to think about is how we manage these very large data centres in a much better fashion, how you combine them with renewable energy, how you combine them with batteries and so on. People have to start thinking about this, and how you are actually using the data centre itself. This is a big issue, but again, AI can actually help AI there.

To reassure you, I believe that at the level of the large language models, we are going to converge at some point. The progress is going to slow down tremendously there because we have already trained them at massive scale. So don't worry so much about that. I agree that we have to start looking at how we use water, but we also have a lot of other options, which are actually optimising the way you use these very big machines,

and they are very big. But at some point, I think there is not much more that you can train them on compared to what they have already been trained on.

Q2: Thank you. One of the stress points that seem to be happening are deepfakes, particularly in finance. And I'm wondering, what do you see as the future of the solution of this. A huge number of people, a very smart deepfake that can simulate people, simulate voices, simulate everything.

Pascal Van Hentenryck: I was very happy that Her Excellency actually mentioned that. And this is why I talked about Tasmania. This is why I showed you the example in healthcare. I think provenance is the real issue. For everything that's going to be produced, we have to understand what data has been used. It's also what has been used in the process of creating something, and you need to be able to validate that. If there is no provenance, you should not allow the system to actually publish these things. Now this is an academic answer. We need a regulatory framework to actually do this, but the technology is there to do that. We have to be able to say, this is where and how this has been produced. This is the data that has been used. This is why you believe this is truthful and trustworthy. And this is kind of like, you have a publication where you have essentially the sources that have been used. And if there are points where you don't know, you're not sure, they should be highlighted. Obviously, you need the regulatory framework to do that. I hope people are going to do this.

Sally Cripps: Watermarks and images?

Pascal Van Hentenryck: Yes, absolutely. Absolutely. And to see this image, this has been created using this, this and this and so on. And the technology, you can't do things like this. It's possible to do this from a technology standpoint. So it's a great question. I'm glad that I could answer it because it gives you a sense that the technology is there, but you need the regulatory framework to do that.

Sally Cripps: Thank you, Pascal. I know there are so many more burning questions, but in the interest of time, I'm going to have to shut them off. You will be getting a chance with Pascal in the very last session today. So hold those thoughts and you can ask him those questions then. Thank you.

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, pp. 68–82. ISSN 0035-9173/26/01068-15

2025 Royal Society of New South Wales and Learned Academies Forum: "AI: The Hope and the Hype"

Panel Session 1: AI and the Law[1]

Emeritus Professor Rosalind Croucher AM FAAL FRSN FACLM(Hon), NSW Information and Privacy Commission

The Hon Dr Annabelle Bennett AC SC FAAL FAA, retired judge, Federal Court of Australia

The Hon Justice Melissa Perry, Federal Court of Australia

Professor Lyria Bennett Moses FRSN FASS, Faculty of Law and Justice, UNSW Sydney

Rosalind Croucher

I'm delighted to be able to chair this session on AI and the law. I thank you, Your Excellency, once again for hosting Ideas at the House and the Learned Academies Forum. This is such an expression of your own commitment to the work of the academies, the reach to community, and essentially all that we hope to achieve and aspire to do in the Royal Society. In introducing a panel on AI and the Law, a very necessary first panel if I may say in today's proceedings, I want to observe regarding Sally's framing of the day in terms of being "at the crossroads of rampant hope and rampant hype," that lawyers are not prone to rampant hype, and we keep our rampant hope for private conversations!

But it does provide a perfect prelude to a sage and sensible session on AI and the Law. And in embarking upon our conversation, Pascal, I have to be a little concerned with one thing that you said about AI and lawyers — that it would make lawyers "more productive." Usually words such as "productive," "productivity," and often their associated KPIs, send a chill down the spine of most lawyers, particularly judicial officers of whom I have two on the panel, but it can be a more finely calibrated assessment of what that "more productive" means. But I do bring it to the table as part of our thinking.

The panel that I'm chairing represents Fellows of the Australian Academy of Law. I'm delighted that the AAL takes its place among other learned academies in this forum today. As a quick background, the Academy of Law is the newest sibling, if you like, in the community of Learned Academies. It was launched in July 2007 at Government House in Queensland, when Her Excellency Quentin Bryce was then the Governor of Queensland. And it was the outcome of a process begun by the Australian Law Reform Commission in a landmark report on Managing Justice, a Review of the Federal Civil Justice System. The idea for the AAL was a key recommendation to bring together the three strands of the legal profession — the judiciary, legal practitioners, and legal academics — united in promoting high standards of learning and conduct and appropriate collegiality across the profession. (I note productivity isn't in there.)

The Academy of Law already holds an annual symposium with the Academy of Science, and today we expand that connection with other learned academies. The panel includes three extraordinary AAL

1 This is a transcript of a presentation, available at https://www.youtube.com/watch?v=6fEN3zWI-NM

fellows of great distinction — a tiny insight into which you can see in their cameo CVs that are contained in the programme. This panel exemplifies the eminence of those who are elected as fellows and the nature of the intersections of excellence across the branches of the wider legal community. The Honourable Dr Annabelle Bennett, AC SC, is also a fellow of the Academy of Science; next, the Honourable Justice Melissa Perry is a foundation fellow of the Academy of Law; Professor Lyria Bennett Moses, who is also a fellow of the Royal Society and of the Academy of Social Sciences. You are right to be in awe — as am I.

AI and the Law, the framing of the critical question for this forum sets the scene perfectly for our panel, namely: is AI truly different from previous transformative technologies, like electricity or the internet, or is it simply another innovation riding waves of hype? In 1970, there was a critical case in the High Court, *Mount Isa Mines Limited v Pusey*, in which Justice Windeyer spoke then of the law "marching with medicine, though in the rear and limping a little." The same can be said of any new technologies. In the rear? Limping? I pose in a rhetorical questioning way.

Each of the panellists will provide a short presentation. Dr Bennett will be looking at the tools we have now, particularly in judicial decision-making. Professor Bennett Moses will take us into a world when the law, or common law in particular, is not enough — and the tools of regulation. And Justice Perry will complete our wonderful triptych, in defence of the human. Throughout that, I'll add a few interrogative questions of my eminent team and then provide an opportunity for you similarly to have that opportunity at the end.

May I please invite Dr Annabelle Bennett, my very esteemed co-panellist to make her presentation.

Annabelle Bennett

Thanks, Ros. Yes. First, I have absolutely no expertise in AI, which is probably why I've gone first, but I do have part of my lifetime in the law. Before I start, I wanted to share with you an analogy that came to mind. Many years ago, there was a law conference in Alice Springs and we were all climbing Uluru, as was then permitted, which was then called Ayers Rock. And as I was coming down, the person walking in front of me was a rather large man. I took a photograph and obviously it was from behind him. And when that lovely man, and he was charming and became the Chief Justice of Western Australia, I sent the copy of the photo to him and I said, "What's this worth?," he replied, "It depends on whether you consider it a means unto an end or an end unto itself." And that came to my mind when I was thinking about this topic, because it really is a question of whether AI is the end unto itself, which would be replacing judges — for example, I am a sci-fi fanatic and used to watch Buck Rogers in the 25th century, Dr. Theopolis, who did everything and made all decisions and all judicial decisions — or whether it's a tool.

If it's a tool, then you've got to see how you use it. And of course, I'm picking up a lot of the other themes that have been mentioned because they are the common themes that come to mind. Let me just give you a couple of examples. When Dragon Dictate first came in, there was a question whether we could use that for draft judgments. It was a real question of editing because of course it didn't always transcribe things

perfectly. And if it, for example, missed out the word "not," because you spoke quickly as I do, then the answer that when you say the appellant succeeded or did not succeed, you can see that there can be a difference in the outcome of a case if you don't check it very carefully. That has happened to me once. I was walking into court and I was reading a little conclusion (this wasn't even AI) and I saw the word "not" had been left out. That had been through five edits and I only picked it up at the last minute. That's the sort of thing that can happen.

I would like to know how many of us have sent off a text message or an email where auto-correct has rather embarrassingly changed what you intended to say and something was totally wrong. So it becomes a question of editing and checking. I think it's a fantastic tool and we can use it, but again, we get down to two things.

Her Excellency raised the question of evidence. Of course, for lawyers and judges, evidence is paramount. You make decisions on the evidence that's presented, but then you have to be sure that the evidence is correct. And so you have to be able to check it out.

Let me tell you the trouble with checking — I'm picking up on your comment about provenance — sometimes you need to know what it is you're checking. Last week, I delivered the Sir Garfield Barwick Oration. Being a very, very good little girl, the first thing I did was to go to ChatGPT because I was going to speak about his extra-judicial activities. It gave me a wonderful list of his activities, and one of them said, with the years and everything included, "Chancellor of the University of Sydney." Now, I happen to know, because I've been in that academic world myself, both at the ANU and at Bond University, that he was chancellor at Macquarie University. So I typed back into my little conversation with ChatGPT, wasn't he the chancellor at Macquarie? And the response came, "Good pickup". But when I redid the search a few days later, just to check, it still told me the University of Sydney, so it hadn't changed.

But had I not known that, I would have delivered this lecture with all of his family present and made the fundamental mistake of saying he was the chancellor of the University of Sydney. I raise that as an example of a provenance issue. There's one other thing that I want to mention — and this gets back to the productivity — I feel like everyone's stolen all my thoughts in advance. There are two issues that come up here. First, if you rely on ChatGPT or AI to do your decision making, that is where the danger is. I would go so far as to say it's an impermissible delegation of the judicial power of the Commonwealth. There'd be enormous constitutional questions arising if you did that: whether you were actually exercising the judicial function. But at the same time, judges are actually very concerned with productivity.

And of course there's that longstanding expression, "justice delayed is justice denied." If you consider people hearing things without an electronic transcript, writing everything out by hand, thinking it through, rewriting drafts, then you can see the time it takes, even after a case is closed to get that decision to judgement. There's a lot of pressure on judges to get decisions out as quickly as possible. Tools have helped us do that: word processing, drag and dictate, all the other searches. Then you get into a balancing question of what's worth what.

What I've come to realise in a number of other areas that I've been involved in, especially well known in the commercial world, is the word "risk." You've got to have a risk matrix. You look at the risks and sometimes you have to accept that there is a degree of risk that is worth accepting in order to take advantage of what it is that a particular technology can offer. But that may mean that in particular cases, that risk will be realised. It's not much use to say to that litigant, "I made your decision because I checked that there were decisions around the world. I wanted to have harmony in the world, particularly in my area of intellectual property. I got a wrong summary from ChatGPT and I made your decision accordingly." It's okay to say, "Well, that only happens in 0.001% of the time." But for *that* litigant, that's the difference between success and failure, between prosperity and bankruptcy, between all of the consequences of that. I leave that with you.

The only other thing I thought I'd raise is interesting. I'm sure most of you would know that in the world of intellectual property, there are real questions to be asked about authorship and ownership. Bearing in mind that the whole idea of statutory rights, intellectual property rights, is meant to stimulate invention, to create an incentive for inventors to invent and to give their invention to the public, which is the quid pro quo for a patent. There have been cases all over the world now about whether a machine can be an inventor. Of course, some judges said yes, but ultimately most appellate courts have said no. And they've gone to the statute to say that the inventor still has to be a person. But the more interesting question has arisen, for example, in copyright.

There are two cases — one in China and one in the United States — that are most interesting, both of which, I want you to know, involved the creation of beautiful women. That's obviously something that engages a lot of people. Every female in this room would be the role model: all they had to do was take a photo and copy it. They'd be fine, but they went to the internet and took all these options that were available. I was told when I had to present on this at a conference in Singapore that there was a divergence between the United States and China on this. This goes to show, for anybody who's a lawyer or who's ever interested in the case, don't just read the summary. You've got to actually read the decision, because the decisions were identical in their reasoning, but the outcomes were different because of the way that technology was presented and described.

It's very important when you're presenting anything about technology — be it as an expert witness or in any legal case — how you describe and how you present it. Because in one case, the technology was presented as one that had human intervention at every choice, an actual decision being made by a human at each stage of the creation of this beautiful figure. And that was held to be a copyright work and that person was an author. But, in the other case, it was presented that it was really the machine that was making the decision, the machine-learning way, and you didn't have sufficient human intervention. And it was held then that there was no copyright work and there was no human authorship, so the person who was claiming authorship wasn't entitled to it.

I want you to know that, internationally in the legal world, these questions are

coming up. What is interesting is with the speed at which we're talking. Victor, this was part of the points you raised.

A few years ago we were talking about legal issues arising from the internet, simply the fact that something goes up on the internet and it goes all over the world, and where is the infringement taking place? Now it's getting down to a much harder question. It's not just theory now, because this is now impacting us, it is impacting the lawyers, it's impacting the judges, and it's an evaluation process. It's all very well to say you've got to check the provenance, but you need to have a sufficient degree of knowledge to check the provenance or have the tools available to do that.

It gets down to "what you don't know, you don't know." But I would say that I think it should be looked at as a tool for the judiciary and for the lawyers, but you've got to be very careful how you use it and how you check it and how you present it. But it has the possibility of increasing productivity, may I say, Ros, both for lawyers, because don't forget everything you can speed up — bearing in mind most lawyers charge by the hour, you will reduce the ultimate cost to the client, to the community — you'll hopefully get decisions out faster.

Hopefully they are just as good. Just as you haven't been able to do away with radiologists, I'm firmly of the view that you will not be able to do away with judges. I'm hanging onto that thought for grim death, even though I'm no longer a judge.

Rosalind Croucher: Thanks, Annabelle. The productivity question is an interesting one because it's about the use of tools and the distinction between a tool and the decider. I think that was very well put. I'm going to proceed straight into Lyria's presentation, which takes us into the world of regulation and where we can add things. Some of the conversation so far has been focused on the common law and the interpretive role of judges and decision making. We're going to broaden it out and we'll loop back on a number of themes as we go. So over to you.

Lyria Bennett Moses

I'm going to be relatively quick. One can talk about this stuff for a very long time, but I wanted to focus on the regulatory question. I noticed in the introductions a number of people brought it up. The idea of there's going to be all these issues and regulation can be one of the solutions. I want to unpack that a little bit, but what people mean by that, what they think it can do and how it might be approached. I was going to start with the definition of AI, but our esteemed Governor has already done that for us. There are only two things I'd point out. One is that the definition has changed. This shows the difference in the OECD definition of AI between 2019 and 2023. So it has changed quite significantly. The other one is to point out that when people talk about hope and hype, this isn't necessarily the targeted defined thing about which they feel that hope and hype.

I think the hope and hype has been continuous since 2019, even though the definition has fundamentally changed. At the core of the new definition is the idea of inference. These are systems that infer *how*. That wasn't even in the definition in 2019, although the continuity in hype and hope is there. If you think about inference, I don't think it's *inference* that people are concerned about. There are many things and inference can be part of it, but I think if you try to describe the AI that people have hopes for,

concerns about, you'll end up with other concepts coming in as opposed to the way it's currently defined. We're dealing with an emergent phenomenon, in other words. This is not a clear thing that we clearly understand its current existence and future. But I also want to highlight that hope and hype is also something we can feel about regulation. I think there has been a lot of hope in regulation in some of the introductions — that this can be the thing that protects us. Let's see if that is in fact the case.

So why regulate? There's a whole range of issues that come up in AI. I picked one: the algorithmic bias or discrimination problem. Most people are familiar with at least some of the examples of this: these hiring algorithms that prefer men, et cetera, et cetera. The point is that systems, defined AI systems, when they discriminate, do it differently to humans. A human might discriminate say on the grounds of sex because they have a particular view of women as not being able to perform a particular job, whereas systems will typically be black boxed. It'll be probabilistically related based on historic data, influenced potentially by variables such as sex and race or what have you — even when it's not directly impacted by those variables, because those variables are not inputs into the system, impacting people differently by proxy variables based on similar characteristics. My point is that the challenge that we face around discrimination, around fairness, is not an AI problem. AI is part of a broader problem that is all in fact connected to each other, often through the data on which these systems are learning. So what does the law currently do? The core point is that the law, as it is currently written, was for the imagined scenario, which was at a time when the law was written for a human who might be sexist, racist, et cetera.

The goal of the law is to stop that person making a decision, say not to hire a woman because she's a woman or in a whole variety of other contexts. That is what the law is designed for. We have different laws in Australia. It's a bit of a complex mess. There's state law, there's federal law, there's different laws sometimes for sex discrimination and race discrimination. They all use slightly different words. This is not going to be that lecture, but to take direct discrimination as an example, you'll typically find words in the legislation like "on the grounds of." Does someone make a decision on the grounds of, say, sex? Of course, when you're looking at a *system*, it's not that clear. We can't talk about "on the grounds of" as if systems are doing it on the basis of one thing, or that we'll even *know* what the grounds are for a particular system output.

It's a lot more complicated. You have other concepts and terms, but where you often go when you're looking at the kinds of things that systems are doing in the context of unfair treatment is we move into *indirect* discrimination, which is where you say the condition of hiring is being able to go through some hiring algorithm which goes forward to the next stage and has the effect of disadvantaging people, say on the grounds of sex. We're going to pick different acts here, but that's the core concept. The problem of course is that this will almost always be the case, at least in some conceptions of what we mean by different treatment. In other words, it's very hard to think of a system where that will make absolutely no difference at all.

What happens then is the wonderful machine-learning community tries to

enlighten us. And they create a whole raft of definitions of fairness. There are over 30 of these things, and they're very precise because the one thing I love that comes from that community is its precision. When we talk about fairness in law, we talk about discrimination, we have a general sense of what that is. We use words, they use mathematics. It is a much clearer sense of what particular fairness metrics might require. That is fantastic except for two things. One is they're mutually inconsistent. You can't do all of them at once. And the second problem is that how does that fit in with what the legislation is telling companies to do in discrimination legislation? And it is not a great fit.

Statutes are built on words, as opposed to common law, which is a little bit more flexible. Statutes are built on words and those words get interpreted. "On the grounds of" means what "on the grounds of" means. You have this language and legislation built around one problem type, the sexist racist human. You have systems that can potentially solve for some of these problems. And then you have law that's not actually helping organisations to make sensible decisions and may in fact be pushing them to do things like eliminating variables that might actually increase the problem, but nevertheless avoid categories like "direct discrimination" that the legislation prioritises and has fewer exceptions. You've got a misfit. It's a complicated misfit, but you've got a misfit between what these systems are doing and the problems they're creating and what the law is requiring organisations to do.

Ultimately, you have machine learning trying to help, but the law's not talking to those kinds of distinctions and we're not necessarily getting what we want. So what do we do with regulation? Do we do something? And I think regulation often becomes a narrower conversation than it needs to be. This is my very giant picture of regulation. First of all, not all of it is government. Standards organisations are regulatory. There's a whole bunch of other stuff that goes on. But even within government, there is a whole bunch of stuff. It passes laws. It creates ethical principles and in particular has done so around AI. There are guidelines and all sorts of other documentation. It buys things or doesn't buy things based on particular criteria. It has a whole bunch of influence levers.

Let's even narrow down to law. The first point to make is that not all law that is regulatory is going to have the words "AI" in the heading. It's not all about the "AI Act." I've already spoken about discrimination law. Now that isn't a good fit, but it nevertheless influences decisions in organisations about the kinds of AI systems they build. They don't want to fall foul of those laws. We can talk about law being inadequate or not sufficient, but we can also think about the role those laws might continue to play if they are done better than they are now. Of course, you have the possibility of law that specifically targets AI. But there you have problems — because again, you've got to define it. Here's the EU AI Act definition:

> "AI system" is a machine-based system designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.

Again, what you find is that the act targets a technology. It's one imagination of what we're scared of, but it's not necessarily always the right one. The example I like to use to show this is one of the prohibitions in the AI Act.

> The following artificial intelligence practices shall be prohibited:
>
> (a) the placing on the market, the putting into service or the use of an AI system that deploys subliminal techniques beyond a person's consciousness or purposefully manipulative or deceptive techniques, with the objective, or the effect of materially distorting the behaviour of a person or a group of persons by appreciably impairing their ability to make an informed decision, thereby causing them to take a decision that they would not have otherwise taken in a manner that causes or is reasonably likely to cause that person, another person or group of persons significant harm
>
> (*Artificial Intelligence Act* article 5(1)(a))

You can read it if you want, but I want you to imagine leading the text in red. These are about the use of subliminal techniques to manipulate people. We don't want that, is what they say. That's prohibited. But only when you use an AI system, only when you fall within this very particular conception of a particular technology and not in general. But why not? There's nothing in this where I feel that the removal of the red text would make it worse and it would in fact expand it to cover a broader range of the problem of manipulative techniques. That highlights, I think, some of that problem. And where do we end up? We have choices in regulation. We can focus on the technology, we can regulate AI, but if we do that, we're focusing on a current imagined definition of what we think this technology is and we're passing laws about that *particular thing* and imagining *that* is the problem.

We're not seeing it, as I started with in that broader context of broader problems where the technology may be playing part of a role, but isn't the whole story. It will become obsolete because whatever we think AI is now may not be the thing that's coming. There's a whole bunch of problems in doing it that way. One can do it the way the US did under the former Biden administration, which is sector-based. You do something in education, you do something in commerce, you do something in housing, et cetera, and you say, these are the rules for that sector. That's another approach. Or you can — and this is the one I like to push — start with the values. Start with what problem you're trying to solve for. If we say we say want to solve the problem of discrimination or certain kinds of discrimination, then you need to think quite deeply about what are your red lines, what are the kinds of practices and behaviours you want organisations to do when running systems, say for hiring people or anything else.

What are the tests we want them to run, the evaluations we want them to do? What *matters*? Because then we can look at all of those fairness metrics that the machine-learning community has very helpfully created and help guide organisations in making the right decisions. That is not necessarily a case of an AI Act. That's a case of looking back at all of those discrimination laws and thinking about what we really want.

Rosalind Croucher: Thank you, Lyria. The drawing attention I think always is on "what is the problem?" What is the problem we're

trying to solve? And then not doing that lightly because often law can sometimes or legislators can sometimes do things quickly, the rapid response — "we need to have a law on X." If it's not thought through properly, it can be compounding in its consequences — the unintended consequence of laws. I'll give you just one quick example before I invite Justice Perry to speak. In the mid-1980s, it was the period of burgeoning use of reproductive technologies, a wonderful assistance to families and to having children. There was encouragement to use assisted reproduction through donor sperm. There were laws passed out of concern to preserve the anonymity of the donor — and it was just simply that. The sperm donor is not the father except in very, very narrow circumstances.

The sperm donor was irrefutably presumed not to be the biological father; problem solved. But what that did was completely overlook the rights of the children who might be born as a result. Their curiosity, indeed, their rights to know their own biological and genetic history, which has led much later to responsive amendments that are in line with the changes that came in the adoption area, where children were able now to ask questions about their own biological past. What is the problem we're trying to solve? The next two parts of the formula are: what are the possible solutions, and then what ones should we adopt? Melissa, could I invite you to address us?

Melissa Perry

Thank you. In the movie, *Eye in the Sky*, senior military personnel from the British and US authorities were advised that it was lawful under international humanitarian law to order an attack from a drone on a cell about to undertake a suicide bombing. Collateral damage being the likely killing of a little girl, selling her mother's bread from a makeshift stall proximate to the attack was considered proportionate to the harm that would be caused if the suicide bombing were to proceed. However, the agent on the ground was moved by the little girl's plight. And with barely a moment to spare, he acted to save her by bribing a local girl to buy the last of her loaves in the hope that she would leave before the assault.

This example illustrates how an exercise of judgment by a human decision maker may mean achieving not merely a *lawful* result, but one which better accords with the fundamental human values of mercy, fairness and compassion in the individuals and the public interest. This is not a mechanical procedure, which is readily reduced to an algorithm leading to an inevitable result, even if that result is not predictable by a human. The example also illustrates the ability of humans to problem solve in creative and imaginative ways. Such fundamental human qualities as mercy, compassion and fairness draw upon our shared understanding of the frailty of the human condition, and lie at the core of our common humanity, resonating through the ages across cultures and underpinning fundamental human rights. Human qualities of the kind which I've described have long informed courts and government decision makers in the exercise of their functions.

While the width of statutory discretions will vary according to context, discretions potentially afford humans the latitude to make judgments and reach decisions, which reflect community, administrative, and international values, and align with statu-

tory objects in the face of a wide or almost infinite variety of human circumstances.

However, these are qualities utterly lacking in any machine technology. No robot has yet been created with the capacity for sentience, understanding, or independent thought, let alone with a conscience, which is a key, if not often a determinative, element used by humans to make choices and reach decisions. Thus, while generative AI can emulate the process of exercising a discretion and may even provide reasons which give the appearance that it's exercised a discretion or engaged in an evaluative process, this is merely smoke and mirrors. This is because the processes which it utilises are based upon statistical analyses and not upon the nuanced reasoning processes in which humans engage. For example, there are systems available which will assist in predicting the outcome of litigation, and these may operate with a high degree of accuracy and be infinitely more cost-effective than engaging a lawyer, but they may have regard to largely irrelevant factors, such as the name of the judge and the amount of money involved in the case.[2]

This leads me to a concern that the massive efficiencies and cost savings which the use of machine technologies potentially afford in government decision making, especially high-volume decision making, such as determining eligibility for social welfare benefits, may encourage governments to frame laws in such a way as to authorise the use of technologies to make the decisions. This in turn may curtail the scope for an exercise of judgment and discretion by a human decision maker to accommodate borderline or unanticipated cases. These concerns are heightened in the context of decisions which affect human rights and fundamental interests. Indeed, even the use of such programs to assist in the exercise of evaluative judgments and discretions needs to be approached with caution to ensure that the human decision-maker brings a properly independent and impartial mind to bear on the issues and interrogates the data provided by such technologies.

An example about which much has already been said is one form of generative AI being a large language model or LLM, such as ChatGPT, Copilot, and Grok. An LLM is a complex algorithm which responds to human prompts to generate new text, representing the most likely words and the most likely word order based on training from massive data sets. The text, therefore, is ultimately no more than probabilities, but the manner in which generative AI functions has a number of consequences, which are also highly pertinent when consideration is being given to the use of such models in the judicial and litigation spheres. I wish to highlight three short points.

First, as Lyria has already mentioned, the risks of hidden bias influencing outcomes from generative AI are based on the fact that the data on which it is trained is historical, but it goes deeper than this. Because generative AI operates by formulating the most likely sequence of words in the most likely order, it may limit exposure to a range of different and competing views, particularly minority opinions, and further marginalise minority and disadvantaged groups in response to prompts. Second, the outputs of LLMs are unpredictable. Devoid of understanding or concepts of accuracy, the

2 See Johns F and Walton K (2026) Why AI will redefine the law degree. *Sydney Morning Herald*, 21 April, p.21, for a discussion of the continuing value of legal education. [Ed.]

capacity of LLMs to hallucinate or make up information and to convey these hallucinations convincingly is well-documented and already been referred to. Infamous examples in the courtroom context include fake citations of earlier judicial decisions generated by AI on which parties have sought to rely to their peril. Indeed, somewhat ironically, in reporting on a recent report prepared for government by a professional services company using AI, the ABC 7.30 Report identified a number of hallucinations, including a fake citation to an article purportedly by me.

Added to this, the way in which outputs have been generated by LLMs are generally unknown and unknowable, even by their developers and programmers. In other words, the systems operate (as Lyria has explained) as black boxes, having completely opaque decision-making processes. Not only therefore is the reliability of output a significant issue, but the use of LLMs raises serious transparency, accountability, accuracy — and therefore audit and reliability — issues. But the metaphor of a black box is also apposite in another respect, given the inability of LLMs to think outside the box in the way that the example I started with illustrates.

The final point I would make is that it isn't just the outcome which matters, but also the *process* leading to the outcome, especially where the outcome is one directly affecting the individual. In other words, AI may give the same answer which a human decision maker would give, and it may give it more efficiently and more quickly. It may also provide the answer drawing upon a vast amount of knowledge well beyond the capacity of a single human being to master. But the process matters, and the fact that reasons given for the outcome are the product of a human being who listened and considered the evidence and the submissions, and brought to bear the human qualities of the kind I referred to at the start in reaching their decision, matters to those whose lives are affected by decisions which we as judicial officers, administrative decision makers, and decision makers in the private sector make *and* take responsibility for.

So those observations are not merely based on my experience as a judge of almost 13 years, where the importance of ensuring that litigants have a real and effective opportunity to be heard and to feel respected and listened to is brought home to me on a daily basis. Studies based upon interviews with litigants show that confidence in the system depends upon the litigant feeling that they were *heard* and not necessarily upon whether or not the particular litigant is successful.

Taking a practical example which I'll finish with. How might *you* as the victim of a scam in which you lost all your savings feel, if instead of a human judge, you read out your victim impact statement to the sentencing robot judge, would you feel listened to? Would it be *enough* to know that the robot judge was more likely to reach an appropriate sentence in line with precedence and community values than the flawed and fallible human judge with a living beating heart?

Rosalind Croucher: The difference in the process of reasoning: we will focus on the judicial decision making for the moment. The law historically is a large language model. Now we've got a lot of law, but that is also the beauty of the way that common law, which is the sum total of decisions that have been made by judicial officers over many

centuries now, is a large language model and is amenable to the process that large language models use. It's when you get to the finer points of the reasoning, and I think, Melissa, you illustrated that beautifully in terms of the human and the flexibility that Annabelle, you spoke of, where the judge can be innovative within the confines of flexible ideas like copyright and who is the maker and so on. There is room in those spaces for that reasoning that is deeply humanly, personally informed, but within the constraints of important processes.

I was thinking in anticipation of our conversation, that when judicial decision making is based on both the use of precedent, and the possibility of leaps when the precedent runs out. I'll direct this maybe to Annabelle, I think of a case like the Mabo decision, which was one dissenting judge and the remainder of the bench where it was a unanimous or unified decision, and it *shifted* the parameters of the common law in relation to the recognition of native title. That leap, how could that possibly happen if informed or otherwise by AI? Not the actual decision, but I can't imagine a world when that decision might have been made relying on the AI tools that are currently on offer.

Annabelle Bennett: That's a hell of a question, Ros. A couple of things that came out as you were talking, though. The first is, you spoke about the common law, and of course that is one of the great strengths in our system, that it gives the degree of flexibility, I think, very much in the human element. If you've got the common law, part of our history is you *can* develop it. What has happened is with a lot of dissenting opinions, over time, those dissenting opinions have become, as the majority have moved towards them, the majority of views. David, my husband, is here, and we acted for Lionel Murphy. At the time when he first came out with some of his decisions, people were appalled about how extreme they were, but of course, in a lot of cases, things have shifted and they're mainstream.

Coming back to the AI thing, the point that Lyria was making is that judges don't *always* have the flexibility to make the decision that they think should be made in the circumstances, because if there is a statute that covers *exactly* what you're looking at, you *must* give that statute effect. And this is where the difficulties are, the sorts of difficulties that Lyria spoke of, of statutes fixed in time. I'll give you one small example that's not in this area, but in terms of trying to do what you feel you want to do and being constrained by statute, I sat on the Court of Arbitration for Sport, I was at the Olympics, there was a wrestler, an Indian woman wrestler who got to the finals — first ever. She failed the weigh-in and the rules that I had to construe said, "up to 50 kilograms." She was a hundred grams over. "Up to 50 kilograms," and she was disqualified.

I wrestled with that all night. I actually wrote it to give her a win, and I woke up in the morning and I said, "I can't do that," because the words of the statute are so clear that, in this case, the regulation was so constraining.

The first point is that this is the flexibility of our system that you described so well. Melissa has given the example of the beauty of giving judges the flexibility to deal with the facts of a particular case. Where you come in though is another one, which there are hundreds of thousands, if not millions of decisions that one could go to to make a decision. A lot of them are dissenting deci-

sions, which no one usually talks about until you look for them later.

Would AI be a fabulous tool to be able with the right prompts to give you that information, which you would never find, and practically you would never do it? Yes, you could have it — as long as you were satisfied of the provenance of it, but you'd have to check that. You could get the little summaries and go away and read the decision and see if it was right for yourself. AI could be a tool to go across the world to look at decisions. Google Translate can tell you if it's from another language. And, yes, AI could be very useful, but, ultimately, as Melissa said, it is up to the individual judge then to take that information, check it through, think about it, take bits of whichever you think are good, and then synthesise it in your own decision-making process and come up with what you believe is the right answer in that case.

Rosalind Croucher: Thanks, Annabelle.

Lyria spoke of the dangers of shortsighted regulation. Annabelle was speaking of dangers, particularly in terms of AI judges — in other words, relating it to the constitutional framing and delegation of judicial power. The idea of judicial power is something not amenable to replacement, because it lacks the human, because it also lacks that ability to discriminate in the way that judicial officers spend their entire life training for, and all of the legal profession. But the shortsighted, knee-jerk regulatory responses, the dangers that might lie there and building on the other observations of the panellists — these are Annabelle's dangers.

Lyria Bennett Moses: Victor, you might have more first-hand experience of this, but I think that the challenge for politicians is that people are often scared of technologies. AI is a really good example of that. It's very nice as a politician to be able to announce, "We are going to do something about this thing that you're scared of. We are going to regulate AI," et cetera. It's far harder to say, "we are going to clarify some of the challenges of interpretation of discrimination law." We're going to have to look at a whole range of acts, because as I said, it's not in one act, and we're going to really think about how to align those values with what society cares about so that it can apply — whether it's a human, a system, or a hybrid decision. It's a lot more work, and you don't get the punch of saying, "We're doing something about the technology you're scared about." And you have to do it *over and over* because people aren't just scared about the fairness of AI systems. They're worried about accountability. They're worried about a whole range of things. You have to keep doing that. It's hard. Getting things right, or the best possible answer, is more work and less reward. Whereas if people are scared of something, if something bad happens and you can say, "We have a law that is going to constrain this beast that you're worried about — this AI," it's attractive.

Rosalind Croucher: We have room for two questions, and then I'll end with just a couple of sentences allowed to me as the chair.

Q1: In a criminal system, if an AI system commits a crime, how is the crime judged? Is it "we are going with the accountable human" and judge the *mens rea* and *actus reus*? And will there be also a human dignity issue: who faces the crime? How complex is that scenario?

Annabelle Bennett: How long is a piece of string? Lyria might have an opinion on that too. First, judges are used to evaluating

evidence and the weight of evidence. It's one of the things we do all the time. If you had a human giving evidence against some sort of machine output, you would have to evaluate that. In the ordinary course you would do that. I don't know about an AI committing a crime. This gets down to the question of who's responsible. It's a bit like, can you get an AI, can the machine be the inventor? You'd have to look at the statute and you'd have to see whether it's a person who's required. Then ultimately you get to who wrote the algorithms and who marketed it. It depends on the particular crime you're talking about.

But these are clearly questions that eventually could well come before the courts. And I don't know what else Lyria or Melissa might add.

Lyria Bennett Moses: I can have a go at it. So first of all, statutes are drafted as statutes are drafted. At the moment, most of the kinds of crimes you'd think of are basically crimes committed by people. At various times in history, people used to convict animals. We don't do that anymore. Crimes of statutes are written about humans. Should we expand it to systems? I think that is an interesting question, but also a confusing question because systems are not motivated by the same kinds of things that motivate humans. If you're trying to stop humans murdering people, one good way to do it is to say you will spend a really long time in jail if you do that. I don't think that's necessarily motivating for systems as such, or what jails for systems might look like, but we can think differently about how to stop systems doing bad things. And we have to sort of ask those questions rather than just assume current law and just add "or robot" or something like that in it.

Melissa Perry: The question that you raise is a very interesting one because it leads into the question of who is to take responsibility for the outputs from machines. Is it going to be the person who uses the machine to produce the output, or is it going to be the programmer, or is it going to be the person who has responsibility for those who have responsibility for auditing the system? So there probably does need to be some clarity around those sorts of issues. They're very fundamental questions. Second, in the administrative government decision-making context, we do have a lot of provisions now in statutes which raise some rather interesting constitutional issues, but purport to not only authorise the use of machines in government decision making, the exercise of statutory powers, but also to deem any aspect of the decision making process undertaken by the machine to be made by an officer of the Commonwealth, so by a human decision maker. In that way, the statute itself has expressly dealt with *who* is to take responsibility for the output from the machine or any flaws in its processes.

Annabelle Bennett: So waving a magic wand and saying, "I hereby deem this machine to be human."

Melissa Perry: Yes, effectively, and a specific human.

Rosalind Croucher: Thank you. I think we'll leave it at that question. We began today in a world of, I think, hope — the possibilities of AI as an enormously useful tool, the world of hype that can transform complex decision making over large data sets into seconds, decisions that are assisted by analysis produced in seconds. So that's enormous areas of hope. Technology can certainly fix, it can enhance, and it can transform. Where laws and legal frameworks come in is that law

and the possibilities of legal frameworks and regulation can legitimise, and it can also hold to account so long as it is anchored in the world of people, and a world of humans and a world of humans with human rights. With that, I would like you to thank and warmly congratulate my wonderful panel.

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, pp. 83–96. ISSN 0035-9173/26/01083-14

2025 Royal Society of New South Wales and Learned Academies Forum: "AI: The Hope and the Hype"

Panel Session 2: AI in Communities[1]

Professor Sally Cripps, UTS
Jill Bennett, UNSW Sydney
Stephen Hardy, partner, Bain & Co.
Jacky Hodges, NSW Dept. of Education
Dan Hart, AI, University of Sydney

Sally Cripps

The next two sessions are going to be about AI in communities and AI in health. AI in communities is of particular interest to me. I spend a lot of my research doing actually graph theory, and using it to try and figure out how communities can use insights from graph theory to transform their own communities by better decisions under uncertainty. It is a real passion. One of the things about Australia is that, although we are a wealthy country, we are also a country of almost two parts, one of which is urban and the other rural. The reason this is such an issue in Australia is because we've got a big land mass which is very sparsely populated. Often what happens is that getting services such as education, health, employment into more remote communities or less affluent communities can be problematic.

This session is to explore how AI might address that as a sort of a general theme, but we are also mindful of what risks there are in AI. I've noticed that previous speakers have spoken about what it means to be human. To set the scene, I'll tell you a story before I introduce these wonderful panellists. I stood up here two years ago and I had a ChatGPT response to why porcelain — ground porcelain — is good in breast milk: it said because it had the right nutrients etc. Two years ago, this was an example of this terrible hallucination. When I came to do another talk about two months ago, I thought, I'll just ask ChatGPT, do they still agree with their statement?

So I asked, "Do you still stand by this statement about porcelain and breast milk?" And ChatGPT came back and it said, "Oh, let me apologise. I'm very sorry, but I was prone to hallucinate and I'm getting better now." And I thought, "Well, that's really nice." Then I also went to a conference in Western Australia where I heard a lot of people talking about AI as magic: it can now reason like a human. Given ChatGPT was inclined to be honest, I thought I would ask it about whether it can reason. It came back and said, "Oh no, please really. I don't reason. I work on statistical properties about which is the next word. It can sound like reasoning, but it's not human reasoning." Then I said to it, "Well, thank you for that frank and honest response." And it came back and it said, "No, thank you for your questions." And then I had to take a double take.

ChatGPT and I were each giving credit where credit was due, but I had to remind

1 This is a transcript of a presentation, available at https://www.youtube.com/watch?v=jKcJFgBYhcM&t=10s

myself, I'm a human and it isn't. I think that's my concern when I think about AI moving forward, particularly as I'm a grown person who actually knows a little bit about AI. So it's not too hard for me to disengage, but as a prelude into the education piece we're going to talk about, we actually looked at deploying AI tutors in rural areas to help with the education, and we had a lot of success. But one of the little boys was told after he'd been engaging with it that it was a robot, it was not a human. He was a lonely child who struggled to make friends; he took himself off to the bathroom for an hour and wouldn't come out because he was just so distraught that this thing that he thought was his friend was not his friend. There's a tension there. On the one hand, there's huge potential to really deliver services. On the other hand, I think we need to be mindful of the human element in there and think hard about what we should really be doing.

I'm now going to just briefly say what each one does: they're going to come up and give a slightly longer introduction to themselves. The way this session is going to work is I'm going to ask them each question and then I'm going to throw to you to ask questions.

We have the lovely Jill Bennett. Jill is a scientist professor and an Australian laureate fellow who's very interested in ageing and the mental wellbeing of the ageing population. Stephen Hardy has had 25 years' experience in rolling out AI systems. His last one was as the chief AI officer or in charge of the AI of Bain Consulting in the Asia-Pacific. Jacky Hodges is an executive director and the chief data officer at the Department of Education, and I've had the absolute joy of working with Jacky over the last three years on an AI-related project. Dan Hart was at the Department of Education, and now is in charge of AI and innovation at Sydney University, and was responsible (with Jacky) for rolling out EduChat.

Jill Bennett

Thanks, Sally. I lead a research program at UNSW, developing AI companionship specifically to address loneliness and psychosocial needs across a variety of communities, but, in particular, we try to look outside the mainstream. We run tea-and-technology events for older people in the community, particularly people living with dementia. This is specifically to redress the balance in youth-oriented technology markets, because we recognise that AI is pushing its way into aged care. We really need appropriate products. We work with a number of aged-care and home-care providers in our ageing program. With home care, we focus especially on regional, rural, remote communities, where often the providers report working with people who don't see anyone for two weeks, other than the care worker who comes maybe for an hour, an hour-and-a-half a fortnight. There's a real interest in integrating AI to expand what they would see as person-centred care, so not just automating.

We take a very bottom-up, community-based approach to all of our work: working with people in everyday settings; working from lived experience; and co-designing with people living with dementia. There are people in aged care using the AI companions via an iPad. See one of our first developed AI companions, Viv, on the YouTube video, who was developed with and for women with vascular dementia. She is now used in training, also as a peer-support companion.

What is AI companionship? It is essentially personalised chatbots that are driven by LLMs, but generally used for support and/or sustained relationships. In a sense, you define the personality and the attributes and the knowledge base of the character through an app. The market is currently dominated by so-called romantic companionship, which is a whole topic in itself. You might have seen there's a big focus on addressing loneliness, which, again, is far more complex than just perpetuating low-level chat.

Some of the problems that are widely reported — sycophancy and over-validation — are problems that OpenAI have been tried to address in the current ChatGPT-5 iteration. Essentially, tech companies want to prolong engagements for as many turns as possible. The best way to do that is to make chatbots as agreeable as possible, but then you run into problems when people are airing negative or even suicidal thoughts and the chatbot is primed to ultimately agree rather than to gently challenge. So that is quite a skilful process. Chatbots are also largely reactive. They don't have great capacity to hold in mind the personality and to engage in planning. This is something that our programs have tried to address.

Moreover, chatbots are very stereotypical: on screen, that is what you get. I put in that I was a fictional person over 65 who wanted a companion. You always get these fresh-faced, beautiful, idealised youths. Not to say that older people can't enjoy having a gorgeous youthful companion, but clearly there is something missing that we need to address.

There are some very viable cases that we've been testing. We have trials running across aged care at the moment, so chatbots are very good. They are patient and responsive, attentive, and validating in good ways. As mentioned, we focus on peer support. What that means is not so much expert support, but creating relatable characters who have a knowledge base and access to expertise, and a kind of everyday lived-experience filter. We do a lot of work developing strategies for emotion regulation, and you can imagine that when people are distressed, dis-regulated, confused, then often the best thing would be sitting calmly talking with someone for several hours; but care providers mostly can't provide that; families mostly can't provide that.

So we are thinking about how we can integrate chatbots into a care ecology rather than simply creating something separate. Our work on loneliness really focuses on understanding the context of loss, late life changes, struggles with diagnoses, and grief, which are not something that you solve, but something that has to be processed and worked through and discussed at length. We also, of course, focus on cultural specificity. We do quite a bit of work in forced migration. We have a Lebanese chatbot that we're testing at the moment. All of these conversational agents, we design and develop through lived experience consultation and a lot of co-design. It's very important to anchor in people's needs.

To talk about when a chatbot might actually be better than a human, I think there are some cases we can talk about. See one of our collaborators, Gwenda Darling, on the YouTube video. At the recent NSW government inquiry into loneliness, where we spoke about the importance and potential of AI, Gwenda said this: "In the middle of the night, when I'm pulling stuff out of the wardrobe, or getting distressed, an AI companion would be really helpful."

In the middle of the night, what does she need? She doesn't want to call her son. She wants to preserve that relationship. He can't speak to her for two hours at 2 AM. Also, she says she wants to vent very often, and she understands that you can talk to a chatbot and let off steam in a way that you actually wouldn't want to necessarily burden your family with. She said, "There's no way I'm ever going to ring a helpline. For me it would be great to have Zoe [her bot] when I want to talk to someone. I don't have to make it a cup of tea. I'm just totally in control."

We learn a lot from these extended engagements about the way that people adapt to technology and take up technology and make particular demands of it. That enables us then to design for the context. We found that in all sorts of instances, people are able to ask questions that, for whatever reason, they might not want to ask family and friends.

Gwenda has actually fashioned her conversational agent in her own image, which we don't always do. The important thing is that we need to not only build these things from the ground up, but draw on lived-experience data. We use this data to actually drive the design of AI modules. We do obviously test with psychotherapists. We use top-down data as well, but I think the concept of designing within people's everyday life context and using their skills and knowledge is super important. This is what really enables us to create products that are anchored in communities. What that means is not just extracting people's private data about their lived experience, but implies working with the way in which technology is incorporated — a sort of dynamic relationship to everyday living context.

I think that's a really important note to end on, and an important place to start the discussion of how we introduce technology that is capable of psychosocial support into care systems and into people's lives, so that we're not actually flooded with inappropriate foreign, off-the-shelf solutions. Thank you.

Sally Cripps: The other panellists are going to introduce themselves.

Stephen Hardy

As Sally mentioned, I initially spent the last 25 years building machine-learning systems, and then the last three years building AI systems to basically change the way people work in employment. I've been at Bain and Company running their AI delivery team across the Asia-Pacific, a team of about 160 people. We basically go into large corporates across every domain you can think of from — not government — but healthcare, retail, consumer electronics, mining, financial services — across the board. We look at how organisations can change the way that they use AI, from being a tool that helps a person doing a task towards a tool that allows a task to be done with appropriate human supervision. That's the key of the productivity gains that are inherent in AI. It's not that you give the person a tool, but you replace a task being done with a machine that does a task.

That obviously has profound implications for a couple of different things. We'll cover that later. That is, how to keep the human part of that is really been important to me over the last 25 years as this has developed. Looking forward to unpack that as we go.

Jacky Hodges

My name is Jacky Hodges. For another few days, I'm the chief data officer at the Department of Education before I move across into the Cabinet Office. My background was from the Australian Bureau of Statistics, using a lot of thinking about how you'd use AI machine learning in terms of data analytics. Coming into education, it has been very much around: how would you think about the use of data in different ways? One of the projects (working with Sally and also with Rebecca from Western Sydney University) has been thinking about how kids finish school well. What are the things that actually make kids finish school well? What are the contributions? Often through traditional statistical techniques, you only get so far. We've been exploring a whole lot of things there.

The other big project is around the NSW EduChat, which is like a chat facility that we've rolled out to schools. Last Monday it went live to students in our higher years and we're about to go live on Monday with the rest of the students. That's been very much in a pilot trial stage, thinking about how you prepare generations for the future where they're going to be using things and tools like this. I think we'll talk about the burning platform. Moreover, when you know that so many other tools are landing in their laps at home or even that other organisations are trying to get things into schools. How do we do that in a safe and considered way?

Dan Hart

My name's Dan Hart. I'm currently the general manager of automation and innovation at the University of Sydney. It's actually a super cool team. I've just joined the team. They've existed for six years now, and it's a team that anyone, any staff member in the university, can reach out to. If they are doing something that's repetitive or boring or they just wish they didn't have to do, they can come to our team and our team will design an automation to do that part of the job for them, if it's possible. Early in those six years, they were pretty much focusing on a type of technology called robotic process automation: that's where we get a computer to use the actual interfaces of systems that humans use. Usually, it's just simple copying and pasting information between different systems; there are a lot of our jobs that involve that kind of activity.

As AI got better and we introduced machine learning and generative AI, more of the solutions have involved artificial intelligence, either to make decisions about how we route different things or to provide written responses in some cases. What's really cool about the team is there's a free service provided by finance, not by technology. We don't call back like the cost savings that they get from implementing these solutions. The teams choose how to reinvest that time into the things that they think matter most to them. I think it's quite a nice way to introduce automation into an organisation.

Prior to that, I was two years as the head of artificial intelligence at the NSW Department of Education. It was my team that created NSW EduChat. I'm going to talk a lot more about NSW EduChat in the panel discussion, but I thought I should just explain what it is for those who haven't seen it.

NSW EduChat for teachers is like an assistant. It's very similar to using ChatGPT, but the main difference is that it's hosted entirely inside the Department's data cen-

tres and tenants, which means they can use it for use cases that they're not permitted to use ChatGPT for, like using student names, for instance, or asking for help writing an email where you're including personal information. The other way it's different for staff is that it also incorporates school policies and information about the curriculum. That means that when they ask a question with a pretty basic prompt about either a syllabus, a piece of information, or a specific detail in a policy, it's less likely to hallucinate. If you do that kind of thing with ChatGPT, it's likely to hallucinate because there are many thousands of schools across the world and it doesn't know, unless you've specifically named it, which one you're talking about.

It makes introducing generative AI to teachers a bit easier. For students, it's actually a completely different experience from ChatGPT. It looks the same, but it behaves very differently. As a student, if you were to say, "Hey, can you write me my essay about a fellow?" then it will say, "Oh, that's great. Let's write this essay together." Then it would do that thing that teachers do: asking questions back and forth, really getting you to think about the content yourself and to provide most of the content yourself. Before you know it, you've written it between you and the AI. The other way it's different is that it's got a ton of safety filtering to restrict it to certain topics, to stop it from going into areas that we don't want students talking about, and also make it very clear that it is an AI and not a person as well.

Sally Cripps: I'm going to do the questions in reverse order. First, there are two questions that relate to EduChat, to Jacky and Dan. If each of you could take a different part of a question, then that would be great. The NSW Department of Education is actually not known for its alacrity or agility. And yet, Jacky, you or somebody managed to get them to roll out what is literally a world first with EduChat. The question to you is: what did that take? We've heard Victor talk about Government 2.0 and Government 3.0, and how difficult it is to get governments over the line.

Then, relatedly, Dan, you are in charge of the technical rollout. If you can talk about how you did that, are there any case studies that stand out to you? Two questions on EduChat.

Jacky Hodges: Going to the bigger question of how do you get a massive department like Education to actually adopt and take on the use of AI? When the rest of the world has ChatGPT rolling out and at its fingers and there's a burning platform to say, "Oh my goodness, we don't know the impact that this will have on the students in a classroom," and "How do we actually prepare the teachers to engage in something like that so that they can also be advising future generations on the way they should or shouldn't be engaging with this type of technology?," I think that then brings a bit of focus to say, "Are we going to just ban this type of technology in schools and just say, sorry, you can't use it at all, or what are we going to actually enable it?"

I think the enablement comes with my previous colleague here who was able to develop something in a safe, innovative and controlled environment and then go, "Ta-da, you know what? We're actually ready to go. We could actually do this ourselves." He can talk about the program there. But it wasn't just like I was one piece of a puzzle to make this happen. My other role is co-chief AI officer, but that came about with the need

to think about how you do these things in a safe and ethical way.

My team's been very focused on how you assess different products going to schools with that safe and ethical perspective in mind. A whole lot of work is happening across NSW government as well as around the assurance frameworks and other things where the impetus was for us to do it right. The risk when you're thinking about children is exponential. You're actually putting into their hands something that you have to consider in a really risk-controlled way or think about how you actually use this as a platform to prepare them for the future.

In terms of the network to make it happen, we had a very multidisciplinary team: across our teaching and learning space, our technology side, and the data safety side. Initially, we also had a much bigger group, and then we had different groups to actually make it happen — Dan will talk about that — but we had the Teachers Federation who were very keen on making sure that anything that went into school was done in a safe way. It was done too with the support of the teacher and the humans in the loop. Having that backing meant that we could take more control over what vendors were doing in schools. There has been a real cultural shift. A lot of vendors used to put whatever they liked into different schools, while the schools made choices. Now, the schools have direction from the system: this is new.

It can be quite exciting. It's quite powerful. Having all of those stakeholders involved in that journey is really important. We learned at the same time as Dan made the technology happen.

Dan Hart: It wasn't just me. We knew that some private schools were already starting to provide ChatGPT, so we had this urgency of not wanting to get behind. I was new in the role too. I was in there a month and my director had been managing data science projects for a while, but none of them had really stuck. He gave me the approval to dream big and make something and then just go around and sell it. That's how the first version of NSW EduChat was created.

It was actually called ChatEdu the first time. That's probably why I get confused with the name sometimes. One thing we're not very good at is naming. All of the names we picked for this product at different stages were cumbersome to say; we've now shortened NSW EduChat to just EduChat, it's easier to say.

We made our first version of it in August of 2023. That was pretty late, eight or nine months after ChatGPT came out. What we really wanted to show to people was that EduChat didn't have to look like ChatGPT. It didn't have to be an answer provider: we could make it enter into this dialogue of questions for students. That could be the foundation of a tutoring-style piece of software.

The first version was very basic. It couldn't actually do multi-turn conversations. Every time I demonstrated it, I would sneakily start a new conversation or purposely ask questions that were fake so that you didn't see that it wasn't working — it was really like the definition of a proof of concept. That's quite rare in government because normally we start a project with a list of requirements and then we build for a long period of time and then we go ta-da. That's not always possible when you don't have requirements like in our case. Having that early version was great because we always

take it out and show people that AI could be different.

Over the next six months, we made much better versions. We designed testing suites to test the accuracy of the responses and the safety of the responses to students. At the same time, we incorporated knowledge bases for staff, so the staff version was more compelling for teachers to use.

We entered six-month negotiations with the units. In February 2024, we were ready to trial it in 16 schools. We rolled it out to 16 schools — both students about age 10 and up, and staff. It was a great success. I don't think I've ever done anything as an insider in a technology department that has been as great a success. I think the reason why is because the education team and my team were not separated by organisational distance. Previously, you literally had to go all the way up to the secretary of the education department, and then come down to the other team. We completely bypassed that by physically co-locating the teams and spending lots of time talking together. So, we made a product that was actually useful for teachers rather than getting requirements and then delivering something. The fast, iterative process was the secret. The governance committee was meeting weekly. We had the autonomy to make decisions, and that really helped massively speed up the process.

The trial was a great success, involving 16 schools. Around April 2024, there was great demand and we had spare compute capacity, so we expanded the trial to 50 schools, both staff and students. That allowed our team to get great feedback on how the tool was being used. It got so popular with those 50 schools that we had no choice but to roll it out to all 2000 schools between August and October, just for staff to begin with.

And then this week, starting rolling out to more students. By next week, it will be available to half a million teachers and students across NSW, which makes it probably one of the biggest deployments of generative AI systemically to schools worldwide.

I've got time for an anecdote. Part of working with schools closely continued after we delivered the project. Some of my team went out to the schools, to find good practice, record it, and then turn it into videos that we could show to other schools. Because in technology, teachers sell to teachers, and they never pay any attention to what techie people say because we're not talking the right language. We realised early on that the way to get higher usage was for teachers to tell their stories. We went out to a rural school in Hannon Vale, which is, I believe, on the mid-north coast of NSW. It's a small school, so it has about 25-ish students. In a small school like that, it's not possible to have separate classes for each year: they lump years together, which means for the teacher, it is a great deal of work to make and teach content for lots of different year groups in the same class.

There is this technique they call differentiation, where they take an activity and change it slightly for the capability of students. When you've got multiple years in the class, doing so creates a difficult teaching environment. This story is about a girl called Myra, who had just moved to that school from a metro school. When she arrived, the teacher noticed that her literacy levels weren't up to the level that they normally are. There was a creative writing activity in the class, so she differentiated it down to a level for Myra to take part. The class had to

write about their favourite place at home. Myra wanted to write about her mango tree, but she wasn't confident enough to start writing straight away. The teacher said, "Why don't you draw a picture that describes the story?" (this is a typical way to differentiate that class). So Myra drew the sketch.

The next part of that differentiation is usually the teacher sits down and says, "Okay, now write some bullet points that describe the key parts of that story" to eke a bit more writing out of Myra. So she did that. But then unfortunately, the teacher had a class of other students that she needed to provide coaching to. She couldn't spend that entire class with Myra, so what she did was to set Myra up with EduChat and those bullet points, and then she said to EduChat, "Myra's writing a story about her favourite place at home. These are the bullet points she's come up with so far." Then Myra continued that conversation with EduChat as it talked her through ways of adding more content into the story that she was writing. The outcome of that was amazing. We were there hearing about it from over a catch-up and they said by the end of the day, Myra had been able to write a whole page about that story. So we were like, "This is incredible. This definitely is showing it's working."

We sent the video crew out to record. And while we were there, she said, "No, this is even better than you thought." Because what had happened was Myra went home, was so excited about being able to write a story that she continued doing it in her own time, and she wrote a three-page essay on the mango tree in her home. What was really interesting about her writing was not just the length of it — her handwriting quality had improved, the spacing between the words had got better, she was using more punctuation. The thing was littered with adjectives because EduChat was really helping add more descriptive content to it. It was good because it wasn't just about teaching Myra to write more, but it was also getting a real fuse back into writing as well.

Sally Cripps: That was a super cool story. I mean, it was amazing. Congratulations. It's just a phenomenal achievement. I do hope you write it up, the two of you about how you did it and motivate it with that wonderful story.

Stephen: work and how AI is changing work practices. You talked about keeping the human in the loop, just observing the task, which sounds okay, but then does that actually work? What is this change in work practices doing and how is it playing out in communities?

Stephen Hardy: First, the technology is moving very fast. If we look back to December 2024, the very first example of a commercial release of a language model that could approximate reasoning could actually give you a conversation that looked like it was thinking and reasoning. It wasn't actually thinking and reasoning, but it was trained to mimic reasoning. It's not just the next word now: it's actually been trained to mimic human reasoning, maths, computer code, in many different ways. That's one set of changes that has been happening. The technology is getting better and better. The AI we use today will be the worst AI we ever use, it's important to remember that trend.

The second change in work practices is social uptake. For every company that we've talked to, it's clear that all of the people who work in that company or their kids are using AI at home. That isn't true for all communities across Australia, so that's part

of the divide that we've got to think about, what the impact is going to be of these technologies. Pretty much all the people in the companies we work with are using the technologies. They know that they'd be useful in their work, so they want to use them in their work. In almost every case, you'd find examples of people using open ChatGPT to do work tasks that shouldn't be done because the data's possibly moving overseas. The incredible security risk around not providing tooling to people inside organisations is that they're going to bring their own tools because it really does help their day-to-day work. These are two of the things that are driving the adoption of AI in a workforce.

The third thing is the promise of productivity gains, which essentially is going to increase profit or more work. Which of these actually occurs says how much job displacement is there going to be, versus how much new work is going to be created. This is a tricky task to unravel. It's like a global experiment we're doing at the moment. Stanford recently put out a paper saying that there's likely to be the first signs of impact of AI on hiring of junior people across the US economy. In the younger group, just the rate of employment in senior ranks went up 13% and the rate of junior employment went down 15% on AI-exposed tasks.[2] It's a really good bit of work, but it's not causal, it's correlational. It doesn't show cause and effect, but it's suggesting that something's happening now.

Radiology was mentioned earlier today. Radiologists had a great improvement in productivity from 2000 to 2010 when all radiology services were digitised. Rather than using film, they started to use screens and digital technology, which made the task of radiologists much faster. But, that did not cause any job displacement for radiologists because the latent demand for radiology was huge. There are twice as many radiologists employed as there were before, who get paid basically twice as much. The question then is around different types of work: which ones are going to be actually displaced; which ones are actually going to enable more work to be done.

Software engineering is something that people are basically predicting the demise of. I think there's a great deal of latent demand for software. If you think about companies, there are always thousands of things they'd like to do, but don't have the resources to do. If you make the software engineering process three times more efficient, I think that means you get three times as much software.

What that means is that across the economy, across businesses, the impact of AI is going to be very jagged. Call centre operators are going to have a hard time finding work: how do you retrain them? Sometimes those workers, if you think globally, are often in less wealthy countries. Those sorts of jobs are typically going to dry up. They'll be fewer, at least. Then the question will be: how do we retrain these people? What is the replacement work that they will be able to do? There is an obligation for us to think through this before we actually change the parameters of our economy. So that's some of my thoughts on that matter.

2 Brynjolfsson E, Chandar B and Chen R (2025) Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence. Stanford Digital Economy Lab. https://digitaleconomy.stanford.edu/app/uploads/2025/11/CanariesintheCoalMine_Nov25.pdf.

Sally Cripps: Excellent. I think it's wonderful. If we want to really scale things up so that it reaches more people, it inevitably leads to systemisation because we've got limited budgets.

To Jill: what do you think the problems are or what might the problems be if somebody were to take what you do and ask, "Look, there's such a demand for it. We're going to scale it." What would your response be and what do you think a few of the risks might be?

Jill Bennett: This is why we need to actually do the co-design at the outset and not just import systems. This is the essence of our work in these areas like dementia support, where nobody has really worked with people with dementia. What are the conversations that they want to have? I think we can generate a lived-experience database, but we do have to be careful and diligent around the use of this kind of qualitative data and how it's translated.

What happens at the moment is LLMs are driven by a whole lot of work with psychologists, but usually the kind of mental health psychology, like cognitive behavioural therapy, that is more procedural and more aligned with computational logic. Nothing wrong with that, but I think going into these areas that need to be looked at through a psycho-social lens. We need to understand the dynamics and the qualitative nature of the relationship. That work absolutely has to occur.

The same goes for the cultural and place context. This is the essence of the work in regional Australia: what do people want? Understanding the points of resistance is also really instructive. You just can't deliver new technology to someone who is 70 with mild cognitive impairment or early-stage dementia, living on their own. There's going to be a high rate of attrition. Understanding how to design into that context is an important discipline in itself.

The big picture: we're talking about how disciplines change when AI is added. What does design become? What does psycho-social support become? This is also my answer when people say, "but a human psychologist is better." We are talking about context where a human psychologist isn't available; in Australia (and indeed everywhere), the vast majority of people who might be deemed to be in need or desire of mental health support cannot access it or afford it. This is the problem of resourcing and delivering services into regional Australia: there aren't enough care providers, so you're never going to get that. We are actually talking about something that doesn't exist. We might want to restrict it to the sort of sub-clinical level. Obviously, there are some complex cases we would not want to leave to chatbots, but we know that loneliness increases the risk of dementia by 50%. If you are not talking, if you don't have that mental exercise capacity to offload and work through simple everyday worries, then there is a health impact that's not being addressed. Thinking about the newness of this, what we don't have, and then most importantly, thinking about the equity issues and how we design for these very particular living contexts across Australia, is the way to go.

Sally Cripps: Thanks, Jill. I'm now going to open it up to the audience for questions.

Q1: Gabrielle Upton. I used to be in the state government. I am really interested in the AI product into schools and for teachers and how embedded is that in the curriculum, because the example you used was something that was a way to develop

skills in a challenging situation in a remote or rural school. How much of it is actually embedded?

Jacky Hodges: What we've tried to do is to build it into our curriculum reform and the way that the actual tool has been trained with the new curriculum. In the rolling out and training of teachers and working in schools — in particular the pilot schools — it has been for teachers to think about how they use this kind of tool as part of their teaching of the curriculum. How successful that has been is probably dependent on the innovation of the teachers and how comfortable they are. When we think about equity, the reason for having something like EduChat is because — in the same reason you talk about in the health field — we have something to provide across schools, as well as the training and the support for the teachers to feel comfortable engaging and using it.

I think there's more to be done around how you build that into the curriculum itself. Having something where we could pilot and then roll it out to all students, that will result in almost the whole population using ChatGPT or others. It will bring that innovation and other ideas back into the system on what else needs to change to actually make that successful. Yes, I think it's still a work in progress.

Dan Hart: I'll say from a technical point of view, the curriculum integration that we did in — I think it was January 2024 — is now over a year old. And Stephen's saying, AI is advancing so fast that probably it's time for a revisit of how we've done that to integrate it a bit deeper. The curriculum's also changing over a period of years. That in itself is an activity that the team can keep up with. We work very closely with the curriculum teams inside Education and they get early releases of EduChat. So they get to ask, is it behaving in the same way that they expect? Are we addressing any of the issues that they previously raised? They also come up with prompts that are specific for the subjects that they're covering to help teachers integrate AI into practice a bit more easily. I would say it's definitely work in progress.

One of the good things about owning your own tool, though, is that you also have access to the data. While teachers are interacting and students interacting, we're collecting all of those conversations. We have — in some instances, where there's been low-performing responses on a subject — gone through those de-identified conversations to work out: where is it not working well, and what extra data do we need to give it to strengthen the answers?

Jacky Hodges: The data scientists are providing that feedback back to the curriculum people too, to say, what else do they need to do?

Q2: I'm an AI professor. We have lots of teachers who need training and so on in AI. We are so happy as university professors to help you to organise the courses and to help you to train teachers with AI technologies and how to use AI. Contact the Australia AI Institute at UTS.

Q3: Branko S at UNSW. I want to ask a question. Can EduChat provide the users a companion for the students as well? If it can, do you have any feedback system whereby you do post-processing, backend processing whereby you find students that clearly are developing cognitive or emotional or mental problems and can you then feedback to try to identify those students to help them?

Dan Hart: This might be a very disappointing answer to that question: we don't allow like personal advice-type questions to occur and that's enabled through several technical mechanisms. One of them also has an alerting feature. If we do detect something is not working and it has provided what we term advice, then someone will look into it to work out — first of all, does the student need assistance or anything, but then also to change it so that it doesn't do that. We're making it very clear this is an AI tutor to help you with education.

Jacky Hodges: We strengthened that recently to refer back to the importance of the teacher or a parent or someone else to try and bring that human back into the loop. There's a lot that we've done around well-being questions with the actual release on Monday and then following on to really test and say, "Well, if they put this type of thing in, what response would they get?" That's a developing area.

Dan Hart: Can I share a bit about technically why it's difficult because I'm super passionate about this? The way we do this is by having another AI checking the AI. As the response is building up and streaming back, we take the partly streamed response in its building up form and send it to another AI and say, "Is this okay?" So we're going, "Is this okay? Is this okay? Is it okay?" — all the way through that streamed response that's coming back. To do that in a way is still fun for the user, we have to use a smaller bot. Unfortunately, it's not as clever as the big AI that's making the responses. That means sometimes it doesn't really understand context. It's always a balance between filtering to make it safe, but then not filtering so much that we block educational-related topics. That's particularly difficult with history and English literature.

Sally Cripps: I'm going to two more questions, but they have to not be about EduChat. It's fascinating, but I want to hear some questions from the other panellists.

Q4: I'm Helen Salmon, director of the British Council here. I had a question for Jill. I'm fascinated by the companions. When you're working with patients with dementia or short-term memory loss, can you talk us through the ethical considerations that you had around someone potentially not being able to retain the information that they're working with in AI over the long term, and how that works with consent and some of those conversations?

Jill Bennett: Yes. we work all the time with people with dementia. It is really a question of prioritising their subjective experience rather than measuring them against the norm to see if they meet our inclusion criteria. The question is, how do we design for them? The initial work that we did was with different groups of people with different kinds of dementia where they really shaped the characters to accommodate their perceptual world. In every case, other than with very advanced dementia, people always had an understanding of the AI status of the agent, whether they called it a robot or whatever. I think that's a very important ethical consideration. They will have fun, they will play with it, they will potentially even talk nonsense with it from time to time, but understanding the nature of the thing that they're playing with is important.

One of the ethical limitations that we observe is we get a number of people asking us if we can develop a clone of them because they can't visit their mother or father: wouldn't it be great if my parents

just had a version of me, the AI version of me? We say, "Well, no, it might actually be nicer if they could have a relatable sort of friend to whom they could talk about you and missing you and what they want." That, I think, rather than feeding a kind of hallucination, is important. In general, I would say we adjust to where they are rather than measure against a norm.

Q5: My question is an extension of the question that was just asked something about EduChat. I chair a large regional and rural aged-care provider. One of the things with the new Aged Care Act coming in and this real emphasis on government on ageing in place and ageing at home rather than being in homes, tools like this are going to be absolutely vital. The thing we're wrestling with is when someone in the middle of the night is having that conversation with their AI bot and feeling great, but if they're not feeling so great, what's the feedback loop? Or is there a feedback loop where AI is saying there's a real problem here, there needs to be a connection to a human, there needs to be some really, really big help. Obviously, there are privacy issues around that.

Jill Bennett: Yes, absolutely. I think we have that conversation with different partners and there's a different answer in every case. The counterbalance to that is that older people with dementia like everyone else often use the chatbot to make disclosures that they might not make to other people, so there is also the counterbalancing privacy issue. Gwenda says, "No, I want to scream and shout and vent at it." I wouldn't be doing that if it was a real person. What is the trigger? There certainly can be a feedback mechanism. At the moment, we are involved in designing systems that can accommodate a feedback mechanism. This involves paying attention to the interface and the accessibility and what is actually viable in the middle of the night when you're already dysregulated. You can't have a three-stage login.

What would be the agreed protocol with people like Gwenda, who has actually raised this issue directly? As far as possible, we would have that conversation with family, care provider and so on. I think that that's scalable as well. There could be a bespoke system with an agreed protocol and even individual sign off for all of the people using that system.

Sally Cripps: I'd like to thank the panellists once again. I think you can tell that this AI is here to stay, and it's becoming part of our jobs, of our education, of our healthcare, and it's happening for real people at a very grassroots level. Thank you very much, everybody.

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, pp. 97–110. ISSN 0035-9173/26/01097-14

2025 Royal Society of New South Wales and Learned Academies Forum: "AI: The Hope and the Hype"

Panel Session 3: AI in Health[1]

Michael Blumenstein, Flinders University
Amandeep Hansra, GP and Australian Digital Health Agency
Simon Carlile, Google Research (Australia)
Richard Taggart, eHealth NSW
Sean Baldwin, NSW Health

Michael Blumenstein

I am the chair of what is likely to be one of the most controversial and exciting sessions of the day: the AI and health session.

What you might have noticed was a really positive vibe around AI leading up to this session. I've been researching AI for 25 years. And we're in the summer, as Pascal said, but I've seen the winter of AI. We're maturing in this country: when we talk about AI a decade ago, the then chief scientist, Alan Finkel, was quoted in saying that AI would be a threat to society, would bring down everything.[2] When you listen to predictions like that, back then AI was very different to what it is now. That that posture really created the cloud that has sat over Australia for a long time in relation to how it reacts to AI.

In an AI arms race across the globe, we need to be a contributor to the innovation, not posturing, waiting, and perhaps not doing anything. The danger of not doing anything will probably outweigh our worries and fears. It is really positive to see the change that has occurred: how people are reacting and responding to the absolute explosion of AI in every part of our lives. People don't know that AI is in their washing machines, in the brakes of their cars. It has been for decades because AI has been different over different periods of time. The future of X, as Pascal said this morning, is AI plus X. This is the reality: nothing is moving forward without AI.

It's about how we respond to it and how we embrace it, rather than worrying what might be because someone isn't treated as well as they could be in terms of their healthcare, or whatever. We've got to look at the risks for sure; we've got to look at governance; we've got to look at regulation, but we really need to update and evolve our posture, particularly in Australia. When you look at AI plus X — in this case, AI plus health — this is a really big issue, not just for the state, but also for the country and globally.

Recently, UTS commissioned a report a white paper with PwC,[3] where we basically did a survey of hundreds of clinicians,

1 This is a transcript of a presentation, available at https://www.youtube.com/watch?v=PYKhfcc0B7s&t=8s

2 Finkel A (2018) Chief Scientist warns AI could create "parallel society." https://www.chiefscientist.gov.au/2018/07/article-chief-scientist-warns-ai-could-create-parallel-society [Ed.]

3 "Reinventing healthcare. Unlocking the power of AI in Australia's health and wellbeing sector," UTS PriceWaterhouseCoopers, 2025. https://www.pwc.com.au/health/health-matters/reinventing-healthcare.html [Ed.]

healthcare workers, end users across Australia. Interestingly, what came out of it was there was apprehension, but there were also really interesting views from among the community around what the promises could be. What is really important about these discussions today is that we have experts who are sitting on our panels who are some of the best in their fields. They are also people who have crossed over between clinician and technologist, even expert in technology. It's really important, whenever you do AI plus X, that you have a really good understanding of what the X is, because there's no point in talking about AI if it's only from the perspective of a technologist or a computer scientist. It has to be the perspective of the end user, the actual adopter, who is going to make things work.

Today we have a panel that bridges that gap perfectly; everyone here has a background that is going to contribute to a very interesting discussion. I'd like to jump straight into a panel discussion with questions because I would like to leave at the end time for people to ask questions from the audience.

So without further ado, I'd like to introduce each of the panel members. Dr. Sean Baldwin is an emergency physician, Deputy Chief Medical Information Officer for North Sydney LHD at NSW Health. Dr. Simon Carlile is the senior staff research scientist at Google Research Australia, but is also in his own right an esteemed academic entrepreneur and will add a very interesting perspective to the conversation. Dr. Amandeep Hansra is a general practitioner, but also a chief clinical advisor in medicine at the Australian Digital Health Agency and has a significant background in bridging the gap between digital and health. Richard Taggett is the Chief Executive eHealth NSW and Chief Information Officer of NSW Health. Richard will bring a very exciting perspective, as will Sean, around the NSW context.

This morning we had an AI and health panel and there was also (in the Law panel) a lot of attention around regulation and what that means. I would like to get this panel's perspective and start with Amandeep around the issue of balancing regulation and ensuring that is done in a way that doesn't stifle innovation, but with particular reference to health and healthcare.

Amandeep Hansra

I've had a really interesting career journey and I've found myself at the intersection of exactly that problem. In my role with the Australian Digital Health Agency, even though we're not a regulatory body, we are responsible for the national digital health infrastructure. Our stakeholders are all the states and jurisdictions, as well as the Commonwealth government. We need to make sure that the foundations are right. It feels like we're building an amazing house, but I need to make sure that I'm laying the concrete properly so that all the systems are connected, and we're speaking the same language, that we can actually then build and innovate on top of that. But I also work as a GP. So I'm on the frontline seeing patients and seeing my peers start to adopt technology.

Another role is I work with startups. I work with a lot of very early-stage companies with founders who come up with great ideas of wanting to solve big problems and who are particularly attracted to healthcare. We help connect them with clinicians and then help fund them through the doctors'

network. We get great innovation happening, but the problem is that you can't innovate and scale that innovation if you don't have the foundations right. The foundations are not just the infrastructure that we're building at the agency — around connecting all of our systems, making sure we speak the same language, making sure that we can identify patients, clinicians, everyone across the whole system — but we are also making sure that we do it with a safety lens. When things go wrong in health, people die.

I'm not saying that any other industry is not important, but for us as clinicians, our main job is to do no harm. We have seen innovation move at rapid speed, with the regulatory bodies trying to play catch up. The problem is that if you do these things without proper regulation, you can end up with people coming to harm. A lot of the people building AI tools are very well intended, but exactly to your point, Michael, they haven't actually got the health-specific knowledge or they're not connected. They're technologists with great ideas — and I love when they come into the healthcare sector and try and solve problems for us — but they need to bring us in on the journey. So there needs to be a lot of co-design because we understand what the stakes are. We have some great work happening with the Therapeutic Goods Administration (TGA) which has obviously built some regulation around software as a medical device.

If AI has a therapeutic or a diagnostic role, then it must obviously be regulated. The problem is that we're seeing a lot of AI that is in a grey zone. For example, in my practice, I use an AI scribe, where I talk to the patient and then I don't have to look at the screen and type notes. I can directly share eye contact and look at the patient, do my assessment, and the AI scribe has documented what's happened in that consultation. The problem is that sometimes when I look at the notes, they haven't directly translated what I've said, but have made some interpretation and said, "You've listed these symptoms. I think your brain is saying this is what you think the patient has." And that starts to get into diagnostics but these devices are not regulated.

We've got no clear line: what is safe and what is regulated. Many of our regulators are grappling with how they play in that area, but without stifling that innovation of those people who are coming into the sector, really well intended and wanting to solve the big problems that we have. It's a bit of a balancing act, but the key is making sure that you have clinicians involved in the design of these projects, of these innovations.

Simon Carlile

I've got a question for you. One of the things that really impressed me about the AI and Law panel was that I came away with the notion that in some ways treating AI is different from everything else. It's not really the right lens. There are plenty of examples of where — if we're talking about discrimination or harm and those sorts of things — you can just strike AI out and use the legislation that we've got or rather insert AI in. There are still obvious complications around how you assign responsibility and those sorts of things. It brought to mind the late Daniel Dennett (1942–2024), the philosopher and neuroscientist,[4] who was asked about deep

4 Dennett DC (2023) The problem with counterfeit people: Companies using AI to generate fake people are committing an immoral act of vandalism, and should be held liable. *The Atlantic*, 16 May. [Ed.]

fakes; he said, "But we have the mechanisms to deal with this." And when probed, he said, "Well, look, legitimate currency is only legitimate because there's no way of reproducing it. And people who attempt to reproduce it are hunted down and they're prosecuted." It's an example of where we have a mechanism in place; it just needs to be repurposed. I'm wondering whether the regulatory things that you guys face in your own practice — because I'm a neuroscientist by training rather than a clinician — might be translated more rapidly than having these large open debates that go on interminably?

Amandeep Hansra: It's a great question. When we're building clinical governance frameworks, it's the principles that matter, not so much the technology. If we spend five years writing clinical governance frameworks or accreditation guidelines, by the time we've written them, new technologies emerge and we have to rewrite them. We can't write them about specific technology. My issue is: if you think about a drug that goes through a regulatory process with the TGA, they assess the drug, the ingredients, the manufacturing process, and then say, "This is safe for patients to use." The drug that they've assessed is the same drug that goes out to the public. The problem is when you assess AI through a regulatory process, the AI is constantly changing. So once it's out there, it's not the same as what was assessed.

Sometimes that's where we've fallen into traps: where there's been an update or they've been dependent on open AI models that have changed, and then you realise that there's a drop in quality or there's a bad outcome and there are no real post-market surveillance processes in place. So we have to develop some new mechanisms because it is unique, I think, to our traditional device and drug regulation.

Richard Taggart

I would love to build on your question and comment on Amandeep's. I'm really proud to work in the NSW health system alongside some 180,000 people. We provide care in over a thousand locations, from the city to the bush. If you're a consumer coming into one of those services or you bring your family member into one of those services, you expect empathetic care that is kind, appropriate, evidence-based, safe. You expect that you get the best-quality care that you can. The way that we do that, the way that we ensure that we can be as good as we can be — and we're not perfect — is through the governance, the regulation, the policies, the procedures, the structures, the frameworks, and the foundations that Amandeep just mentioned. AI has to coexist with some of these existing well-established foundations and build upon them and enhance them where it makes sense.

In fact, AI itself has to be a regulatory tool. I think that's the opportunity: you wouldn't expect as a patient that one of our clinicians would provide you a treatment that didn't have evidence or didn't have some level of security or safety around it. AI needs to have the same lens on it.

Michael Blumenstein: That is a really important point you raise about the evidence, Richard. When we're talking about health, we're talking about people's lives. What's at stake is significant. It's different from someone just browsing on Netflix and getting the wrong recommendation by AI; the damage isn't huge apart from being bored to tears for two hours. But we have a big issue if the AI is wrong and the impact is huge.

I'm going to touch on evidence in a second because I think that's a really important part of the discussion. Sean, what do you feel are the misconceptions that people may have in relation to AI-enabled healthcare that sometimes bring fears and concerns up in individuals who are pursuing the healthcare system with AI-enabled processes?

Sean Baldwin

Great question. Where I work in emergency departments, our patients come to us and they're unwell, so they're worried already. They walk in the front door, they're worried about what's wrong with them, about what's going to go wrong with their health. And will their clinician appropriately diagnose them? Will they treat them? Will the treatment have side effects? Is that going to cause them harm as well? And now they've got all these extra worries. They're worried about where their health data is going. Where's it stored? Is it secure? Is it going to be part of the next data leak reported in the media and who's got access to it and who can learn all about their private data? So they're worried on top of that and they're thinking the doctor might make a mistake.

What if the AI makes a mistake? What if it hallucinates during my presentation to the emergency department and now it's decided I need a brain biopsy? That's what they're really worried about. So your patient's worried, they're scared, they're unwell, and trying to explain to them, "Oh, but I'm going to use this AI tool." In that situation, having them understand what that AI tool does, how it integrates into the healthcare system, and how it affects them personally, is really difficult in a short time.

The AI scribes are a great example. All clinicians love them because they return us to the bedside. They reduce our administrative burden, they give us more time to deliver healthcare and do the jobs that we love. But when this makes it into the media, all the commentary that you see back is patients saying, "Oh, I don't want the AI making decisions for me. I want my doctor to make those decisions for me." But that's not what the AI scribes are doing at all, although, as you say, the boundary is blurry.

Simon Carlile: This is something that really gets me going. I think this is a case of poor cognitive hygiene. The late Danny Kahneman (1934–2024) got the Nobel Prize in 2002 for, amongst other things, identifying what a range of what he call cognitive biases, which are ways that we think about things that are in daily life actually very effective. One of the things that he identified was that the ease with which you can bring something to mind equals your spontaneous assessment of the probability of that happening. That's when I stopped watching CNN because I renamed it the "Crisis News Network." This is about taking the narrative and making sure the sorts of things that we think are dreadfully important are really what are front and centre. That's what gets repetitive. Poor cognitive hygiene; we need to wash our hands of CNN, for sure.

Michael Blumenstein: I would say that news media and other sources of information have inflamed things. In the case of AI, there's a crisis around the corner or that there is a potential for real concern. And yet, when I talk to people like the esteemed people on this panel, I'm actually quite confident that there are a lot of things in place — whether it's governance or whether it's the ethical requirements in health and medicine.

I want to go back to the point of evidence. Would you both, Richard and Sean, and

maybe Amandeep as well, see a base emerging of evidence that there is improvement to be had when you use AI in the sector? And could you extend to the data that's used? Do you have some examples or would you like to comment on that?

Richard Taggart: I should caveat and say that I'm a proud public servant and work for NSW government. One of the things that the NSW government has done in their wisdom — because this is an emerging area — it has added new levels of bureaucracy, and nominated me as the person in NSW Health who has to fulfil that level of bureaucracy on behalf of the system, which means I often have to stand in front of clinicians like my good colleague Dr. Sean here and say, "Hang on just a second, show me the evidence of what you're about to do before we allow use in the system." That means that I've got a good view of what's going on out there across NSW Health. There are many hundreds of really quite exciting cases, but I would have to say the quality of the evidence that's showing the ones with real meaningful clinical impact is still at the beginning of the journey.

There is some really good solid evidence where things have actually gone through to fully regulated devices — there are some really great radiology use cases, for example — but the vast majority are still in that early phases. The projects I'm most excited about are where NSW Health clinician researchers are contributing to that body of evidence. But the opposite side of that coin is it isn't all just about healthcare and clinical outcomes. If I look at what's actually going into production, more than 50% of these use cases are in a category that I would call increasing the time for care. So, taking away the admin burden — we spend almost a quarter of our time in the health system just on administration, keeping that thing running — anything that we can do to reduce that percentage of time — every minute that we give back to a clinician to provide frontline services — is a benefit.

There is good evidence that the time for care benefits for AI is pretty strong. There's some strong evidence in some of the administrative tasks, the finance functions and submitting forms and getting the quality of the submissions back to the various government agencies. That's why you're seeing very rapid adoption of those perhaps less attractive, but equally as important, use cases.

Sean Baldwin: I can definitely empathise with these discussions. I have them on a smaller scale — just within a health district rather than throughout the whole state, but all speaking the same language. I think that's the real problem I see with AI. Any treatment that I'm going to give to a patient is going to be supported by high-quality scientific evidence, but the high-quality scientific evidence for most of the AI programs that we see is lacking. There are a couple, but they're few and far between. When I see a large-scale randomised multi-centre trial evaluated using AI and evaluated not using AI and see that there is a patient-oriented outcome difference — which is morbidity or mortality to do with that treatment — then I'll adopt it and I'll use it. I think we'll see medicine shift that way as the evidence arrives.

I have no doubt that it's coming. The problem's going to be keeping pace with all the evidence, but I'm probably going to use AI to keep me up to speed.

Amandeep Hansra: I think that we also need to look again at the evidence in this new evolving world. Historically, we have relied

on the multi-year randomised control trials. The problem is that, for example, from the moment when you identify the actual target of a drug to going into production, it's something crazy like 17 years. We can't do that with AI because by the time we do some of these trials and we publish them, the world has moved on and we're already onto ChatGPT 1500. We're moving so fast. I'm not saying that we shouldn't have evidence. I think it's just that we need to look again as an industry at what is acceptable evidence in this world where we are moving at such rapid speed. Are we comfortable with that? How could we use AI to help accelerate how we're generating some of that evidence?

The other thing that I'm finding is that we're getting a lot more real-world evidence. When you think about the evidence that we have from some of the trials — for example, we haven't had women represented in a lot of research trials over the years. It has only been more recently that we've really pushed to have gender equity in terms of representation. The way that I was taught to pick up a patient having a heart attack was based on mostly men's symptoms but women present differently.

We haven't always had cultural diversity, First Nations representation. But what we're finding, with the world of AI and the way we're capturing data and how we are able to take some of the meaningful insights from that data, we're actually finding more representation of our communities in the data. So we're actually able to see what that drug or that intervention looks like in the real world because I'm actually now checking what it's doing to my heart rate. I'm checking what it's doing to my sleep. I'm getting a lot more patients who are coming back to me with actual data of how that intervention affects them individually because they've got access to technology that they didn't have before.

To me, that is actually much more important than what randomised control trials have shown in some ways because it's personalised care. It's about that individual and how it's impacting them — not a whole group that wasn't representative of the patient in front of me. I'm probably going to get shot down by a lot of my research colleagues who say, "What? How dare you say this?" but I do think we need to stop, and take a moment and ask, "What are we able to accept as evidence moving forward?" Health is the second largest industry in terms of publications in AI outside of technology. So we're doing great in publishing a lot of things about AI trials, but the problem is we're not also scaling those. We've done the trial, we've got a bit of evidence, but then how do we take the next step and actually put it into the real world and say, "Are we feeling comfortable enough to trial this on a state basis and a country basis and a large population?"

Richard Taggart: I think that's just an Australian problem. We're very good at producing the evidence, but very bad at actually adopting it.

Michael Blumenstein: I really think that the perspectives here are so exciting to hear, given the fact that there's a representation of someone who can oversee what's happening in NSW Health around technology and may making sure that it's going to be for the betterment of patients in healthcare. I really want to emphasise what you said, Richard, about the backend, the non-direct-patient-affecting AI processes that are occurring. The single largest expenditure in the NSW budget is Health. So there are opportunities around, and there is actually a lot of

advantages outside of just direct health and patient care around advancing the operations of facilities, of hospitals, etc.

The other thing is Amandeep's point around the question of what we are looking at in terms of evidence. What is the data telling us? And now there are so many places for data to come from. I want to ask Simon: how can we scale up what we're doing? What are your thoughts around the role of universities, the tech industry, healthcare providers, and governments to contribute to that scaling? What else is there that brings together important institutions to drive this innovation, that can contribute to the future?

Simon Carlile: I'm very happy to speak to that. I wouldn't want to lose the previous track, though, so we can come back to that because I think there are some real changes afoot that change the focus from illness to wellness. That is going to happen. AI is going to play a big role there, so I'm happy to chat about that later. But, okay, returning to universities. I've been a university professor. I've worked in the area of hearing health and researched in that area for most of my life. I've had spinoff companies and I'm now working for a large technology company, so I've got a little bit of visibility on each of those things. I think that universities in particular — and it was wonderful to hear about EduChat or ChatEdu or whatever they decided it is today, because that is an example of where we need to drive continual pedagogic reform in the way that we approach education, and no more so than in universities.

I had the pleasure of working with a number of brilliant clinicians and educators in the '90s where we designed a medical course that was based on problem-based learning. The questions that we drew out were, how do you do differential diagnosis? How do you know what the questions are that you need to ask? Education should be about the process of learning and discovery, not the regurgitation of somebody else's facts, because by the time you get out to practice, they're irrelevant anyway, because the world has moved on to pedagogic reform that really focuses on the way that students and professionals can leverage AI when they have left the university.

In terms of university research, they're making an effort in this regard, but they're still very discipline-bound. As you would know — I was a pro vice chancellor as well — there are the baronies and the deans are the barons and there's a lot of fragmentation or silo-ification in universities. And while there are attempts to bring together much more collaborative in that direction, I think that can go a lot further. There are also Commonwealth incentives that could be put into that — not just dropping a little bit more into the National Health and Medical Research Council Fund for a so-called collaborative research project. I think the most salient thing that struck me in the years that I've been in technology is they've got to use a lot more user-centred design, particularly engineering companies.

I'm a scientist and a time-zone engineer as well, and I fall in love with my own solution, but is it the right solution? Is it actually moving the needle on the thing that counts? Anything that's involving the end user is super critical to make sure that your effort isn't being wasted. I think tech companies have got a lot to learn there.

I think health providers as well. I just came back from seven years in the United States, and that was an interesting experi-

ence for a whole variety of reasons, trying to navigate the health system there. What really brought into very sharp focus for me was the way one needs to really balance patient-centred care and — obviously in the US health system — a profit motive. And it's got to be win-win, not zero-sum. There really has to be thought about that. There are increases of efficiency but that shouldn't just hit the bottom line. How do we turn that back into increasing and improving patient care? I'm quite sure that while we're not purely profit-driven, we are budget-constrained. Many of the same arguments pertain at state and commonwealth levels here in this country. So I think these are still very pertinent arguments.

Coming back to regulation, it was raised this morning by an ex-minister. Slow is a problem. About three weeks ago, I had an internal briefing in my company about hallucination — we don't use the word "hallucination" anymore, not because we don't like the word, we call it factuality, because there is a huge amount of research that is going on, on determining whether a statement from a large language model or any kind of foundation model has a factual basis. We heard one example: for instance, we ask another LLM, is that true? Designs now are moving towards looking at this in a parallel process, so you don't wait for the answer and then ask someone else, "Is that answer a good answer?" It's all built into the process. So factuality is one thing.

The other thing that was raised today is that you don't understand what's going on inside this black box. There's no way of looking at the logic. Well, reasoning is also a hugely, hugely influential area in foundation model design and implementation now. They're not there yet, but to your point about the way things move so quickly — and that's another real difficulty as well — is at what point do we start making judgements on the data that we have. Because we know that I'm not going to wait for the next super-fast CPU to come out before I buy another computer. I need a new computer now, and so I've got to make some risk assessment of where I am. They are just some quick thoughts on the four areas, universities, technology, health providers, and government.

Amandeep Hansra: Simon, I just wanted to comment on one of the points that you raised. I'm going to mention the word "productivity," sorry. One of the things that happened when we were being sold AI scribes is that the tech companies would come and pitch to us and say, "You can see this many more patients if you use an AI scribe." Obviously that will translate to billings for us in general practice. They thought that's what sold us and that's what made us pay the quite steep prices initially. The thing is that we didn't see more patients. We love the scribes, but we didn't fill that spare time with seeing more patients. We actually just spent more time with our patients. Sometimes we measure the wrong things — when you couldn't spit out the data and say, "This is how much money you're making extra or how much extra productivity." What happened was that the doctors went home happy, the patients went home happy, and the doctors wanted to come back to work the next day.

That is a measure of productivity, though, because we have a lot of doctors and other healthcare providers in the system who get burnt out and leave. And we've got about a 10% attrition every year of medical doctors who leave because they just don't enjoy

their work. Sometimes we think that we're addressing one thing, but, if you actually look at it, yes, there has been a productivity benefit, but it's not where we thought it was going to be. It's reducing that what they call in the US "pajama time," having people be happier at work and actually provide a better service to their patients and ultimately have better patient outcomes, which also leads back to productivity. I think we need to think more broadly about the benefits of AI rather than this direct return on investment (ROI).

Richard Taggart: Amandeep, you just touched on something that's at the top of my problem list at the moment, which is AI scribes. And if you're not familiar with the term "scribes" — effectively you allow the AI to listen to your consultation, then it writes up the note in a format that you want. So you can put it straight in your clinical information system and do various tasks with it. Actually they're becoming more of an agent now than they were. They're doing other tasks and there's a lot of very rapid adoption. This technology has been adopted by general practice, by consultants in their private rooms, by allied health professionals all over the world. The demand has been sensational. It has been really quite exciting to see that, but the implementation of a piece of software or a new tool or a workflow change inside a general practice, even a large practice, compared to implementing something in an enterprise as complex as NSW Health, are two very different scale problems.

There has literally not been one single day since I started my role on the 20th of January that I haven't had a request for a scribe of some flavour to be rolled out at scale in the NSW health system. I'm very eager to do it because I can hear demand from clinicians a potential benefit, but it has many layers of complexity in the scaling of that technology and in implementing it safely. I'm sure this is something that comes up in your world as well, Sean.

Sean Baldwin: Absolutely. I'll send you an email later.

Michael Blumenstein: You guys just sit next to each other, pass a note, no AI. I would have liked to keep this going because you probably agree with me that this is a very interesting area: AI and health, because health touches everyone, no matter what age you are, no matter what size of family you have or how old your parents are.

The last question I'm going to ask (and it's terrible for me to do this). For the sake of productivity, I'm going to ask each panel member to give just a one-sentence response to this: someone mentioned the fear in the context of jobs with the rise of AI, in the context of health and the health sector. What do you each see will happen to jobs in the future with the adoption of AI?

Richard Taggart: Augmentation is the word I would use. We are not going to be able to meet the demand for the health system over the next 5 to 10 years. We've got a rapidly ageing population, which is wonderful and that we're living longer and hopefully healthier lives, but that increases the demand on the system — there aren't enough workers whom we can bring through the system, there aren't enough hospitals that we can build to satisfy that demand. We have to embrace AI so that we can get some level of productivity — to use that word — to meet the demands so people get high-quality care at that scale.

Sean Baldwin: Completely agree. I guess we work for the same organisation. I don't know anyone in the public health sector who's worried about losing their job to AI. I think if anything, we need to get more done with what we've got, the resources we have, and hopefully AI can help us do that.

Amandeep Hansra: Yes, I definitely don't think I'm losing my job anytime soon. I think I'm actually going to go back to providing care as opposed to doing lots of admin tasks. I'd love for AI to do all the stuff in the background that needs to be done so I can just sit there and focus on the patient in front of me, provide the human touch, provide the care coordination, be their partner in their journey through the healthcare system. I think AI is going to give me the opportunity to do that. So I'm hoping we can all fall back in love with our jobs.

Simon Carlile: I think that the biggest impact of AI is going to have nothing to do with jobs. I think it's going to have to do with this shift from illness to wellness, as I mentioned before. Everyone who's wearing a digital watch probably has a range of physiological sensors on that watch. You've probably also ticked a box when you put it on, which said, yes, I want to use the Fitbit notification of whether my heart rate's going up or down or blah, blah, blah, blah. So you've just given away all your data. (Well, not given it away because there are very strict privacy rules around that as well.)

The agents that are being developed by the companies that make those watches are looking at that data are not providing you a medical diagnosis, but they are providing you advice. They're also providing you indicators when you should go and seek some medical diagnosis. What they're building is these large models that can look at your data and determine when it is starting to deviate from your standard norm and starting to indicate that you do need to talk to a medic about how you manage your behaviour or your diet, or whatever that's driving you out of the norm.

That's not going to impact anybody's job other than maybe decrease the load of chronic disease that is hitting our medical systems, increasing the rate at which people are able to provide diagnoses and treatment, because early treatment is always a much better way than waiting till it's stage four. So, yes, that'll be the biggest impact. It's happening now and there's no government regulation standing in the way and people are happily signing over their data and taking the advice of those agents.

Sean Baldwin: Certainly I see someone every day whose Fitbit or Apple Watch has told them to come to the emergency department because they might be dying. It's extremely rare that I find much wrong with these people. It does happen. Most of the time it's a measurement error or something's gone wrong, but it has made them so, again, anxious and stressed about their health because it's such a constant thing that they're getting exposed to every day. They're going, "Oh my goodness, what could this be?" And then they come in to be reassured by me.

Michael Blumenstein: I think that the formal panel session has ended in a very positive note, before we go to Q&A to the audience. I think we have time for a few questions.

Q1, Susan Pond: My question is to Simon, but it may overflow to the others. If you're at an industry conference — not one that's a more for-profit-driven conference — and you are a part of the research team now in Google, what are the big data companies

(if I can call Google that) thinking about in terms of making profits from AI in healthcare?

Simon Carlile: That's a really interesting question. If I knew the answer, I'd probably be a vice president. My own research is actually applying machine learning to re-imagining and computational modelling to re-imagining the hearing aid. I'm a neuroscientist and a hearing scientist by training, so that's unfortunately what my focus is. So I'm not sure that I can give you a cogent answer from an industry perspective.

Q2, Mark Pesci: There's a new device that's just come out from a company called Kohler in America which manufactures most of the toilets that Americans use. It's called the Dekoda.[5] It's a camera that points down into your toilet tank. So when you've peed or had a number two, it will do an analysis. It uses edge AI and uploads it to the cloud. There's now going to be a galaxy of devices like this that will be flooding all of our medical systems with AI data around patients' wellness. How are we going to be able to incorporate that into our medical systems workflows?

Richard Taggart: I might try and tackle that question, although I'm fascinated by the experience. You must hold your breath when you use that restroom. I think there's a lot of opportunity here and a lot of upside. One of the things that I think about as a leader in the public system is how could we possibly keep large numbers of people away from seeing the lovely Sean and potentially have a regular appointment with the lovely Dr. Amandeep instead? There's one risk factor that I think would really move the needle on that scale — we could use AI devices and technology — and that is hypertension. About 30% of people have hypertension, so I divide this room into three. Maybe you look particularly healthy, maybe you've got hypertension, and about half of you would not know. But if we could tell the half that doesn't know that actually you might be at risk, go and get your blood pressure checked by a health professional, we could potentially prevent a whole heap of downstream health issues. That's a relatively simple use case, but if we can scale that, I'm very happy.

Amandeep Hansra: I went to that health conference — I think it was in Vegas — and I saw that toilet. I was too scared to sit on it. But I think we're in this period where we've got too much data. People are being given all of this extra information and they don't know what to do with it. The clinicians don't know what to do with it either. I get patients coming in all the time with, "Here is my output of all my wearables, my Aura Ring, my watch data, spreadsheets of information." People who are not diabetics are wearing continuous glucose monitors and saying, "Hey, I went for a walk and my sugar spiked. What's wrong with me? I must have diabetes." I'm like, "That's normal." So the part that's missing is actually the actionable insights and getting those accurately.

We're going into this prevention space where we're going to arm people with their own data; with that, we have to educate them. We have to tell them to help them figure out what is something to be alarmed about that needs to send you to the ED or send you to your GP. I think people should be allowed to manage themselves. That is what I'm hoping as we start to give them access to their own data. One of the things that happened a couple of weeks ago

5 See https://www.kohlerhealth.com/dekoda/ [Ed.]

is we just changed the legislation for the MyHealth record — I hope everyone in this room has a MyHealth record — where we've now implemented "share by default," which means when you go and get a blood test, you should be able to access that pathology test result the same time as a doctor accesses it. If you get your blood done in the morning, by the afternoon you should be able to see that report yourself. A lot of people panicked about this and said, "Oh my God, people looking at their own blood test results, God forbid they're going to get panicked and they're going to end up in the ED." I thought this was a very paternalistic view — only I can hold your data, even though it's about you, only I'm allowed to look at it and I'm only going to give it to you when you sit down with me and I can bill Medicare and that we are going to work through it together. In some situations if it's complex, I agree, but I think we need to move away from being the custodians of all of our patients' information and giving it back to them, and educating them in the process.

We're going to go through a period where there's going to be a bit more utilisation of the health system before we get to this period where people are going to start to take responsibility, accountability, and learn about their own body in ways that they've probably never known before. So I'm excited about that toilet — a little bit scared, but also excited.

Q3: How important is context to diagnosis? At times you're saying that all the data being collected is a bit too much or it spooks the patient, etc., but how important is it to the clinician to get the right amount of data? I did a little bit of research on snoring. Why does somebody snore and what kinds of snoring are there? I'm just wondering, we found it important to have something on the person to collect all the snoring patterns and the condition of the environment. That was quite a useful tool for the clinician to decide there is a problem not with the person, but in the way he was living. Same thing happens in aged-care centres, the CO_2 levels just go way up because all the doors are closed.

Amandeep Hansra: Look, I think that is part of our role in this world, where we're going to have a lot of data, where I say, "I'm not going to go out of a job." It's because that is some of the nuance that I'm going to have to figure out as a clinician what the context is. When was that information collected? You had a terrible sleep, but I knew that you'd been delivered some bad news that day at work. Maybe the AI doesn't know that, and maybe that is my job, understanding the human context, also understanding what's happening in your social environment with your family, your kids. This is why I think that we as clinicians are going to go back to doing what we're doing better, because we're not going to be gathering objective information that can be gathered by AI and wearables and systems; instead, we're going to be gathering the information that is more about the human and the context.

Going to your point about the environment, that is also static information that can be measured. We're now starting to look at smart homes, particularly for aged care, where I've seen people develop mats where you stand on the mat every day to brush your teeth and over time it can tell whether you're starting to develop osteoarthritis because you're starting to not have equal balance on both feet. We're going to start to see these environmental data collectors that

are going to feed into what the clinician gets to see in terms of your home living conditions, your context. How often do you open the fridge? How often do you move in the house? The problem is that at the moment we're not set up to receive all that data and interpret it. At some point, I would love it if the AI systems could pull that together and give me a really succinct summary of "this is what I think is going on," but then contextualise it with the human factors that AI can't figure out. It's going to be a fine balance, I think, getting there.

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, pp. 111–126. ISSN 0035-9173/26/010111-16

2025 Royal Society of New South Wales and Learned Academies Forum: "AI: The Hope and the Hype"

Panel Session 4: AI in Practice[1]

Mary-Anne Williams, UNSW Sydney
Ros Harvey, Australian Conservation Foundation
Dan Jermyn, Commonwealth Bank of Australia
Salah Sukkarieh, University of Sydney

Mary-Anne Williams

This panel is the Academy of Technological Sciences & Engineering (ATSE) panel. We are celebrating 50 years. Our panel will focus on AI in practice, and we have three extraordinary panellists to share their actionable insights into using AI in practice from robotics all the way to financial engineering, in a good way.[2]

Ros Harvey[3] was the founder of The Yield[4] and really focussed on how to help farmers make faster, better decisions. Dan Jermyn[5] is actually a closet physicist, just to give him a bit of credibility in this crowd. He's the chief decision officer at CBA — you might know that CBA was in the top five last year, and now the top three AI leaders globally. That's pretty extraordinary. Salah Sukkarieh[6] is our token roboticist. Does amazing stuff, also related to what Ros does in agronomy. The way this panel's going to run is that each panellist has five minutes and no slides, to introduce themselves and their perspective and point of view on AI. Then we'll settle into a few panel questions and will have 10 minutes of Q&A. Please start thinking about your questions early on because we want to fill that 10 minutes. So, Ros, you'd like to go first.

Ros Harvey

A quick bit of background about me and The Yield and AI. The Yield was a technology company I founded in 2014, focussing in agriculture. Basically, it was a very early AI business. We were focussing on microclimate weather predictions to use in agriculture to combine with agronomy knowledge. Later, we also grew it into doing microclimate predictions that farmers could use to make better decisions about how to grow their crops from irrigation, when to fertilise, when to harvest. Then we grew the product more into predicting yield as well. We did two things. I sold the business last year to Yamaha Agriculture, and at the same time they bought a robotics company. Salah and I have worked together for many years as well, so that we could stitch together the data and close the loop. So we're predicting

1 This is a transcript of a presentation, available at https://www.youtube.com/watch?v=jP_Rp6Aypno&t=20s

2 See https://research.unsw.edu.au/people/professor-mary-anne-williams for Mary-Anne. [Ed.]

3 See https://www.acf.org.au/about/our-organisation/board-council/acf-board [Ed.]

4 See https://www.theyield.com [Ed.]

5 See https://h2o.ai/ai-100/winners/dan-jermyn/ [Ed.]

6 See https://profiles.sydney.edu.au/salah.sukkarieh [Ed.]

the sort of things that you should do to grow crops more efficiently, and then predict the yield.

If you can add the data that robots actually see and what they do and you can automate, then you can close the loop with data, which allows you to create more powerful AI models. We were a very early innovator out of Tasmania. The Yield now holds the patents in Australia, US, Europe, Singapore — major markets for using AI to predict micro climates. It's a very foundational underlying technology. The patent application date was 2016, a long time ago.

I'm now president of the Australian Conservation Foundation, which is what I'm doing post-The Yield. I advise Yamaha strategically. I also have the right to take the underlying technology into other verticals, which I'm experimenting with.

I was asked what the big challenges of this were in the real world, of which there are very many, particularly in agriculture, which is a challenging sector, as Salah and I both know. I think the really headline thing I'd say is, look, CSIRO says right now, today, 80% of AI projects fail.[7] Some other folks say 95% fail. Why? People don't know what they're trying to achieve. What is the actual impact commercially? They don't have the data. I work in quantitative AI, not in general language models, and getting the data in the right format structure — so that it's actually AI-ready — is an enormous challenge. I don't care which sector it is from: finance, which is very digitally matured, through to aviation, through to agriculture, which is probably the least digitally enabled. It's very, very difficult.

The other challenge is the human side. A lot of organisations are just not ready. And when I say not ready, not only is it data, it's their people. At the end of the day, in my experience, the technology's the easy bit. The hard bit is the organisational and people bit. There's a lot of enthusiasm out there for AI, and there will be a lot of money spent, and there will be a lot of failed projects if we are not very, very careful about how we focus its on solving real problems with real people, solutions that are then capable of being adopted.

The final thing I want to say is we also need to be very mindful about the cost and the environmental impact. This comes from bitter experience in terms of the cost, because, at the moment, the industry is being really driven by the major vendors who have no interest whatsoever in reducing cost. Effectively what this means is that they're encouraging architectures and data-management strategies that are not optimising for cost-effectiveness.

I like to say that you can choose three things: you can choose flexibility, you can choose performance, or you can reduce cost when you're actually engineering. But what happens is the vendors don't care about cost. Flexibility is great because it drives up the compute and the amount of data that you're using, but it's very bad for customers and very bad for startup businesses.

It's also very bad for the environment because we know that every one of those ChatGPT queries draws two watts of power and half a litre of water in data-centre costs. This is going to have potentially a very significant impact on our energy use and our

7 Hajkowicz S (2025) Evaluating and prioritising artificial intelligence projects: A guide for better decision making and investment outcomes. CSIRO, June. https://www.csiro.au/en/research/technology-space/ai/How-to-choose-and-invest-in-the-right-AI-projects [Ed.]

net-zero transition. AEMO (the Australian Energy Market Operator) is predicting that up to 80% of our data use could be taken up with energy from data centres, which is a really big issue.[8]

These things are solvable because technology is human-mediated and we get to choose. We can choose: the way we architect, the way we design, the way we use technology that both solves for our community. I'm a techno optimist, I love data, I love technology, I love AI — but I think we've got to be very mindful in the way that we do it in the environmental, social, and economic benefits and impacts.

Mary-Anne Williams: Dan, there's a lot to dig in from what Ros was saying. I'm sure you're all realising that she is a sort of volcano ready to erupt.

Dan Jermyn

I'm from the Commonwealth Bank. We care a lot about responsible AI and trust and accuracy. I have to correct a few things. We're fourth, not third, and it's just in financial services.[9] I'm also not a closet physicist — I'm a failed physicist, but it's a shame, I was quite good for a while there. I'll talk about finance broadly. I am a representative of the Commonwealth Bank as well, and I was reminded of that, as I walked in casually, I didn't know where this place was. I didn't realise it was quite so grand. I feel terribly underdressed. Also, dangerously dressed, because if you can squint closely, this does say powered by CommBank AI. As soon as the first person who spoke to me realised I work for CommBank, they wanted to know a lot about a lot of things that I do not know an awful lot about.

I think it's important to understand one fundamental principle in AI: it is evolving so quickly. It's very, very complicated and it's very multifaceted. I've been working in it now for a very long time. I would say there is still a huge amount of unknowns about almost every aspect of it.

Notwithstanding that, maybe it's worth explaining a little bit about who I am, where I've come from, why I'm standing here today, and what I want to talk about. As you can probably tell from the exotic accent, I'm from old South Wales. I grew up in Cardiff, and I'm from a family where nobody had ever been to university. I did physics at university. I thought I'll do the hardest subject at the hardest college I could get into. I was doing quantum physics, and then I did a PhD, which I subsequently abandoned in computing to do a startup, a technology startup. That went quite well. We sold that.

I ended up in a bank — the Royal Bank of Scotland — and subsequently here in CommBank, where I've been for eight and a half years. Very pleased to say I'm a citizen of Australia now. Growing up in Danwood, never in a million years did I think, "Well, I'll end up in a bank." Actually, whilst I was being a failed physicist at university, I did an internship at an investment bank. Not for me ever again, that's not a place I want to spend any time. And yet, as I look back now where I am now, I feel an unbelievable

8 AEMO's updated forecasting methodology targets rapidly growing electricity loads, following industry consultation. 6 Aug 2025. https://www.aemo.com.au/newsroom/news-updates/aemos-updated-forecasting-methodology-targets-rapidly-growing-electricity-loads [Ed.]

9 See https://www.commbank.com.au/articles/newsroom/2025/10/commbank-among-best-banks-ai-maturity1.html [Ed.]

sense of pride and privilege that I got to work at CBA.

I think it's important to recognise that a bank's not necessarily the most obvious place to do AI, certainly not to do it quickly. It's not a glamorous place, it's not always a popular place (as I discovered with that first conversation), but it is an important place. Almost everybody in this room has a bank account — many of you have had mortgages, loans; you'll have insurance policies and so forth. It's a vital part of the fabric of society and it's vital that we encourage people who care about that society and that community to be a part of the conversation on AI.

When I joined eight and a half years ago, our CEO, Matt Comyn, who was then the retail CEO, had this idea of the customer engagement engine. The idea was that we would serve our customers well. We would just try and be better for them. It's apocryphally one of the most brief business cases in history. The stories have grown in terms of myth, but it's written on the back of a napkin: just do good stuff for customers and trust me.

But it's kind of true. I was never asked to do anything other than do good things for customers. And we thought about our multi-channel approach. We have customers in branch and a call centre, huge numbers of customers coming to us in digital. How do you stay relevant? How can you be helpful? It turned out that AI was a really critical part of that story because it allowed us to try to apply meaningful interactions at a scale that simply wouldn't be possible without that kind of technology. Over the course of that, we've strayed into all sorts of really complicated aspects of AI and how you get it to scale. We've made mistakes, we've made breakthroughs, but we've done some incredible things that simply would not have been possible without applying AI with a purpose-driven workforce. We've returned $1.3 billion to customers through our benefits. We've dramatically increased the speed at which we can respond to emerging bushfire activities. AI plays a critical role in all of these things.

Moreover, the way in which we're able to defend our customers from the rising scourge of fraud and scams is absolutely critically underpinned by AI. The bad actors are using it. We have to be in a place where our communities can be defended. I'm really proud that for all the difficulties that we have to work in this industry we succeed, and I want to maybe explore some more of that within the Q&A.

Mary-Anne Williams: Brilliant, Dan. It's great to have a leader like CBA in that industry, leading. And we've got little pockets everywhere, but we need to scale and franchise very fast. So it's good to have CBA leading the way.

Salah Sukkarieh

My name is Salah Sukkarieh from the University of Sydney. My background's in an area called field robotics, which is the application of robotics in outdoor environments. Just by the nature of being outdoors, the environment, biology, geology, and chemistry all change. The processes that we have to go through in developing these systems to work outdoors 24/7 in all weather conditions in different terrains is quite complex. Coming from the university, I generally focus on where we can push the boundaries on many of the algorithms and the sensors and the systems that we have to build. In the context of what we're talking about here, if we have to look at it from that

lens, we think about it from a physical AI perspective. So how does AI work outdoors? How does AI work and adapt and change as the environment changes? There are lots of challenges around there.

My particular background expertise — or the area that I like to focus on — is environmental robotics. It's really in that extreme natural environment space that I work. We've done a lot of work in industry engagement with the use of robotics and AI and so forth with mining. Recently, we finished a project with the NSW Rural Fire Service on using AI and robotics in first response into bushfire scenarios,[10] and a lot of work in agriculture, both grains and wheat and horticulture and exploring those different domains; more recently, interplanetary systems as well. So rovers on planets and satellite systems, etc. So it's a wide range of activity.

When you look at robotics 20 years ago, the focus was: how do you take a platform that already exists and is operating, then add some computers on board, add some sensing on board, and make it work. You treated each one of those subsystems as individual units: the navigation system will tell you something; the control system will do something; and the sensing system will perceive something out in the environment.

Where robotics is going now, it's a lot more embedded. There's a lot more of that activity happening between the different subsystems. We're starting to realise that the way you move or control the robot affects what you sense and how you sense, and that becomes important. That quality of information then dictates how well you're going to control the robot as well, so there's that underlying principle. I think one of the key, really novel areas that's being pushed at the moment now is in agriculture. The focus there is on biology.

In one particular project, we're looking at phenotyping. In phenotyping, you're looking at millimetre accuracy changes day in, day out, on a platform that's bumping and moving around and slipping because of mud, so there's quite a lot of complexity around there.

In another project with the cattle industry —the grazing livestock industry — it's a much more complex scenario, where you're dealing with changes in pasture, but also moving animals. What we're trying to do with this system is couple the farmer's response and the agronomist's response into that. So you get a much more complex picture of, not just hard data and AI systems that can work with data, but also soft data — like large language models and how they interact with each other in that form.

The challenges that we're seeing as we go forward are lots and lots of data coming through, processing that information, trying to do it in real time — starting to ask questions as to whether or not you keep pushing down the AI path or whether there are other spreads. For example, optimal sampling in the environment may get you better returns than worrying about large collection of data and how that works within that operational context. These are some of the questions that we're asking.

Mary-Anne Williams: I'm sure that ending on that point — it's not lost on the audience — the connection between the importance of collecting the right data and

10 See https://www.nsw.gov.au/ministerial-releases/innovative-technology-supporting-faster-smarter-bushfire-response [Ed.]

the smart toilet that we were talking about in the last session, right? Now consider the toilet with arms and legs and robotic capability.

We've got to try and find some common ground here. We've got a few threads happening, but across all of these domains — robotics, digital agronomy and finance — where has AI really delivered value? If you could also identify one or two blockers that you had to get over. Where's the value and what are the blockers?

Dan Jermyn: There are lots of places where we've certainly added value. I've been at CommBank for eight and a half years, been doing AI for most of that time. Lots of different ways that we've spoken publicly about what we're doing with AI for our customers. If you fast-forward to today, one of the things that we've talked about recently that's really powerful is in terms of fraud and scams. I alluded to it earlier. Customer losses to fraud and scams for CBA — I think I'm right in saying —from a high in 2022 have been reduced by 76%. It's a whopping amount of money. There's still a lot of fraud and scams out there. We can't be complacent. We're dramatically reducing impacts on customers.

One of the things that we're most passionate about —actually one of the most interesting use cases I've seen recently — uses generative AI. A lot of the AI that we started with was predictive. Effectively the bank — all banks — are really largely predicated on prediction. You try and work out if you should loan to a customer and you try and work out the way you should set your price points and so forth. Even fraud detection is looking at anomalous behaviour and trying to work out if it's dodgy anomalous behaviour or just someone's doing something a little unusual.

In the generative AI use case —it's a great story for Australia — we partner with a local startup called Apate[11] to use synthetic agents for scam baiting effectively. You've all heard scam calls, people purporting to be from a large online retailer, let's say, and saying, "Somebody's been using your account" and trying to get you into giving them your details. What we've been using is 10,000 synthetic, fake CBA customers who now, using generative AI, can very believably in real time speak to a real scammer and keep them on the line for a long time. They think they're speaking to an actual CommBank customer that they're just about to defraud. Actually, they're wasting their time. The record is, I think, 52 minutes so far. It's very good. We're always trying to beat that. And in one simple sense, every second that guy spends on the phone to a non-real customer is the second that he's not spending on a real customer, really trying to dupe somebody.

Much more powerful is the intelligence that we gather about how these criminal gangs are organising themselves, how the scams evolve over time. This is something that's dramatically increased our ability to be reactive. That's an example of how we're able to use to today's technology, really advanced technology, predicated on a long history of building the muscle to build AI safely in a heavily regulated environment in a large corporate entity that wouldn't traditionally be thought of as a place that could do very rapidly adaptive technology.

11 See https://www.apate.ai [Ed.]

I think you mentioned hurdles. I would just say very briefly that none of this is supposed to be easy. We'd like it to be, but corporate transformations are always hard. I think there was nothing that we couldn't overcome, but there was nothing that wasn't challenging. The challenges persist today. You can't be complacent about it. The space still evolves. If I think about what did we do that was really important, it was around trust, governance, risk, transparency, and thinking long term about how AI would evolve and did we have the structures and the procedures and the tooling in place to give us confidence that we could do that at our scale.

Mary-Anne Williams: That's really great. As I listen to you, I'm thinking, okay, are you already developing the strategy for when the bad actors have their own synthetic AI attacks?

Dan Jermyn: That's already here. And the domains like cyber and fraud and scams are really at the front line of what's happening with AI.

Mary-Anne Williams: Ros, what do you think?

Ros Harvey: Trying to sell is actually a very good indicator of what actually has value. I originally came out of a university; there's a world of difference between a proof of concept and the brutal reality of actually running a commercial business, of having to hit revenue targets. With some confidence, I can tell you in agriculture, the single biggest value is reducing labour costs. That is the holy grail that everyone is looking for, and they're looking for Salah's robots. The problem is there are very few of them, particularly in horticulture. It's different in broad acres, but we both worked in horticulture. The technology is behind the demand.

The part of the business I came from — which is basically quantitative AI – there are two ways you look at it: something that reduces cost, and something that creates value. I'll use a very concrete example: yield prediction. Why does yield prediction matter? If you're running a big operation, say wine — and we had Treasury as one of our biggest customers, a multi-billion dollar listed company. We only worked with large companies because they're the only ones who had the data. The little farmers did not have the data for us to be able to access. So the value is massive. It's documented: thousands and thousands of dollars per hectare, because, if you don't get it to the winery at the right time, it goes off. You've got to pick it at the right time. It's this massive logistics problem of getting the right grey wine at the right time into the winery.

It's a massive labour cost. You've got to move people around to pick it and get it in. And then you've got to also deal with longer-term forecasting around what your blends are going to be to mix the market. Knowing up to 12 months in advance what the actual yield is going to be using AI at the block level, which we do very well, is really a lot of value. The only thing I could ever sell it on was, I was going to get rid of the cost of sampling, because what they do at the moment is they wander around and they pick, and they actually hold a frame out in the field and count flowers. They use a model developed by researchers, basically a mathematical model to project yield. Then they do it twice because it's expensive, so they can only do it twice.

We could sell doing it using AI: using satellite, weather and their historical data,

we could do it for less than they could do it themselves, and then everything else was up side, and that is what sold it. The reason it's very important is when you try to actually get someone to take this technology on in the real world, if you say to a customer, "I'm going to create all this value for you," what happens is their employers push down on them and say, "Great, you're going to create all that value. I'm going to cut your budget because now you're going to actually have all this extra revenue and extra value and you're going to do it more cheaply." Not surprisingly, they're not very enthusiastic about budget-cutting. So it is really important when you go into the reality of doing this is to understand, particularly for a startup, how you do that.

I think it's different in large corporates: I'm now working in other logistics and travel industries where they can more flexibly deal with value creation. For example, reducing stress and things like that. But if you're in a startup business trying to sell an AI product, I think cost is the big driver.

Mary-Anne Williams: I think all universities really focus on innovation and entrepreneurship now. When I left university way back — I won't even tell you what decade — it was a hurdle to learn to program. It wasn't actually a course, but you couldn't get your degree without it. I think we're really moving very quickly to a world where you don't get any degree of any kind without some understanding of the process of innovation and all those things that you're saying. That's got to be good.

Ros Harvey: There's one thing I will say — I think in this audience in particular — it's a challenge managing researchers, because we worked a lot with UTS and I've worked with Pascal in the past with NICTA. Researchers want to push the boundaries of innovation, because that's what you're rewarded for. You get publications and all those things, but often the people who want impact, which is different, they just want one step better than where I am. And this creates a tension when you're trying to do both.

Salah Sukkarieh: I think universities want to create impact too — just a different form of it. I think if you look at robotics and how it has formed over the years, regardless of whether you're working in, say, a bushfire setting or an agriculture setting or environment rehabilitation setting, you throw these robots out there, and the first thing that was always worked on was just to automate the platform, just to make it work, go up and down, do whatever it might be. Where AI — maybe if I could just call it machine learning for a little bit at first — where machine learning really helped was allowing these platforms to now understand and perceive their environment, to understand it's a tree, this terrain I can drive over, something's moving in front of me that I need to avoid. That perception role — how machine learning allowed you to collect all that information, detect things, whether they're stationary or moving, and then determine how to navigate around — that became a big thing. Probably for maybe 10 years, there was a lot of effort and a lot of work in that space.

Over the last maybe five years, the value of machine learning is in the control space: how you control a platform optimally and accurately is really gaining a lot of momentum. What I mean by that is the task space of a robot can now open up a lot more. It can do a lot more things, because now you can learn the policies and the parameters that you want to go through.

Where we're coming to now is probably on two areas of high value. One is the interaction of those two things together — machine learning and perception, machine learning and planning control, and how they come together and work as a dynamic system. When you start to see those systems work and operate, you start to see what you think is intelligence in some way. It's doing something, it's reacting, it's behaving in a certain way, it's doing it optimally, etc., and it can adapt and change.

The other part is the interaction between the human and these robots — working with firefighters or working with the farmer or working with the conservationist and trying to say: if you've got this platform, it's like an extension of a Google search into the physical environment. These robots are going out, they're doing things. How do you interact in natural language? How does that process into real data? How does that come back to you as natural language? You get that synergy through that.

The challenge is to the second part of the question: when you throw into the natural environment, you miss communication systems. Connecting up to the cloud is not an easy thing to do.

How do you get these machine-learning models to actually operate on small embedded systems in the platforms, operate in real time, grab that data, and be able to do those tasks in real time as well?

Mary-Anne Williams: Riffing on your answer, I think it's great that you raise that interface between human and AI or robot. If we start thinking about agentic and autonomous systems or agents or robots, can we just outsource to them now just with some guardrails and a few trip wires? What are you comfortable actually letting these technologies do end-to-end or by themselves? Or what have you seen out there?

Dan Jermyn: There's a lot of ambiguity or what even is meant by agentic and/or autonomous. I can talk to some specific examples within the bank, which I think get the essence of what you're asking about: to what extent do you allow AI machine-learning systems to do things without intervention? I talked about the fact that we've been doing AI for a long time at CommBank. We're in a very heavily regulated industry. Everything we do, quite rightly, gets scrutinised very heavily. One of the fundamental principles around that is accountability. I'll explain a couple of systems in which things are happening autonomously in the way that you describe. It's important to be very, very clear that there is a human who's accountable for that system and checking that it works. If something goes wrong, it's that person who has to explain why, not the system.

One very obvious one is from the very first large-scale AI system that we introduced: the customer engagement engine. To very quickly recap: if you're a customer of the CommBank, and we're attempting to work out if you come into the branch, if you contact the call centre, or use the app, what's the best, most relevant thing that we could talk to you about? What are you likely to want from us? We're trying to be relevant. We can never get that 100% right, but we use AI to try to figure it out. There are hundreds of things the bank could be talking to you about or you may want to talk to us about. It could be: you want to change your credit card, you want to take out a home loan, you want to close your account. It could be anything.

We have several machine-learning algorithms trying to work out what you might

want: the propensity and the likelihood. The customer might be eligible for lots of those conversations. The AI is trying to work out, of all the things that you're eligible for, which is the one that's most likely to be relevant. If we can get that right, more often than not, it's a better customer experience than just doing it randomly, which is what most companies would do. Those algorithms are running autonomously, they're self-learning, based on customer feedback, in real time. The very first was eight or nine years ago.

Mary-Anne Williams: Well, I don't know how you're going to get from four to one now if you're already doing that.

Dan Jermyn: I know. The bar is very high, but they're self-learning algorithms. We saw the power of that during COVID. We think a lot about the risks of AI because there are lots of risks associated with the technology that's evolving so quickly and are so complex. What we found in the implementation of this is no accident. It's by design: where we implement AI reduces our risks. At least that's what we're attempting to achieve. This is a great example. We have this autonomous self-learning system. We have lots and lots of automated checks and balances, by the way. From COVID, one of the principles underlying it is customer channel preference. We saw a couple of days when the model started to really drift in terms of expected behaviour. Some of that was to do with channel preference stuff. Alerts kicked in — bing, bing, bing, bing, something a bit wrong here. Maybe a data feed hadn't turned up. We defaulted to older models.

Mary-Anne Williams: Okay, Ros, what about you?

Ros Harvey: I think we probably are doing it customer-facing and then internal. The only way you can make this technology scale internally is to automate the process, AutoML, to actually create the models and results, because it's way too expensive with high labour costs, particularly the cost of data scientists. We have an automated ML (machine-learning) environment that is supervised, because obviously you need to have serious alerting set up on this and you need to be able to constantly measure your own accuracy of the models. We have detailed feature sets that we've developed through lots of contextual learning. We have auto ML environments that just trawl through different models, different combinations, to actually come up with new model results. That's internally focussed, if you like. We're increasingly using AI for coding internally.

On the external, I think we're basically at an automated system with automated supervision and accuracy checking in yield prediction and weather prediction. It's supervised in the sense that you've got an alerting system that tells a human if things are going off outside.

Mary-Anne Williams: What's the worst thing that can go wrong?

Ros Harvey: The worst thing? Well, we're not in an industry where someone dies if we get it wrong. What could go wrong? Well, for example, say you decide you have too many apples growing on a tree and you're trying to get the right number of apples on it for the season, so that you're not using unnecessary inputs that are going to give you too many little apples. Then you apply a chemical, a hormone, a product that effectively strips the crop to get you to the target number of apples. Then it has to be applied in a particular window, because if you don't get the temperature and humidity

right, it's a little bit like napalming your tree. That is catastrophic. You can lose a whole season if you get it wrong. It is actually quite important to get the weather prediction for this thinning right, and you'd want to have a lot of confidence. You need to be careful on what your use cases are and what your risks are and what your confidence in your model are to do that. So, yes, you can have pretty bad commercial impacts.

Mary-Anne Williams: The idea that you can ahead of time specify how many apples you're going to allow this tree to grow?

Ros Harvey: Very known science. Developed by biologists.

Mary-Anne Williams: Okay. I have more respect for apples.

Ros Harvey: It's more about getting the weather prediction right, because our weather is on these big global forecasting systems, and rich countries like Australia add weather stations and we grid down. Maybe you're going to get to a 4 × 4 kilometre grid, where we use AI to predict inside the growing tunnel, inside the canopy, very, very specifically what the plant is experiencing. You can only do that if you're forecasting weather every 15 minutes, generating new predictions that are going out for every variable, massively out, up to 10 days and running this all through models. You have to automate this. You just cannot do this any other way.

That's where there are significant risks. What that means commercially is you might choose not to release a product that says, "Apply this chemical in this window." You might say, "What I'm going to tell you is what our best model estimates of these conditions are." It's a decision to support, not an automated system. With time, it will become automated, but I think we need these feedback loops so that you can effectively ground truth by what you're seeing.

Salah Sukkarieh: It was a good point that you made at the end, Ros. We used to apply AI systems or machine-learning systems, just black boxes. But the research over the last few years has done a couple of things. One is developing new machine-learning algorithms that also predict the level of certainty of their results. That has become an important element because now the output isn't just, "Here is the answer that I'm going to give you," but, "This is how confident I am about that answer." To follow on the decision support tools, what that also means is that if you keep running those algorithms, then what you give the user is options and different levels of risk maybe, or different levels of uncertainty.

The other part from the algorithm side has been recently how the machine-learning algorithm makes decisions. You're getting that output as well now, not just the final result, but: what it's doing, what it's processing, why it reprocesses something. That thought process, the algorithm itself goes through, helps the end user. Both of those situations have allowed the end user, the operator, to have a lot more confidence in the system for it to behave on its own. In our applications, we've been able to work from TRL[12] (Technology Readiness Level) 1 all the way through to TRL 9 in these different applications. What you find as you get further and further up the TRL is some of

12 Héder M (2017) From NASA to EU: the evolution of the TRL scale in Public Sector Innovation. *The Innovation Journal: The Public Sector Innovation Journal* 22(2). Available at: https://innovation.cc/document/2017-22-2-3-from-nasa-to-eu-the-evolution-of-the-trl-scale-in-public-sector-innovation/ [Ed.]

those machine-learning algorithms are not trusted by the end user for obvious reasons. A 400-ton truck on a mine site, a 20-ton fire truck, a 1-ton tractor, etc. You don't want these things running around spraying anywhere or going off-track. What you end up having is machine-learning algorithms and a change in operational behaviour. The two things are starting to come together. I used to operate in this manner, but now because of these machine-learning systems that are going to allow me to do something more cost effectively or repeatedly or whatever it might be, I will change the operation to allow that system to behave in that manner. So we're seeing those changes as well.

Mary-Anne Williams: I'm going to ask the next question with a little bit of a twist: we want to explore human-in-the loop adoption. Since you're all leaders and doing fantastic innovations, we want you to surprise us and give us an insight, an actionable one that we couldn't get out of ChatGPT. Something that is surprising because you've found it unintuitive when you've been working and building these systems. Just what has made you stop and think?

Ros Harvey: I've got one that I find difficult to believe. It comes going back to the day in oysters. We were predicting salinity levels. Salinity is a proxy for the runoff of water; because oysters are fresh filter animals, the health authorities say you can set a salinity level as a proxy for runoff by catchment. When they go over that salinity level, the growers stop harvesting. It's bad financially to do that. We started off by just trying to monitor real-time salinity. We did this with CSIRO, during the day, and that was very effective. Then the grower said, "Well, couldn't you tell us when it's going to actually be shut?" i.e. predict, which actually was an important innovation that we would predict. This was 2014. We had to come up with a machine-learning model to do this, to predict.

The thing that really surprised us, about after a year of doing this, is that we were taking data feeds like rain, wind, etc. And we were taking the live feed of the salinity and effectively building a model that would predict salinity from these environmental variables. The feed went down unexpectedly, but we continued to predict very successfully without the actual environmental variables. How can this be? I just didn't understand this because I was thinking about it as like an offset thing. And the team said, "Look, at the end of the day, the best predictor of something is itself. If you get enough data and you're feeding the error rates back into the model, it can actually predict itself." And I'm like, wow, that's amazing. We turned It off, but it continued to predict pretty accurately.

Mary-Anne Williams: Okay. Dan, what do you think?

Dan Jermyn: It's continually surprising. And that's both exciting and quite horrifying depending on precisely which job you have. In the example I was giving earlier, of the customer engagement and those models that were degrading, it turned out to be the earliest signs of people modifying their behaviour as a result of COVID. So the models weren't wrong. They were right. Our expectations were wrong. There was something happening at a macro level that the AI had cottoned onto that we hadn't. That was remarkable. That was not an intended consequence. That's not what the machine was for, but we were very pleased to see that it had "worked" and is accurately telling us what was happening out there.

More recently, the way people are using the tooling and what it means, I suppose, is it's evolving so quickly. If you had asked me a couple of years ago, we would have said a lot of what's happening in code generation AIs is very exciting.

It's going to democratise things, it's going to lower the bar, we're going to have a lot more people being a lot more productive with coding. What we've actually found as it got more and more sophisticated is the really elite engineers are being wildly more productive as a result, but the incremental benefits for an average coder are not really there in the same way. Rather than thinking about a world in which we're going to have lots of loss, more of these engineer specialists, actually where we're getting the real value now is from folks out in lines of business, maybe in more traditional areas, like risk management or in finance who are creative problem solvers and deep thinkers. A lot of the very, very best use cases that we get are being driven from areas that we definitely wouldn't have seen even a year or so ago.

Mary-Anne Williams: Salah, you have the final word.

Salah Sukkarieh: Maybe two examples very quickly, both of them related to how the human, the end user, has reacted to the output of the system. I'll take Ros's previous example around the apple industry. We used to work on perception systems on robots that could identify and count every single flower down a row, and then build a flower density map across a whole paddock or an apple farm. Six months later, you come back again and you do the same thing, detecting the individual apples. The idea is that this longitudinal study and data collection allow you to predict yield, so you keep going through that process. Really great farmers would stand on top of the hill, look down at the paddock, look at the approximate flower density, and estimate how much yield they would get at the end of it all. And we would go through the process and calculate the amount of yield as well, and we'd come within plus or minus a percent. So his estimates were perfect; that showed that the machine learning also could do it in real time and over that period of time, but you got roughly what the farmer was predicting by looking at the paddock. The difference was, we noticed that there were density changes across the different rows. So he was very good, repeatedly predicting the estimate of the yield because he would look at the flowers and he would then place the beehives for pollination purposes and get the yield that he was predicting. But you could then optimise where you placed the beehives to get a greater yield. That was not known up until that point, but there was a lot of reaction in terms of not wanting to do something. And that was one learning.

The other one was work that we did with Qantas on flight path prediction, fuel prediction, and trying to get optimal trajectories for the aircraft to fly, given the winds. That was another scenario where we were predicting behaviours of where the aircraft should fly, which looked completely different to the last 20 years of dispatching of flight plans. But once the upper management forced the dispatchers to fly the routes that they were going through, we were getting routes going right over Antarctica, for example, following through paths that you would never follow across to the US. They were getting 1%, 2% fuel savings on these flights, which is quite significant — another one where, from our perspective, we're

applying machine learning and optimisation algorithms etc. to a problem.

We feel like it should do better, but then you've got— I don't want to call them the gatekeepers —the individuals who are so used to certain behaviours and processes; how they react to that result is quite novel.

Mary-Anne Williams: Okay. Well, I've got to apologise for the audience now because I've clearly got three optimists here. Well, all we got was serendipity. We didn't hear about all the disasters, but we're now going to turn to questions.

Q1: I'm an AI ethicist working in the federal government at the moment. I think this is probably for you, Dan. In 2019, our government published eight AI ethics principles,[13] and we have a body of work from, I think, Fifth Quadrant called the Responsible AI Maturity Index that looks at the relationship between effectively 500+ organizations' stated commitments to the eight principles versus how well they think they're operationalising those principles.[14] I just wanted to get a sense at CBA, how effectively do you think CBA is operationalising those AI ethics principles? What have been the main learnings from the challenge of attempting to do that in the real world?

Dan Jermyn: It's a great question. Yes, we're familiar with the Fifth Quadrant stuff and I've been through that process. I think it's a great piece of work. Actually, it's an involved answer because I'm very proud of the work that the team does on responsible AI. At the same time, we're very much grounded in the fact that we're working very quickly into a space that's evolving very rapidly. We need to be on our toes, and I think it behoves us all to have a certain humility. One of the things we've done is partner with a lot of very smart people externally, such as Bill Simpson-Young from the Gradient Institute.[15] There are other many great academic and other industry partners that we've very thoughtfully worked with, specifically because they have an independent and external lens on what we're doing. I think that's been really key. I certainly would like to think that we invite and receive — and quite rightly so — a lot of scrutiny. It's okay. It's good. Should be that way.

I think we're doing well. I certainly think that we've put a lot in our research agenda. We actually have a lot of really good muscle in this into matters of transparency, unintended bias detection. This is really what our work was predicated on from the very start. There's a lot of well-honed expertise at managing risk well. There's a lot of execution muscle in risk and governance within the bank naturally. What I think we could get better at is probably a good question, where we need to evolve.

Certainly, we're pushing very hard at innovation, at the boundaries of some of the things that are happening in, say, agentic AI, for example. We're very excited. It's actually an example where we have other autonomous parts of processes running and we're seeing better customer outcomes, but we have an ambitious agenda. We know that it becomes an ever more complex system as

13 DISR (2019, 2025) Australia's AI Ethics Principles. https://www.industry.gov.au/publications/australias-ai-ethics-principles [Ed.]

14 Fifth Quadrant, The Australian Responsible AI Index 2025. https://www.fifthquadrant.com.au/responsible-ai-index [Ed.]

15 https://www.gradientinstitute.org [Ed.]

it scales. The way that we're thinking about evaluation frameworks are the safeguards, the guardrail models that we've put in place for our customer-facing generative AI. This is something where we've done an enormous amount of work, but as the models get smarter and smarter, the work that we have to do to keep them safe becomes harder and harder too. So I am pleased with where we are so far. In fact, in that evident AI index, we were ranked number one in the world for responsible AI leadership, which is very good for Australia. Australia in general has a great tradition and a name globally for responsible AI.

Q2: If we'd had this meeting 50 years ago, we would be talking exactly the same way about the start of the digital revolution. What was interesting back then — because I actually started my career a bit before that — was that all the software had embedded assumptions. Of course, understanding the impact of these embedded assumptions was critical to understanding how you should use the software and in what particular case. Does that apply today in what we're talking about in terms of the embedded assumptions of all the software and tools that we use? Are these easily found and understood so that they can be used properly?

Dan Jermyn: If you want to gather your thoughts for something more erudite, I think it's a brilliant question. I think it's a more fiendish challenge with AI than it was in the digital revolution. I think it's implicit in all technological shifts, but I can give you a tangible example of some of the work that we've done with CommBank, because we worry about that. We worry about large foundational models that allow us to do incredible things, but they're trained on data. We don't always know what, we don't always know how, and therefore we don't always know exactly what the outputs may be. So how can we answer or at least investigate some of these questions? We did a study. One of the very best parts of my job is where we get to work with our behavioural scientists and the human-centred AI work. That's really not about the technology itself, but how people react around it and what it really means.

We did a study using large language models. Most, if not all, of you will be familiar with the late Daniel Kahneman and *Thinking Fast and Slow*[16] and Prospect Theory[17] and lots of well-published studies about what people will do if you say there's a 50% chance of you getting $5 or a 1% chance of getting $1,000: people will make "irrational" choices about that. And we thought, well, what LLMs will do, so we created a survey. We surveyed a thousand real Australians. We created a thousand personas like the real Australians, and we asked them the survey questions are for this very repeatable, very frequently run study that always gives roughly the same results.

The results were remarkable. So we ran it and the first question gave almost identical responses. The Australians answered it the same way everyone always answers it; the LLMs answered it the same way everyone always answers it. Then we asked question two, and I think this was a mistake. The team told me they did this deliberately, but they got the numbers wrong by a factor

16 Kahneman D (2011) *Thinking Fast and Slow*. NY: Farrar Straus and Giroux. [Ed.]

17 Kahneman D and Tversky A (1979) Prospect theory: an analysis of decision under risk. *Econometrica* 47(2): 263–292. [Ed.]

of a thousand. The real Australian people answered in a wildly different way to what the traditional survey would say, but the LLMs answered in exactly the same way as the non-mistakey version would. It's a really visceral example of how easily we can be duped into thinking that, well, this is a reflection of society and the LLM is humanity talking back — well kind of, but also kind of not.

Salah Sukkarieh: I was just reflecting over the processes again. There's a lot more software being written now, so there's a lot of activity happening. You think you're on your favourite model and then next week there's a new model and you think, okay, we've got to try that one because, for whatever reason, we've got to keep up with the curve and get that going. But the same processes still lie from requirements all the way through; anything you write, you bring that through. When it comes to machine learning, usually we're asking questions such as what is the data that's being fed and that goes to that learnt model. We're a little bit more controlled, in the sense that we have that control to be able to do something like that. I think it's the LLMs where there's a lot more uncertainty.

Even in the context of the research of applying that — not so much doing research on LLMs, but using them in the research context – there are a lot of questions still there around: what data was it trained on and how can I use it in that context? That's probably where the greatest uncertainty is at the moment.

Mary-Anne Williams: Well, let's wrap up. But I'd like you each to identify one skill NextGen needs to thrive in this world we're rapidly entering.

Salah Sukkarieh: Co-designing.

Ros Harvey: Communication. That means listening as well.

Dan Jermyn: Creativity.

Mary-Anne Williams: Co-designing. Communication. Creativity. Well, there you have it. Thank you so much for this stupendous panel. You have been amazing. Thank you.

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, pp. 127–142. ISSN 0035-9173/26/010127-16

2025 Royal Society of New South Wales and Learned Academies Forum: "AI: The Hope and the Hype"

Panel Session 5: AI — Future Directions[1]

Hugh Durrant-Whyte, NSW Chief Scientist & Engineer
Jie Lu, University of Technology Sydney
Pascal Van Hentenryck, Georgia Institute of Technology
Toby Walsh, UNSW Sydney
Lina Yao, CSIRO Data61

Sally Cripps: Our final panel for today is AI and future research. It's chaired by the NSW Chief Scientist and Engineer. He's going to give some closing remarks at the end of the panel.

Hugh Durrant-Whyte

We have the world's best people for what's going to happen in the future, which is a great subject to end the day on. We have Jie Lu from the University of Sydney; Pascal van Hentenryck, who gave the Keynote, from Georgia Tech; Toby Walsh from UNSW; and Lina Yao, who's currently at CSIRO, but is about to move to UNSW as a future fellow. I'd like to talk about the future of AI. There are lots of diverse views, and I'd like to do it in two parts. First, I'd like to talk about what's going to happen in the next one or two years. What is the current thing going to look like? Should we all be out investing in AI shares or should we all be careful and buy a bunker and put it in our garden somewhere? Second, what should we be doing in more like 20 or 30 years? How will it change humanity? I think that's a different take on these things. Jie Lu, I'll start with you.

Jie Lu

I'm Jie Lu from University of Technology Sydney. Currently, I have two roles. First, I'm the director of the Australia AI Institute at UTS. We have 30 staff members covering almost all areas of AI: machine learning, computer vision, natural-language processing, and robots and so on. And 15 postdocs, and 220 PhD students. Quite a large AI institute. In the last four years, UTS has been ranked top in AI in Australia and one of top 10 globally. Second, currently I'm an ARC Laureate Fellow; my topic is autonomous machine learning to support decision making in complex situations. This will be finished early next year and immediately I will move to my industry Laureate Fellow to work on AI for women's health.

The first question: one or two years. That's relatively easier than 20 years. In one or two years we will have some directions, particularly from a research point of view. The first direction is autonomous AI systems. You have lots of apps in your computer, on your mobile phone. What we are doing now is AI agents. An AI agent system will deal with the tasks, and close all the apps to minimise human supervision. At the same time, AI with autonomous machine learning will

1 This is a transcript of a presentation, available at https://www.youtube.com/watch?v=RKsxXSYTWOg

autonomously pick up relevant data to support human decision-making. The second direction is AI for science. Science has lots of unsolved problems: in mathematics, physics, chemistry, and health and so on. AI will drive new science discoveries. Two very important directions, I think.

Hugh Durrant-Whyte: I certainly agree about the science one. Last year, two machine-learning researchers won the Nobel Prize in chemistry. That shows you what's changing in scientific discovery. Lina, give me a brief introduction and what do you think about the next two years? What are the big problems?

Lina Yao

I'm Lina Yao. I was with CSIRO Data61. I was a senior principal research scientist and science leader for translational machine learning. Very soon I will move to UNSW as a professor and carry out my ARC Future Fellow project. My background is machine learning and data mining.[2] We wrap up everything as AI. My research has been gradually evolving from future learning, generative modelling, and reinforcement learning, to deep learning. Currently, my focus is lifelong agentic AI for multimodal decision-making. I will focus on addressing three critical research challenges. The first is alignment: it's a multifaceted concept to align the LLM agent with human values, as well as physical laws, for example, to handle some knowledge-intensive tasks. The second focus is how to equip the agent with the capability to self-involve by learning from the interactions and the dynamic and contextual environment. The third focus is how to further improve reasoning ability.

In the next one or two years, I believe we will witness the thriving of AI agents, especially LLM-powered agents. I can summarise it in two words: first, *verticalisation*, because everyone has been making great efforts to transfer the general ability of LLM agents to adapt some domain-specific tasks. That means how to transfer the ability from a general purpose to vertical domains like eBusiness or other things. The second word is *personalisation*. The agent will involve and learn to customise itself to fit each individual. That will be of great benefit, especially for education and e-commerce.

Hugh Durrant-Whyte: Toby, you've been in the area quite a while now. I guess you've seen ups and downs in the AI area. I'm really keen to get your view as to where you think things will go in the next couple of years. Are we going to see huge growth? Are we going to end up with a fall? And a bit of background, too.

Toby Walsh

A little bit of background is that Hugh used to be my boss when he was CEO of NICTA[3] and Pascal used to be my boss when he was a group leader at Data61. I'm pleased to be on a panel with my ex-bosses who still want to talk to me. I think it's easier to predict 20 years in the future than a few years in the future. I think it's easy to say we'll be using AI to accelerate the way we do science, transforming the way we do science, to do things that look at data sets that humans just don't have the capacity to look at, in order

2 Hawkes H (2021) This award-winning engineer turns big data into useful AI applications. *Create Digital*. https://createdigital.org.au/engineer-wants-to-help-ai-become-more-human/ [Ed.]

3 See https://en.wikipedia.org/wiki/NICTA [Ed.]

to transform genetics and medicine. In 20 years' time, you can be confident that we'll have completely transformed the nature of those pursuits.

It's much more difficult to predict the next couple of years. Until recently I thought that the AI bubble was yet to burst because we are only now talking about it being the "AI bubble." I remember we were talking about the dotcom bubble for several years, in the early 2000s, and it was only after we stopped saying it was a bubble — but maybe this is a new type of economics — that the bubble really burst. I was thinking maybe since we've only just started talking about the AI bubble it might last. But then yesterday I read that Google is proposing putting AI data centres in space. It's madness. Why would you spend money trying to get GPUs up into space? It really cannot make any sense other than the fact that you can.

Hugh Durrant-Whyte: 24-hour-a-day solar panels?

Toby Walsh: Well, yes, but we're a vast continent. The sun shines most of the time. We have batteries as well if we want solar in the middle of the night. It sounds to me like the worst excess is to think that we should waste the planet's resources by sending rockets up into space to have data centres up there. But let's move on.

It's unprecedented that we're spending billions of dollars a day, the scale of investment happening. Clearly some of this is very speculative. A number that caught my attention the other day: 2.5% of the USA's GDP is being spent on AI data centres. That is almost again without precedent; there's only one precedent I could find, which was at the height of the nineteenth-century railway boom: when they were building the railways and the industrial revolution in the US, they were spending two and a half percent of GDP.

But there's a useful historical lesson to be learned there: lots of people lost money. There were lots of public bonds being invested in the railways. A few people — the Vanderbilts and the Rockefellers — made a lot of money in building the railways, but most people lost everything. But building the railroads laid the infrastructure for the economy that became the United States. The same I think is true here: the bubble will continue to expand for a couple of years; maybe in two or three years' time, it will burst. A lot of people will lose a lot of money. A few people will make a huge amount of money, but that's irrelevant in the long-term scheme of things because that's laying the infrastructure for the new type of economy — the AI-powered economy — that we will have in 10 or 20 years' time.

Pascal Van Hentenryck

I'm going to disagree with you guys. I think it's much easier to predict the next two years than the next 20 years. I'm going to tell you why. Toby started stealing what I wanted to talk about: the dotcom bubble. What happened in the dotcom bubble was very interesting. It came to the public consciousness that there was this thing called the Internet, and everybody started building websites and e-commerce and so on. 99% of them failed. One of them wildly succeeded and is now one of the largest corporations in the world. So this is what's going to happen in the next two years: everybody's going to build an AI agent for X. Everybody's going to build a chatbot for X, and 99% are going to fail, but some of them are going to succeed and do something very interesting.

We'll talk again in two years, but I think that's going to happen.

Toby Walsh: Yes. You're right. Many of those businesses failed in the dotcom bubble, but two — Webvan online groceries and pets.com — are now proper businesses. Those things became real businesses but it took a while.

Pascal Van Hentenryck: That's the first thing. And there is another thing, a big one. If you've seen eight-year-old girls playing soccer, they bunch together all the time. They just don't know how to play. This is going to happen in science. Everybody is saying, oh, LLM is the new tool. I'm going to do LLM for X, for Y, for Z, for everything. All scientists are going to do that. It's obvious: 99% is going to be irrelevant. But some are going to do some interesting things. They're going to try to make them more trustworthy. That's what's going to happen.

And the third thing that's going to happen is that there will be a lot of good work in AI for engineering that is already getting close to being mature. I completely agree with Ros: academics are good for doing proof of concept, but you have to bring things to market. I think that's what's going to happen in the next two years. Finally, I'm going to agree with some of the things that you have said. What's going to happen in the next two years is AI for scientific discovery. The Nobel Prize in chemistry was an exception because they had massive amounts of data. What's going to happen in the next two years is accumulating a lot of data on many different scientific fields such that you can do the same thing in those field. I'll stop there.

Hugh Durrant-Whyte: Let me start by disagreeing with the last thing you said. One of the things that's interesting about that Nobel Prize is they had almost no data. They only had 200,000 protein sequences and they used it to predict 20 million of them.

Pascal Van Hentenryck: Yes, but they had secrecy. They used —

Hugh Durrant-Whyte: A good model.

Pascal Van Hentenryck: And I think that's important. And it was not an LLM, so this was very special.

Hugh Durrant-Whyte: I want to unpack that slightly before we move on to the long term. There is a big focus on LLMs. Like the engineer who said something a little earlier, I do have a concern that you cannot say or provide any guarantees about the use of LLMs in almost anything. This means that you can only use them in things where there is a person involved. There are no known rules to transfer one thing to another thing, if you see what I mean. I wonder whether we're putting all our money on a horse that isn't even going to finish the race.

Pascal Van Hentenryck: Yes, but I'm not putting my money on LLMs. Let me tell you what I see happening. LLMs use a very general machinery. People are starting to build foundational models for natural language processing. There are a lot more foundational models appearing: foundational models for the grid, foundational models for X, for Y. They are going to try to capture some physics or some chemistry or some biology and so on. I think by definition and the way they are trained, they will have limitations. The next step is to combine that with techniques that can guarantee things, and there will be a lot of progress there. I think LLMs are going to be useful as a first step, but they will be complemented by reasoning techniques and things that all of us have been working for like 40 years, and

they're going to start merging in interesting ways. It's already happening.

Toby Walsh: Think you're unduly pessimistic about the failings of LLMs, which are many, because mostly they're also the failings of humans. But we put humans in delicate places. I think medicine is a wonderful example: we've constructed these institutions around the humans — medics — that we put in these life-or-death situations, so that we can put our lives in their hands.

The challenge we have is not that LLMs are a bad technology that can never be fixed, it's that we need to construct the institutions around them so that we can put them in those sorts of places.

Hugh Durrant-Whyte: That comes to a key thing in the next couple of years. Now let's spread ourselves out to 5 or 10 years. What do we need to be doing now to make sure we are in the right position to use AI properly and actually make it apply in engineering, in human situations, in science discovery, in all these areas? What problems at a science level do we genuinely need to start addressing, and what social problems do we genuinely have to start addressing?

Toby Walsh: Let me mention one of the elephants in the room, which is that I don't think we can continue to move fast and break things, when we're breaking a young girl's body image, when we're breaking our democracy. Given the "experiment" we did with social media and the tech companies, I think we've discovered hopefully that maybe we have to take a little more care and that maybe we have to be a bit more careful. It does worry me that, for example, we're going at speed now with agentic AI: giving this stuff some agency to do more things in the world and do more harm. Yet we're still not putting in the governance and guardrails. The other problem — and again, we saw this with social media — is that sometimes we're talking about really small harms, but multiplied by billions of people.

We're already starting to see this with AI chatbots being used as therapy bots, being used as companion bots. We're already seeing that people have committed suicide: I was listening to a very sad interview with the parents of someone who sadly committed suicide after long sessions with ChatGPT. At one point, ChatGPT said, "Can I help you write your suicide note?" It blew my mind. There were no safeguards to prevent the chatbot from actually saying that.

Hugh Durrant-Whyte: Okay. Yes, I agree. Jie, what are the problems we need to be working on the next five years?

Jie Lu: Well, we are at Industry 4.0 at the moment. Industry 4.0 has two main key technologies. One is automation, the other is big data. Consider the challenges. After five years, probably in 5 to 10 years, we will be at Industry 5.0. The key technologies for Industry 5.0 are not big data, not only automation, but human-machine collaborations and personalisation or personalised AI. Human–machine collaboration is a challenge because humans and machines will be team members. There will be machines including software and robots and AI tools and humans. Consider you have a team — including humans and AIs — and you want the machines to understand the humans. Humans should control AIs, how to manage the team, how to work collaboratively. Currently, humans mainly use AI as tools, but later tools will not only be your team members. How will you manage them and collaborate and work with them to achieve something each side couldn't

achieve separately? That's the key aim for human-machine collaboration. I think that's a challenge.

Hugh Durrant-Whyte: Okay. I'll ask you about the science one in a minute. Lina, what do you think? The next five years — you've got a fellowship. What are the problems that really need solving so that we are not in these sorts of positions with chatbots and so that we can genuinely work effectively with AI?

Lina Yao: My focus would be to explore how to develop more effective reinforcement-learning methods for verified rewards, because, as we know for the current LLMs — large language models and foundation models — reinforced learning from human feedback is one of the mainstream post-training methods to further align the LLM agent with human preferences and knowledge laws. But, for many tasks, it's hard to define the reward model. It's hard to define the verifier. For mathematical questions, we know the answers are usually unique, so it's easy to verify the model. If anything can be verified, then it can be solved using AI. But for some tasks — for example, asking AI to judge if the generated methods are novel, and the tests are satisfied or not — it's actually more subjective. It's hard to verify them. The foundation that I'm very interested in is to make sure the agent can be employed in a safe manner, and how to define a comprehensive verified framework.

Hugh Durrant-Whyte: I come from a robotics background, and, for me to get a robot into the field, I have to guarantee it will never go wrong. It goes into the Pilbara or wherever and operates for 15 years and it can never go wrong. It's very hard to use some of the now-standard AI techniques to build systems like that. I have to come up with — if you like — not just guardrails like you were talking about, but genuine ones that provide guarantees. And the conversation we had — if you remember, Pascal — was about optimising energy grids. Again, you have to make some kind of guarantee for the AI that's used in that. You can't just let things rip, if you see what I mean.

How do we genuinely develop a new approach to even developing these things that provide useful guarantees? What will happen and how will they behave?

Pascal Van Hentenryck: In the space you worked on and the space I'm working in these days, you have some notions of the physical work, the infrastructure and the way you control things — your robots, Hugh, or my grid. I think the key point is that we should not rely on something like an LLM and assume that they know the rules of physics or the control systems that we have to use. We also need fundamental understanding of the physical world or the biological world and that has to inform the system. We can still use predictive learning techniques for approximating things, but at some point we have to satisfy those rules — the rules of the physical world, the safety constraints — that are hot constraints. I think we can do that. That's one of the things that I wanted to mention.

You can do that and we should do that, but that's very different from what is happening now. In a sense, what I'm talking about here is that what we need to define is not large language models or large foundational models. Instead, we should define smaller domain-specific foundational models where we understand how we apply the technology and can provide the guarantees.

I work in an area which is nice: optimisation of control systems. Fundamentally, we have guarantees in the systems that we are using right now, but they're just too slow. What we have to do is preserve those guarantees and speed up the technology by orders of magnitude. For me, it's like using domain expertise and then using AI technology, then merging them together. That's the winning proposition, in my opinion.

Another thing to address is the question that you asked before: how do we do that? We need better education. We need people who are actually more generalist instead of being really highly specialised. Everyone should have AI classes when they start undergraduate education, they should have them in high school as well, so that they understand the technology and the limits of the technology and also the potential uses of the technology.

Toby Walsh: This is exactly the weaknesses of LLMs and the strength of science. In science, we have a model of the world and we have physics equations that describe the evolution of that model. LLMs are remarkable in that they manage to achieve what they do without any of that.

Hugh Durrant-Whyte: The thing that I am most interested in in my research life at the moment is what machine learning can actually do for scientific discovery. My focus interestingly has not been in robotics. It has been in biology, and biology is not physics. There was a wonderful article published three days ago in *Science* (the front cover of *Science*) about building a machine-learning-based model for a cell, metabolism, everything.[4] They've been working on it for 20 years. It's not a trivial thing. The ability of that, to be used in terms of predictive outcomes for disease and things like that, is potentially remarkable. AI is more useful in domains which are not physics, but have opinions and so are hard to model. In some sense, data domains model better.

I'm interested as you said in the science discovery. I'm going to come to Jie in a minute but just I'll let you have one or two words.

Pascal Van Hentenryck: Yes, one or two words. The book I'm reading now is the life of the cells in biology. This is like the century of the cell. This is super interesting. As I said before, my daughter is a doctor and we talk about cell biology and systems all the time. I agree with you. I think what we need to do is understand the systems better, and that's where machine learning can help. That's where simulation techniques can help and learning them at higher speed is helping. I completely agree that scientific discovery in the biology space is the next frontier. Australia is in a very good position, given the biodiversity that you have and the scientific basis that you have. It's a massive opportunity. If I had to put money on something, that's what I would put my money on.

Hugh Durrant-Whyte: Jie, you talked about AI applied to science discovery at the start. Tell us where you think it's going.

Jie Lu: I can give you an example. My team and I have been working with 23 students in Australia at a Sydney gene industry company. We have been working together for 7 or 8 years through many projects. What we are working on — one important thing — is

4 Leslie M (2025) The silicon cell. AI cell models could transform biomedicine — if they work as promised. *Science* 390 (6772), 30 October. [Ed.]

to try to discover new associations between genes and disease. We do lots of genetic data analysis and have generated some new associations for some significant diseases like strokes. Currently we are working on indoor metrics which will be used for the company to put into their service reports and also provide to clinical staff to support their decisions and to give an early warning to people who do gene tests. That's a guarantee, as you mentioned. We not only improve the product activity — not only improve the accuracy of prediction — but we also consider guarantees. We use mathematical models to prove some algorithms, but, more importantly, we use more data to test our models.

In Australia, there is genomic data, but we did not think it sufficient, so we also talked to the UK Biobank — the UK's biggest gene data site — about data access. We also talked to the US, which also has very big gene data sites. Each has millions of people's records. We developed advanced transfer learning, and transferred knowledge learned from the UK Biobank and the US to support our predictions for Australia. This also makes more evaluation, more measurements to improve the accuracy and guarantees of our models to support industry to put into their service reports.

Hugh Durrant-Whyte: Lina, what about 20 years from now — when I'm going to be retired by a long way? What is this society going to look like? Let me put a bit of context on that. Sometimes I look back 25 years ago, when even then we were building autonomous vehicles and we were doing AI. In fact, there was a period where robotics did AI and nobody else did AI, if you see what I mean, because it had fallen off the edge with expert systems. In 20 years, what do you think we will be doing?

Lina Yao: In 20 years, I assume that we will already have overcome the limit of energy. You see nuclear or any other … We'll still —

Hugh Durrant-Whyte: Be arguing about whether to have a nuclear plant?

Lina Yao: I think AI will be embedded in ubiquitous devices. Maybe not our mobile phones anymore, but maybe our glasses, our watches. Each very tiny device can run very comprehensive and powerful AI agents, so I can see VR and MR also will be thriving. I hope my daughter will enjoy the future world, when AI is everywhere, and maybe when they go to Mars as well.

Hugh Durrant-Whyte: And should we be worried?

Lina Yao: No. I think the government and the researchers and practitioners, in parallel, are working on how to better regulate and provide foundations for AI development. I think they can be developed in parallel, so I'm not worried.

Hugh Durrant-Whyte: All right. I'll come back in 20 years and ask you how it's going. Toby and Pascal have their own, and they can fight over it. You might be retired too by that point, right? What do you think it's going to look like?

Toby Walsh: I hope it's going to look good. We'll have solved the energy problem. It's a temporary problem, the energy problem, because at the moment commercial firms are getting technical advantages by outspending each other. We know that intelligence doesn't require huge amounts of energy: intelligence requires 20 watts of energy. A third of the energy your body produces goes to your head. Interestingly, your brain uses more energy than your heart. To use so

much energy, it has to be on your shoulders to radiate the excess heat.

You can do intelligence in 20 watts and we want to do intelligence in 7 watts because it can then run on your mobile phone and you don't have to share it with anyone. We're already starting to shoehorn stuff into mobile phones. The energy issue is a temporary thing, where people are trying to buy competitive advantage by outspending each other on it. We should be outraged that they're burning more carbon dioxide and using up all our water. There's no necessity for them to do that. We will get to a point where —

Hugh Durrant-Whyte: Send them out to space. Check.

Toby Walsh: Yes, we'll be sitting in our self-driving cars. I'm pretty sure the cars will be doing the driving, and that will be fantastic. The thousand people who die every year in road traffic accidents in Australia will not be dying in accidents. There will still be deaths because still trees fall onto roads and other random things happen, but it will be a significantly smaller number. We will look back and think, "It was the wild west. We let people drive cars. Can you imagine how mad that was?" For the petrol heads in the room, I'm sure we will still have people allowed to drive cars.

Hugh Durrant-Whyte: I'm old enough to remember Carnegie Mellon's No Hands Across America project in the 1980s.[5] They drove 98% of the way across America without ever touching the steering wheel. I agree we've got to 99.99% now, but it's interesting how little progress we've made in that area, given the length of time. That was 40 years ago.

Toby Walsh: Yes, the final few per cent are quite tough.

Hugh Durrant-Whyte: You can have long-distance trucks doing most of the long haulage. That's going to transform our economy. And what about negativities? What are we going to watch out for in 20 years?

Toby Walsh: That's the positive future. All the possible futures are still available to us. That's the problem. But there's a future in which we're starting to be in a world where everything that we see is chosen for us and is actually even made for us by AI. The algorithms are selecting what content we see, the algorithms are increasingly generating that content, and what sort of world is it in which we're all living our own different filter bubbles, and we're all being told different untruths, and we can't tell what's true or false.

You see a little video and you're not sure whether it's true or false and what sort of world it is where we lose touch with truth because we don't believe anything. The price is paid by the people telling the truth because we don't believe anything we see anymore because so much of it is fake. Already truth is a pretty fungible idea, a pretty contested idea. What sort of world might it be, if we're not careful, in which everything is contested?

The things that we should be believing in — like vaccines — that have saved millions of lives, that have transformed, got rid of polio and things like this from our society that used to afflict us, like the climate emergency, that we're still debating — we've still got politicians who debate these. The veracity of the climate is it's happening. You

5 https://www.brownstoneresearch.com/bleeding-edge/no-hands-across-america/ [Ed.]

can see it. What sort of world it will be if all of our truths are contested?

Hugh Durrant-Whyte: I'll talk to you afterwards. 20 years from now.

Pascal Van Hentenryck: I should never talk after Toby, but anyway, 20 years from now? I agree with Toby. Energy is not going to be an issue, that's going to be solved. I think the big changes that you're going to see are in education and healthcare. Those are two things that everybody cares about and are major expenses on any budget on any country. Education is going to be fundamentally different. You will not recognise education. I think we're going to go back into more one-to-one education, offloading most of the learning offline and making sure that it's fun again. When Montaigne[6] was young, he had fun learning. We have to restore the fun in learning and we have to have more one-to-one breakout rooms interactions. That's one.

Healthcare is going to change completely — the way we deliver healthcare, the relationship with the doctors, and all the advances that are going to come from genomics and cell biology and so on. It's going to be completely different. I'm looking forward to that. It's going to be much more affordable. The system is going to completely change. It's going to be much more personal, much more dedicated. It's going to be based on your data. Of course, there are protections to put into place such that you don't penalise people who may not have the best genes, but I think it's going to be completely transformed.

I love driving but I don't think there will be any cars. Even though we are 99.9% there and we may never get to 100%, I think we will be accepting the fact that there are no cars and there is no congestion and we can move anywhere. You'll get around with micromobility for the first and last mile, and the rest is going to be automated in a safe fashion. Scooters. Especially in Australia, the weather is good.

The other thing that I think we need to take care of — and this is aligning to what Toby has said — is that we need to rethink that if we survive, we have a political system that is actually taking into account the human condition and making sure that we don't live in a virtual world, because evolution has made us live in the real world, in the physical world. We react to body language. We have inhibitions, when we interact in the physical world, that we don't have when we interact in the virtual world. We have to make sure that kids spend more time talking to each other, face to face. There is more communication. We'll have more time on our hands, most probably. We have to make sure that people can devote attention to arts, to sports, to anything that they will see as valuable.

Hugh Durrant-Whyte: You brought up a very important point — one of the most important points — which is: AI offers the opportunity to make things personalised, whether it's health or education or transport and so on. That in the end will deliver real benefits for individuals and will help, I think, with the inclusion problem and diversity in everything else. I think that's a real plus.

Pascal Van Hentenryck: So when I say "if we survive," I think there has to be an inclusive way of integrating AI. I work with Panama, and we have to make sure that AI is going

6 Michel de Montaigne (1533–1592) [Ed.]

to be for the city and it's going to be for the rural areas. If you don't do that, it's going to be a disaster. That's what I meant.

Hugh Durrant-Whyte: Jie, you're the last one up here. 20 years from now —

Jie Lu: Okay. My previous response was: Industry 5.0 will be implemented between 2030 and 2035. About 10 years from now, expect two key technologies: one (I mentioned) is the collaboration between humans and machines. At that time, I believe AI's software and hardware will be helping everywhere and you will be used to working with AI.

I'll give you an example. You take Sydney Trains to travel. In NSW, we have 240 stations and over 500 platforms. Every platform has a screen that my team and I developed. From early this year, the screen shows the eight carriages of the train coloured yellow, red, green. Green means more room available, red means full. By using machine learning, we have a data-stream prediction and we can use a fuzzy model to consider morning, afternoon, Saturday, Sunday, and weekdays and public holidays, and even sometimes special events.

This is real-time prediction. The train is irrelevant. We now have a prediction about Central. You have been used to enjoying the benefit from AI because this significant benefit is every traveller's experience. Another very important function is that it will significantly reduce the cost of maintenance of carriages. Maybe you don't realise it, but this is the main purpose of Sydney Trains. That is one part. After almost 20 years, we will be used to working with AI collaboratively.[7]

The second key technology is personalised AI, in personal education and personalised medicine. What I'm doing at the moment is finding associations between human genes and diseases; that's all about personalised health recommendations. Different persons should have different types of healthcare programs because we are all different. How to provide personalised medicine, Medicare, and that is after 10 years.

Moreover, the university will become an AI university. Why do we need you and me to mark assignments that AI can do? Every student should have personalised education. Different students have different goals, different targets. We organise different assignments, different materials, different books to read, different lectures to attend, different personalised tutorials. Why not? That's what I'm thinking after 10 years or 20 years.

Hugh Durrant-Whyte: Good. I still think there's a role for proper mentorship. I think there's a lot of personalised everything, but there's also something which is more about the human element, which is going to remain important. This is my opportunity to open it up to the audience.

Q1: There has been a lot of damage done and we've got the most polarised political and social system we've ever had in the Western world. I blame AI for a fair bit of that, in the way that it reinforces and creates echo chambers in the way that people communicate through all the on-line systems that exist today. When will AI start to reverse that, so that when you have an echo chamber, instead of having an echo, you actually are challenged so that you are opened, that

7 It's a "go" for innovative AI that revolutionises commuting. https://www.uts.edu.au/news/2023/11/its-go-innovative-ai-revolutionises-commuting (Nov 2023). [Ed.]

your perspective is opened, not closed, that you have more socially enabling attitudes developed through AI, not reinforcement of negative attitudes?

Toby Walsh: You've put your finger on a real problem: I dislike listening to music on Spotify because it only plays me the songs I liked, which means I start to dislike them because it's played them so many times. It doesn't challenge me with new music that will actually move my musical taste. Similarly, with my politics: if we live in an AI-filtered bubble that reinforces our views, it doesn't challenge us, doesn't confront us. That was because it was commercially successful for engaging us.

We're doing the same again with these language models, which is why they are the worst possible thing is to sit down and have therapy with: because they never say, "Oh, Toby, it's three o'clock in the morning. You've been talking to me for four hours. Maybe you should go to get some sleep." Because they actually just want to sell you more tokens, so they only say open-ended questions. "Oh, that's very interesting, Toby. Let's talk about this more." We're going to repeat it if we're not careful because the incentives are actually the wrong incentives. They're not actually worrying about meaningful engagement.

Pascal Van Hentenryck: If you look at history, when the printing press came about, there was exactly the same negativity. It was used for creating tensions between different communities. The key point is: can we resolve this and what are the mechanisms to do this? We have been there in history and we have solved it. It's more complicated now because the technology is much more powerful, but I think we really need to think about this correctly.

Toby Walsh: It's not a technical problem at all, it's a regulatory problem. We have to set up the institutions. Actually in Australia we have one great thing. We have the eSafety Commissioner,[8] who can say, "Okay, no, this is how you're going to actually have to regulate these things." The thing that fills me with some hope is that we're a little bit more on the front foot. With social media, it's after the event. We've seen the harms that happened. We already have a generation of young children who, sadly, are much more anxious, have more body image issues and have paid the price of that natural experiment. At least with AI, we're now asking these questions and we're just at the beginning. I'm hopeful in that respect.

Q2: What are your thoughts on edge intelligence and inference models? These days, you can build little smart devices and you can build that inference model after distilling the model from a larger cloud service, and then that model has the same capability to do things at the edge, so you start getting intelligence at the edge. What kind of effect is it going to have?

Lina Yao: That's a good question. To generalise the intelligence to the edge devices, the mainstream solution is federated learning. Usually, we train the parameters on the central server, but only pass the well-trained model parameters to the edge devices in the edge network. In this way, the edge devices don't need to spend computational resources to train the large-scale parameters. The edge devices can collect some data locally and

8 Inman Grant J (2024) How a single letter changed the world: W×3 — the World Wide Web (we weaved), *JProcRSNSW* 157: 266–284. https://www.royalsoc.org.au/wp-content/uploads/2024/12/157-2-08Grant.pdf [Ed.]

also maintain the data locally; the data they collect from each local device doesn't need to be transmitted to the central data. In this sense, the private data can also be very well preserved locally. In this way, only the parameters are transmitted between the server and the edge devices, and each edge device will also form a dynamic network to exchange and integrate their parameters and inference results to the server to keep synchronised.

Q3: I've got two threads in the conversation about education. You can go too far with personalisation of learning. I think a lot of what we've learned about learning over recent years is that it is fundamentally more effective when it's done in a social situation. I also want to pick up on the comment that Pascal made this morning about how our relationship with knowledge is going to change. I wondered if you could weave those two threads together to talk about the evolution of learning.

Pascal Van Hentenryck: The first thing is that people have to understand what ground truth is — we've talked about that. It's very important that we build what are called authority networks — things where we can go and be sure that the data is actually representing true facts and true knowledge which have been vetted. That's the first thing. Education is going to change completely. I agree with you that having groups of people learning together communally is much better. This mentorship relationship is what humans have been doing since we became a species — this is very important. This is how we have evolved, the fact that we communicate with each other, that we share knowledge and so on.

We have to rethink this, and think a lot less about evaluating people to make them learn, and make it fun again. We can do this, and I think we are going to be forced to do this, because of the budgets and the cost of education. This is a way to rethink everything from scratch. I really look forward to that. COVID for me was very helpful because I had to completely change my classes. I built a studio at home, I built videos, I had very professional videos, but when I interact with the students, it was one-on-one. This made a huge difference for that particular set of students. We can actually replicate that in a physical infrastructure. If you organise things as breakout rooms, but in a physical space, then you can have this mentorship relationship and people working together.

Toby Walsh: I think we have to be very careful that AI doesn't decide what is true for us. We're already going down that road that when you do a Google search, it's the AI overview that tells you, and most people only read the AI view. I mean, it's frightening to see Grokipedia, where Elon Musk is giving his version or his algorithm's version of what is true. We might easily, if we're not careful, end up in a world in which AI is generating the things that we believe in, whether they're true or not.

Hugh Durrant-Whyte: That is interesting. Yes, I agree. I think the whole area of education in the future is going to be a big challenge, but a real opportunity too, to get some different things right.

Q4, Judith Wheeldon: In the last hour, we have got to what really matters. I think it has been a terrific day. I've learned a lot, but I would like people to do a little mental experiment: how many times has the word "engineering" come up? How many of you are engineers or teach engineering? Yes, just about every hand has gone up. How fundamental is engineering to the whole area

of AI? Obviously, the answer is very, very fundamental. That's good: I'm not criticising that in the least, but the most important thing that I want to put forward and ask for a reaction to is something that is fundamentally different from the discussion we've had about AI — that is education.

The reason it's different is because it is not engineering. It is the humanities. It is what makes us human. And I, as a still-active teacher, see very, very clearly some of the things that are happening. This is really a follow-on from your very well focused question.

It is important to use AI for personalising education to some extent, but that's just fancying up what we've already been doing. What I maintain that we've been doing, at the same time without noticing, is deskilling and dehumanising our students. I'm talking about preschool, kindergarten, primary school, high school, university — say whatever you like — people are older, and different issues come up. Students today are looking for someone else to do the work for them. The HSC helps in that. It's a school ground for how to get someone else to do your work, but vocabularies are shrinking. AI is being used to write essays, to submit what you've written and get it prettied up, or to see if what you've written is good because you don't have judgement yourself to decide if your work is good. I believe we are fundamentally taking skills away from young people.

One of the reasons that's happening is because the teachers don't have the skills themselves — because they don't know, because they haven't been taught. They're working hard. I'm not criticising the teachers. It's a fix we've gotten ourselves into. How can we look at the human values of education, stop engineering, start developing young people and giving them the confidence to learn?

Hugh Durrant-Whyte: Thank you. I think you made some very valuable comments, but I'd like to also change the question if I might, and recognise that we will or do have these challenges. The issue is — and I'll let you speak to it, Jie — how do we actually do this in a way which is going to end up with those human type of results?

Jie Lu: For this, consider the AI university after 10, 20 years, with human interactions. Consider team assignments: if you have 10 teams, maybe each could have different assignments to meet the requirement about their future jobs. We can still have one-to-one people, people-to-people communication. More importantly, because we have online systems, you can not only talk to people in your classroom, you can talk to people in the US, in Europe, Belgium. That is, you have more friendships to build relationships.

Hugh Durrant-Whyte: Okay. So I might get us all to say something and finish on this. Go for it.

Toby Walsh: I sometimes jokingly worry that in 20 years' time, the machines are smarter than us, not because the machines are smarter, but because we're more stupid that we've dumbed ourselves down, so I do sympathise with many of your concerns. I think that this is an opportunity for us to rethink why we have education. What is the purpose of education? A lot of what I think we need to do is actually quite old-fashioned, which is to go back. It disappoints me so much as a scientist, as an engineer, to see the way that we've devalued the humanities, because they teach many of the things that

I think are really important. The critical thinking skills, the communication skills, those are the things that humanities have always taught. I wish I had had more of that when I was studying my science and engineering, but I wasn't allowed to choose it.

I think that we have an opportunity now to think and to use the tools in ways that can personalise and improve. I'm incredibly jealous of the tools that my daughter has. She started studying for the HSE. She was sitting in the back of the car the other day and her computer was asking her questions in French and she was replying in French. I remembered when I was revising French, I didn't have anything I could practice my French on. I'm incredibly jealous of the tools that are available. But, equally, I'm very aware that we are now starting to think, well, one of the purposes of education is work. Work is going to be different and we need to think what critical skills will be needed. We also need to think about the bigger purpose of education, which is to prepare us to achieve things in our life, to be good citizens. That may mean going back in many respects to the things that we used to do.

Lina Yao: I understand the concern. My daughter is also in high school. My strategy is I allow her to use AI, but under my supervision. She must do her homework without AI's assistance. At her age, the child definitely needs supervision and guidance. But later I think AI is more like a companion, not stealing our skills from humans. It's more like a companion that can work with us and learn with us together to automatically improve us. With the appropriate guidance and supervision, children should be exposed to an AI environment.

Hugh Durrant-Whyte: Thank you. And, finally, Pascal: start with Pascal and finish with Pascal.

Pascal Van Hentenryck: A couple of points. The first one is with respect to what is happening in engineering schools and computer science. I ran an experiment two years ago using all the assignments in my class using ChatGPT — programming projects. ChatGPT would get 50% correct; the other 50% were too difficult. So what did I do? I just took the first 50%, made them harder and made the other 50% even harder. Now, students couldn't use the AI tools and they really had to start thinking. This is an opportunity. We have an opportunity to rethink and to make sure that the students are actually thinking about what they are doing and using the tools that they have at their disposal to actually solve problems that are more complicated. That's number one, and this is purely engineering.

I come from Europe, so I was educated in a European university, which meant that for the first two years I had to study law, sociology, psychology, philosophy, in addition to statistical probability and so on. I value that education tremendously.

Two things that came to mind just now. The first one is that the Law Panel was amazing because, in a sense, the way I was thinking about this is that they were describing the law as it has been for centuries. While no technology can actually benefit from that, it can also inform what we can do because the new tools that we have could actually do fairness in a better fashion or equity in a better fashion. I think there are a lot of things that we can do if we have a broad education and merge the two areas.

In the Health Panel, the general practitioner said something very interesting.

They actually adopted AI to become more productive and see more patients, because that's how they get revenue. But in the end, it didn't work that way: the doctors spent more time talking to their patients. What happened was that the patients were more happy, the doctors were more happy. What I want to focus on is when we build engineering systems, when we build computer science systems or AI systems, we can still choose what the objective is. We can choose the way we are going to build this system and what it is going to be — not the return on investment but the overall goal of the system: we need to think about that and how we are going to deploy AI, so that it matches the value of the society we live in.

Hugh Durrant-Whyte: I want everyone to thank the wonderful panel.

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, pp. 143–144. ISSN 0035-9173/26/010143-02

2025 Royal Society of New South Wales and Learned Academies Forum: "AI: The Hope and the Hype"

Closing Session[1]

Hugh Durrant-Whyte, NSW Chief Scientist & Engineer

Hugh Durrant-Whyte

This is one of a series of events that the Royal Society of NSW has hosted over 11 years. They've always been amazingly challenging, bringing together a great diversity of people. It's just a fantastic place to be, and I have to thank the governor, Margaret Beazley, because she offers her House, she offers her venue, the wonderful gardens, and everything else to really top off an event like this in the sense of making everyone come together and for everyone to value, I think, everyone's opinion. That's been great.

I will say I have a personal love-hate relationship with AI. I think I might just mention this. My previous employees laugh at this, because I have spent 40 years essentially working in AI, building some of the most advanced robotics and autonomous systems there ever been. The Pilbara, Port Botany, you name it. Basically, I've been involved in all of these things.

Yet, people who know me well know that I do nothing at home that has anything to do with AI. I don't even own a smartphone. I never use Google. In fact, I spent half an hour yesterday removing Copilot from Windows. (And I will tell you how to do it if you ask.) I have deep suspicions about that kind of intervention in my life.

I read books. I talk to people one-on-one. And I think I'm getting a bit old in some sense, old in the mind — not about AI itself, the technology, because I think applications and science and everything else — but old in the sense that I'm just not someone who really enjoys the human condition. One thing I'd leave you with — and I think this came out in the last talk — is that the technology will genuinely assist and build a better humanity going into the future. But we do need to work on making sure that happens. We need to make sure that the benefits in education, in health, and law, and all of these other areas that have been discussed today, are genuinely realised.

We need to understand the role that AI will play and the role that it won't play, in areas like mentorship, reading a book ,or even trying to edit the document I'm writing. We need to basically understand how that happens. In the end, discussions that we've had today are critical for us all understanding what's possible, what perhaps is not possible, and what shouldn't happen, as well. I think that's important.

I think it has been a fantastic day all round in doing that. I want to end by thanking a few people. I've got to thank Sally Cripps first, my wife, who has been talking about this to me on and off for the last year and how it'll all be fine on the day. And it has been all fine on the day, as it always is whenever Sally's involved in these things.

I also want to thank all the other people who have been part of the committee, Lindsay, Vince, and Graham Town. There's

1 This is a transcript of a presentation, available at https://www.youtube.com/watch?v=HOBgg7YzT_g&t=20s

a lot of work that goes into organising these events and making sure that it all happens. I want to thank Pascal for coming back to Australia all the way from Atlanta for us to benefit from your wisdom. I'm so sorry that we lost you in the first place. It's one of those things that we really should say. I want to thank again the governor, Margaret Beazley, again for letting us use this wonderful venue that really makes all the difference to these events. Thank you everyone also for coming.

Christina Slade

I again thank the governor, and thanks, Hugh, who as the chief scientist has actually helped fund the event, as well as lending us his wife and running those fantastic panels at the end, which were very, very dynamic. I want to thank Lindsay Botten and the team, Vince and Graham, and particularly Davina as well. People have been working very hard, and I think it has been an extraordinary day.

Just trying to think our way through it: we began with Her Excellency, placing a quite wise set of understandings about how we need to think about governance in this space. We've almost come full circle. Victor Dominello urging us to move faster, particularly in government. Pascal made this remark, which really struck me. He said he wouldn't, as an academic or a corporation, have released ChatGPT. Nobody commented on it at the time. I asked him about it later and he said he meant it: it was not ready.

These large language models, which we've now heard critiqued and thought about quite a lot, are a bit of a worry. We're moving helter-skelter towards where we're going. But we have to think; we've listened to the legal panel, which was a fantastic panel, as Pascal said, and somebody commented that to have a group of women talking about the law was quite extraordinary. We got to that notion, which Sally raised quite sharply in her session, about the human. If you recall, Jill Bennett talked about Gwenda, the dementing woman, who was so reliant on her AI robot. After all, what human could do that?

What's so great about being human? We've got to this point now where we have to think through what it is that's great about being human and how it's going to work out. And also what's left of being human, if, after all, the programs, the algorithms know more about our own likely intentions and future actions than we do ourselves. Not only do they know more about our bodies, they know about what we're going to do next. It is a very difficult space. I think we've had positive news: EduChat, the banks working very hard to help us. We've heard about agriculture, really exciting, and we've heard about how it can go wrong. I'm not in quite such despair as Judith about the loss of the human in humanities, but I think we've got to do some pretty hard work thinking about what being human is now.

I want to thank you too and everyone who spoke, thank you very much and for everyone's attention. It has been an extraordinary day, Sally, and I hope you're actually going to get away and have a weekend at the farm with your husband, because I fear that you might not have had one for a while. Thank you all very much.

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, pp. 145–146. ISSN 0035-9173/26/010145-02

Thesis abstract

Formulas of Chineseness: tracing daigou between Australia and China

Jessie Liu

Abstract of a thesis for a Doctorate of Philosophy submitted to the School of Sociology, The Australian National University, Australia[1]

Over the past ten years, the buying and selling of infant formula across national borders has attracted concern and controversy in both Australia and China. This practice is commonly called daigou, a Mandarin portmanteau of "to buy" (*gou*) and "on behalf" (*dai*).

This thesis examines the infant formula daigou trade between China and Australia, studying both the work involved in doing daigou and the imaginaries that shape and are shaped by it in both countries. The study challenges the dominant accounts of daigou by telling different stories about the trade: stories that critically articulate how gender, migration, labour, and racialisation have shaped it. To accomplish this, this thesis takes an interdisciplinary approach, bringing literature from Asian diaspora studies and China studies into conversation with sociological scholarship on racialisation, migration, labour, and gender, as well as science and technology studies. Empirically, it draws upon interviews with daigou shoppers in Australia and buyers of infant formula in China, participant observation of daigou shops in Australia, and discourse analysis of news media, infant formula advertisements, and scientific and policy literature from Australia and China.

The thesis opens with an analysis of Australian news media to show how daigou has been constructed as a foreign Chinese phenomenon happening *to* Australia (Chapter 3). Part 2 (Chapters 4, 5, and 6) draws upon interviews with daigou shoppers to argue that daigou should instead be understood as an Australian migrant labour practice. It demonstrates how the trade's emergence and proliferation have been shaped by migrants' challenging experiences of Australia's gendered and exclusionary labour market, and how their labour has generated profits for companies based in Australia.

Part 3 (Chapters 7, 8, and 9) examines the imaginaries that underpin — and are reproduced by — the daigou trade. First, by analysing the labour of shoppers and infant formula advertisements, it shows how the trade has reproduced and exported colonial imaginaries of Australia. It then discusses imaginaries of Australian infant formula as "clean and green," highlighting how these understandings reflect middle-class Chinese mothers' negotiation of China's development and pollution. Finally, it studies the new biologised imaginaries of Chineseness that are produced in anti-daigou discourses of infant nutrition in China. In particular, it focuses on the idea of "Chinese formula

1 ANU 2025 JG Crawford Prize for Interdisciplinary

for Chinese babies" and its deployment in the cultivation of a racialised commodity.

Through its examination of the intersecting practices and imaginaries that shape the daigou trade, this thesis contributes to critical conceptual and empirical work on Chineseness, racialisation, markets, migration, and transnational practices and subjects.

Dr Jessie Liu
Research Fellow
Emerging Technologies Research Lab
Monash University

E-mail: jessie.liu4@monash.edu

URL: https://hdl.handle.net/1885/733713802

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, pp. 147–148. ISSN 0035-9173/26/010147-02

Thesis abstract

The Relevance View: aggregation and the structure of reasons

Kirsten Mann

Abstract of a thesis for a Doctorate of Philosophy submitted to the School of Philosophy, The Australian National University, Australia[1]

When allocating scarce resources, decision-makers must weigh up the competing claims of the various individuals whose wellbeing is at stake. Many of us accept that there are cases where a number of weaker claims can, in aggregate, outweigh a stronger claim. But we resist the aggregation of weaker claims in other cases. For example, aggregation seems appropriate in a case where one person's claim to the treatment of a broken arm competes with a million people's claims to the treatment of broken fingers. Aggregation seems inappropriate, by contrast, in a case where one person's claim to life-saving medical treatment competes with a million people's claims to the treatment of mild headaches.

The Relevance View is an influential account of our moral duties to assist others in cases like those just described. The view occupies a middle ground between aggregation and non-aggregation, affirming and rendering coherent our intuitions about these cases. It allows weaker claims to aggregate to outweigh a stronger claim just when the claims, compared pairwise, are sufficiently close in strength that the weaker is "relevant" relative to the stronger. No number of claims against mild headaches can outweigh a claim against premature death, for example, because a headache claim is so much weaker than a life claim as to be irrelevant to it.

This thesis explores and further develops the Relevance View, and defends it against prominent objections. My starting point is Alex Voorhoeve's pathbreaking presentation of a version of the Relevance View that he calls "Aggregate Relevant Claims" (ARC). Over the course of the thesis, I develop new versions of the Relevance View that avoid ARC's shortcomings and that extend the view beyond the class of simple cases to which it has typically been restricted.

In the first two chapters, I defend the Relevance View against criticisms that have been advanced in the literature. Chapter 1 addresses the objection advanced by several critics that the non-aggregative intuitions that motivate and justify the view are unreliable due to the unimaginably vast numbers involved. Chapter 2 addresses Patrick Tomlin's arguments that ARC suffers from serious failures of internal logic and fails to apply to cases resembling real-world healthcare allocation problems.

Chapter 3 sets out an objection to ARC that I argue cannot be addressed within the confines of Voorhoeve's approach. If the Relevance View is to offer plausible guidance in non-binary choices, I argue, its proponents must accept a "binary contrastive" account of the moral reasons associ-

1 ANU 2025 JG Crawford Prize for HASS

ated with claims for assistance. Chapter 4 explores whether the Relevance View could be applied in population ethics, to choices that affect which future populations come to exist. I develop a choice rule for populations where relevance serves as a limit on the aggregation of possible people's wellbeing.

The Relevance View offers a compelling account of our duties of rescue. The chapters of this thesis together demonstrate the view's robustness to the criticism it has received, and its considerable promise as an account of our moral obligations beyond the simple binary cases that have figured most prominently in this literature to date.

Dr Kirsten Mann
School of Philosophy
The Australian National University

E-mail: kirsten.mann@anu.edu.au

URL: http://hdl.handle.net/1885/314318

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, pp. 149–150. ISSN 0035-9173/26/010149-02

Thesis abstract

Indium phosphide nanowire arrays for gas sensing applications

Shiyu Wei

Abstract of a thesis for a Doctorate of Philosophy submitted to the Research School of Physics, The Australian National University, Australia[1]

Semiconductor nanomaterial based chemiresistive gas sensor is one of the most promising technologies for the development of next-generation chemical sensing, with applications ranging from explosive detection and environmental monitoring to industrial safety, healthcare and the Internet of Things (IoTs). Recently, there have been growing efforts in developing room temperature chemiresistive gas sensors with high sensitivity, selectivity, long-term stability, and low power consumption. In this thesis, we design and demonstrate indium phosphide (InP) nanowire (NW) array based chemiresistive gas sensors, as a new and promising platform for achieving high-performance gas sensing for real-world applications.

Firstly, by carefully engineering the NW array geometry (i.e., diameter and pitch), we demonstrate a NO_2 sensor with superior sensitivity (limit of detection 3.1 ppb) at room temperature, with outstanding selectivity and long-term stability. Kinetic analysis and electrical simulation reveal the array geometry correlated sensing mechanism. To minimise the device power consumption, a novel axial p-i-n homojunction self-powered photovoltaic (PV) NW NO_2 sensor was further designed through numerical simulation and experimental optimisation. The fabricated InP NW array PV sensor achieved 84% sensing response to 1 ppm NO_2 with a record limit of detection down to sub-ppb level even under <5% of 1 sun illumination. With high fidelity to the nature lighting conditions, the sensor was integrated onto a commercial microchip interface for self-powered, dynamic monitoring of on-field motor vehicle exhaust.

Not only for oxidising gas such as NO_2, the InP NW array can also be fabricated into sensitive sensors for reducing gases such as acetone. By incorporation of chitosan as a surface-functionalisation layer and a Pt Schottky contact for efficient charge transfer and photovoltaic effect, self-powered, highly selective acetone sensing is achieved. The sensor has exhibited an ultra-wide acetone detection range, covering concentrations found in the exhaled breath of healthy individuals (300–800 ppb) as well as those at high risk of diabetic ketoacidosis (>75 ppm). Our acetone sensor was then integrated into a breath testing apparatus, named Ketowhistle, and proven to be promising for diabetes monitoring and diagnosis.

Finally, a wearable multipixel InP nanowire array sensor was fabricated on a flexible substrate. A signal decoupling strategy was designed and implemented through nanowire size control, surface-functionalisation and electrode metal selection. This approach has led to the demonstration of a multipixel

1 ANU 2025 JG Crawford Prize for STEM

sensor device consisting of sixteen sensing elements with integrated functionality, including NO_2 and acetone sensing, as well as simultaneous pulse and body temperature measurements, demonstrating potential as a personalised device for health and environmental monitoring.

The thesis work indicates that III–V compound semiconductor nanowire array presents a promising sensing platform for the development of high performance, miniaturised, low power consumption, multifunctional on-chip sensing system for future large-scale IoT technology.

Dr Shiyu Wei
Hefei University of Technology
CHINA

E-mail: shiyuwei@mail.hfut.edu.cn

URL: https://hdl.handle.net/1885/733716229

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, p. 151. ISSN 0035-9173/26/010151-01

Heather Goodall (1952–2026) AM FRSN FASSA

Lindsay Botten FRSN

Email: lindsay.botten@outlook.com

Professor Goodall (1952–2026) was a pioneering Australian historian whose work transformed understandings of Indigenous political history, environmental change, and cross-cultural relationships. She passed away peacefully on 29 January 2026 at the age of 73.

Heather Goodall's academic journey began with the University Medal in History from the University of Sydney in 1975, followed by a groundbreaking PhD in 1982 on Aboriginal communities in New South Wales. This research shaped her lifelong commitment to collaborative, community-centred historical practice.

Over her long career at the University of Technology Sydney, Goodall became a leading figure in Australian historiography. Her research was widely recognised through multiple ARC Discovery and Linkage grants and influential publications. On her retirement from UTS in 2016, she was honoured with the title of Professor Emerita.

Her major publications — *Invasion to Embassy*, *Isabel Flick*, *Rivers and Resilience*, *Beyond Borders*, and *Georges River Blues* — combined rigorous scholarship with deep partnerships, often amplifying voices historically marginalised in mainstream narratives. Her work extended beyond Australia, tracing intercolonial and transnational networks that reshaped understandings of activism and identity.

Professor Goodall received numerous honours, including the NSW Premier's Australian History Prize and the Magarey Medal for biography. She was elected a Fellow of the Academy of the Social Sciences in Australia and a Fellow of the Royal Society of NSW in 2014. In 2024, she was honoured as a Member of the Order of Australia for significant service to tertiary education and to the Indigenous community. She leaves a profound legacy as a scholar, mentor, collaborator, and advocate whose influence will endure across generations and will be greatly missed by colleagues, students, friends and all who were inspired by her scholarship, activism and kindness.

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, p. 153. ISSN 0035-9173/26/010152-02

Stephen Craig Hill (1943–2026) AM FRSN FTSE

John Hardie AM FRSN

Email: jrhardie1@bigpond.com

A former Councillor of the Royal Society of NSW and a member of the Council of the Royal Societies of Australia (RSA), Stephen Hill joined the Society in 2015 and brought his enormously diverse international experience and knowledge to both organisations. In recent times, Stephen made major contributions to the RSA's Pre-conditions of Well-being project.

Stephen graduated with First Class Honours in Physical Chemistry from the University of Sydney in 1964 and then taught chemistry at Sydney and Melbourne Universities. After qualifying in Australia's first MBA program in 1965 with a thesis in economics, he was subsequently awarded Australia's first PhD in Business Administration, completed in 1967 at the University of Melbourne. He was subsequently awarded Fulbright Foundation and Ford Foundation Intersocietal Fellowships and appointed an Assistant Professor in Sociology, Business and Industrial Engineering in parallel positions at Northwestern University in Chicago.

After being awarded a Nuffield Fellowship, he took up a Senior Lectureship in sociology at UNSW and then became Australia's youngest full Professor at the age of 30 as Foundation Professor of Sociology at the University of Wollongong in 1974. He subsequently set up three Research Centres — Multicultural Studies in 1978, Technology and Social Change in 1983, and The Centre for Research Policy, a Centre of Excellence of the Australian Research Council, in 1990.

Stephen was the Principal Director and Ambassador for South-east Asia with UNESCO from 1995 to early 2006. Based in Indonesia and Paris, he served as the UN's Regional Director for Science for Asia and the Pacific, and as Representative or Ambassador to Indonesia, Malaysia, Philippines, Brunei Darussalam, Timor Leste and Special Representative of the Director-General to Singapore. He was responsible for major UN science, education, media freedom, culture, world heritage and peace initiatives and reforms in the region, as well as being formally mandated by UNESCO's Director-General to reform and decentralise the entire UN Agency globally from 1999 to 2002.

Stephen founded and chaired International Networks in Science and Technology Policy (UN), Social Sciences and Humanities (Pacific Science Association), co-chaired Australia's UNESCO Science Network, and was then involved in the founding of APEC, where he chaired the APEC Human Resources and Technical Education Network and Programs. Additionally, he was central to the establishment of Australia's science relations with China, Indonesia,

Korea and Malaysia, part of the basis on which his Membership of the Order of Australia was awarded.

During his time working in UNESCO, he was responsible, as UN Representative, for negotiations for the release of his staff taken hostage by Freedom Fighters into the jungles of West Papua in 1996. He was also responsible for managing the UN response to the May 1998 Democracy Revolution in Indonesia — escaping his house at 2 AM through mobs and fires, evacuating all staff but staying to report to the UN Security Council, and in the aftermath launching a wide-ranging program to reform the law and establish media freedom throughout Indonesia. Examples of work in other Asian countries include: the development of a program to bring literacy to 7,000 village women in the war zones of Mindanao, southern Philippines; reforming Mongolia's entire science system out of its Soviet-inherited past and into contemporary relevance; and introducing the internet to North Korea.

Following retirement from UNESCO, Stephen was appointed Honorary Professorial Research Fellow at the University of Wollongong in January 2006 and subsequently created an Emeritus Professor at that university in 2008.

Stephen published extensively across diverse fields — global economics, social dynamics, and the values driving social change — and has been translated into six languages. His bibliography includes over 20 books (many of which explore his own real-life adventures) as well as several hundred research papers.

In addition to his writing, Stephen was committed to engaging with communities, through public talks on contemporary issues. His storytelling ability, combined with his deep understanding of global affairs, made him a compelling voice in discussions on humanity and social change. To that end, he proposed the establishment of a "speakers forum" for the Royal Society of NSW.

Stephen was honoured with awards from Australia, Indonesia, Vietnam and the Philippines for his contributions to Australia and the Asian region — including Member of the Order of Australia (1996) and Australia's Centenary Medal (2000). He was elected Fellow of several international academies and societies, including the Australian Academy of Technological Sciences and Engineering, the World Innovation Foundation, and the Royal Society of NSW.

Stephen was formally named "Purba" after the last King in a Ulos naming ceremony by the Batak Nation of Sumatra, Indonesia, and honoured by indigenous tribute ceremonies, including from the Dani tribe of West Papua, which ceremonially inducted him as an Honorary Tribal Chief.

Stephen lived and worked in the USA, the UK, Germany, France, India, Thailand, Indonesia, Malaysia and Australia. He was a visiting professor of Doshisha and Kyoto Universities in Japan, where he recently co-led an international symposium and writing program on creative economics.

Stephen's CV also includes a stint as a rock musician. During the late 1970s and early 1980s, he was a professor by day and a professional rock musician by night. Stephen's passions for community empowerment and human rights initiatives were a driving force in his life. He was indeed a veritable polymath with skills and influence spanning chemistry, economics and social science at a global level.

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, p. 160. ISSN 0035-9173/26/010154-07

Emma Letitia Johnston (1973–2025) AO FRSN FAA FTSE

Merlin Crossley FRSN

Email: m.crossley@unsw.edu.au

Photo by *Bahudhara*, licensed under CC BY 4.0

Emma Johnston's life exemplified what was possible for a smart and accomplished person to achieve in what was perhaps a golden age in universities in Australia. She was the product of a superb tertiary education system, and she used her learning and talents to give back and make our world better.

One can see a pattern of dedication and excellence at every stage of Emma's life, with her capabilities and commitment providing opportunities for leadership, and her leadership making significant contributions to education, to the environment, and to the advancement of society, particularly in the area of science, and the advancement of women and of inclusion more generally.

Her early years were spent in Williamstown, a seaside suburb of Melbourne. Like many Australian kids she enjoyed the outdoors, especially swimming, snorkelling, and sailing. Her parents were well-educated. Her father was an applied mathematician with a doctorate in chemical engineering. Her mother was also qualified in chemistry, but had broad interests, including painting. Emma was a gifted student. She attended one of the few selective public high schools in Melbourne, University High School, and ended up being Dux of School. From this early age she was already committed to conservation and reform, and campaigned against a carpark that threatened green space in the school, successfully. Today, one of the School Houses is named in her honour.

The pattern of seeking leadership, not through vanity, but through a desire to preserve and improve the world, continued at university. She studied science at the University of Melbourne, and she engaged in student politics. She was deeply committed to the student newspaper *Farrago* and ended up being elected as the President of the Student Union. Like many students, for Emma the campus was not enough, and she led efforts to improve human rights in East Timor, an international trouble spot at the time.

I have spoken to many who knew her when she was a student and they talked of her charm, talent, drive, focus, purpose, and above all an emerging charisma. As her capability and knowledge developed, this charisma would become increasingly evident and would allow her to engage, inspire, and influence throughout her career.

After graduating in 1997, with a first-class Honours degree in marine biology, that included other subjects, such as the history and philosophy of science, Emma embarked on her PhD at the University of Melbourne. She formally graduated in 2002, and even before the ceremony had been engaged as an Associate Lecturer (Level A) at the University of New South Wales (UNSW), in the Centre for Marine and Coastal Studies, in the School of Biological, Earth, and Environmental Sciences, known locally as BEES. She was to spend 21 formative years at UNSW.

Level A is the lowest academic rank, and often new junior staff are given heavy teaching loads and charged with updating the curriculum. Emma was no exception, and she "dived in" and set about reforming and improving the curriculum, engaging colleagues across the School. It is often hard to quantify and make visible the contributions that top academics make via teaching (whereas research can be measured by agreed though often contested metrics). Nevertheless, I would calculate that perhaps 10,000 students benefited from the teaching reforms that Emma introduced, and from the inspiring dedication to teaching that she embodied in the early years of her career. The whole culture of teaching in her School benefited.

She was promoted to Lecturer (Level B) in 2004, then went, as a Visiting Researcher, to the Marine Biology Association, Plymouth, UK in 2007. She was promoted to Senior Lecturer in the same year. Moreover, in 2007 she received a New South Wales Tall Poppy Award for excellence in both science and the communication of science. Her future would unfold along these lines.

I first met Emma properly in 2010 when I joined UNSW as Dean of Science, a role in which she would succeed me. The first thing I noted was her smile — she always lit up a room. She was always quick to smile and had a playful, energetic, and often cheeky demeanour. I've observed elsewhere that if Emma had been an animal, she would have been a dolphin.

The next thing I noted about Emma was her humility. She was a naturally modest person — I think the Australian culture often tends to embed humility in talented people — but fortunately she was already growing in confidence and in strength, as she enjoyed success after success. Thus, though modest and self-effacing, she would ultimately become a fearless academic leader and reformer.

2010 was an important year for Emma. By this time, she had developed a strong research profile. She became Deputy Director of the Evolution and Ecology Research Centre at UNSW (one of the many impressive local academic centres in the university), and she became an Australian Research Council Fellow. This was very important, as it provided time for research, and beyond that it marked her as a "research star." There are very few of these fellowships, so they serve as distinctions, almost prizes, that signal excellence. Emma's excellence, which was easy to see if you were at UNSW, was about to become increasingly widely evident.

The Sydney Institute of Marine Science was established in Chowder Bay, not far from Sydney Taronga Zoo, in 2005, making use of various buildings on the beautiful Sydney foreshore. It was supported by a collaboration between four local universities (UNSW, the University of Technology Sydney, Macquarie, and Sydney). Emma was

always a prominent scientist there; in 2012, she was formally appointed Director of the Sydney Harbour Research Program. This project was a dedicated and comprehensive effort to capture the biology, ecology and human interactions of our remarkable harbour, a place that boasts more species than the entire Mediterranean Sea. The work required a gifted and collaborative leader as it would involve multiple universities, government agencies, and our nation's oldest museum, the Australian Museum. A large number of influential people came to know Emma at this stage. Her contributions beyond the university were just beginning to unfold.

I want to list and then expand on the range of activities that underpinned Emma's impact. Her undergraduate teaching and supervision of research students and post-doctoral researchers, though perhaps less visible, should be mentioned first. Second, her leadership in formal academic science, which I will detail below. Third, her outreach and science communication, for which she won multiple awards. And fourth, her contributions through formal leadership positions, in universities, and on government and other committees.

As a local group leader at UNSW, Emma had a huge impact in her discipline. The School she worked in — BEES — is very highly respected, with a clear purpose and multiple internationally recognised research stars. They happily co-exist in distinct specialities, so the School is somewhat less prone to the ego-driven competitions that beset some academic environments.

But beyond that, Emma and her colleagues deliberately set out to generate a culture of support and inclusion in their School. The 40 or more post-graduate research students and research fellows Emma supervised attest to the culture of her group, as does the pride they had in seeing their work unfold. Their simple dedication to doing good science properly, often with very small budgets, necessitated the creation of equipment from household waste and repurposing old material lying around the department.

Specialised academic contributions often go unreported, perhaps because contributions to leadership and communication are more visible and more accessible. I will, however, list four of her most significant academic projects here.

First, she was an expert in ecotoxicology: how waste, particularly industrial chemicals in sediments, affects the biology of ecosystems. She showed that dredging up settled waste that had been laid down, long before many restrictions of chemical effluents were in force, could make matters worse. Her work on dredging and waste informs current practices on dredging in Sydney Harbour and beyond.

Second, Emma was an expert on invasive species. Australia has strict border security rules in an attempt to limit such invasions, but ships from abroad arrive every day. Emma noted that invasive species often gained footholds in the most polluted local environments. This is because the species had often pre-adapted to polluted conditions in more heavily polluted regions abroad. Her approach was unique, in that she relied on field observations and combinations of stressors to provide results that reflected the real-world situation. Her insights have ensured that monitoring for invasions focusses on areas where pollution is known, because these are likely to be the first places to be colonised.

Third, Emma studied the impacts of desalination plants. With annual rainfalls diminishing across the south of Australia, desalination has become a priority; across the world with population growth in areas with little rainfall, more and more such plants are being built. But desalination plants discharge high-salinity water into the sea, as they extract fresh water for civic purposes. The question was: would this high salinity water be detrimental to the local ecosystems? Emma's work showed a surprising result: no. Coastal ecosystems were quite resilient to high salinity through flows, but very fast flow rates could damage the environment. This knowledge now informs how desalination plants are planned, with saline discharge allowed, but flow rates carefully monitored.

Finally, Emma helped formalise systems for measuring the health of coastal ecosystems. Historically, labour-intensive observations and collections by individual researchers were necessary, with key indicative species being counted. Emma developed structured processes; significantly, she started using "environmental DNA sequencing," which allows the extraction of DNA from sea water samples and its total sequencing. This can reveal both the range and abundance of different species, large and small, and can give an almost instant and comprehensive snapshot of the ecosystem. Her approach, together with traditional sampling, can be particularly useful in detecting impacts of environmental events such as large bushfires, and Emma's work on their effects was particularly well received.

For her research, Emma won the 2012 Premier's Prize for Science and Engineering for Excellence in Biological Sciences (Environmental Science); in 2014, she won the Australian Academy of Science's Inaugural Nancy Millis Medal for Women in Science; and in 2015, the Society for Environmental Toxicology and Chemistry AU Mid-Career Medal. In 2016, she gave the Canadian Rivers Institute Distinguished Hynes Lecture, and in 2018, she won the Royal Society of New South Wales Clarke Medal for Zoology.

Most people who knew Emma would have known her, not through her actual science, but via her scientific outreach. In 2012, she was invited to be interviewed by the producers of Coast Australia, a documentary series that would be shown initially on the BBC and Foxtel. Her stage presence and knowledge were so instantly captivating that she was invited to feature as a scientific expert as the series unfolded. She was a writer and presenter from 2013 through to 2021, throughout all four series of the show. She managed all this while maintaining her teaching, her research, her administrative roles, and, of course, while supporting her family.

Emma won an Australian Museum Eureka Prize for Promoting the Understanding of Scientific Research in 2015. She was also a member of the team that won the 2018 Eureka Prize for Science Journalism for a one-hour documentary on the Great Barrier Reef that she wrote and presented.

All these highly visible achievements now increasingly led to invitations to apply and be appointed to important leadership roles. It is notable that in each case she assumed such roles with her usual view of conserving and improving the world, rather than, as is sometimes the case, to feed a burning ambition to be in charge.

From 2012 to 2016, she was a member of the NSW State Government Marine Estate Expert Knowledge Panel. In 2016, she was

the lead author of the Australian Government's "Australia State of the Environment 2016: Coasts" report, and later in 2021, the co-chief author of the entire report. She became a Director of the Great Barrier Reef Marine Park Authority in 2016, and she contributed to overseeing another great marine sanctuary as Deputy Chair of the Antarctic Science Foundation from 2017 to 2021. She was also a lead investigator on the Australian Research Council's 2021 Special Research Initiative, "Securing Antarctica's Environmental Future." She made invaluable contributions via both hands-on studies and leadership in Antarctica.

Many of her most prominent contributions to science strategy, policy, and the community came via her roles in Science and Technology Australia, a national peak body that represents over 100 different scientific societies and 200,000 professional scientists and technologists in Australia. From 2012 to 2019, Emma served as a Council member, then Deputy President, and finally as President of this organisation. She gave a speech at the National Press Club in 2018 about the importance of science and support for science in Australia, and she was instrumental in initiating and supporting a range of new projects.

Perhaps the most notable initiative is the Superstars of STEM (Science, Technology, Engineering, and Mathematics). She had a very strong appreciation of the importance of role models, and wanted to ensure that the next generation was inspired by seeing junior women and non-binary scientists succeeding in their professions. Emma firmly believed that "you cannot be, what you cannot see," or at least that seeing your destination helped to get you there. The Superstars of STEM programme started showcasing around 30 up-and-coming scientists each year and was so successful that it expanded in following years. She also worked very closely with Lisa Harvey-Smith, Australia's Women in STEM Ambassador, who did important work, not only on strategies to improve access but on the efficacy of those strategies.

Finally, I come to Emma's leadership roles in universities. She was promoted to Associate Professor in 2011 and to Professor in 2014. In 2015, she was made an offer by another top Australian university; I was charged with retaining her at UNSW. We discussed what might keep her here. Her primary concern was neither salary nor research resources: it was the opportunity to mentor and develop up-and-coming academics. This work lay primarily with the Pro-Vice-Chancellor Research; she explained to me how that was the job that could keep her at UNSW. I understood at once but explained that we already had someone in that role. She confidently said that she was sure we'd find a way. Indeed, through collegial discussions, we did find a way forward, and she served as PVC Research until 2017.

At that stage, the Dean of Science role had become vacant and I, as the former Dean, was again charged with identifying suitable candidates. Again, my conversation with Emma did not unfold as I had envisaged — I considered the role of Dean of Science at UNSW to be one of the best jobs in the sector, because science is such a strong faculty at UNSW — but this time Emma was hesitant. She was well aware of the commitments required and the need to balance bureaucratic heavy lifting with many ceremonial occasions, while also keeping in touch with her own research and teaching. But happily, she was eventually convinced

to apply for the role. She was appointed and served as one of the most loved and admired Deans UNSW has ever had — and it has had some very good Deans.

Now, Emma was one of the most visible scientists in Australia.

When the job of Deputy Vice-Chancellor Research at the University of Sydney became available, she was snapped up in 2022, and did excellent work in the role. I know that the injection of a new face with so many practical ideas and such a positive outlook was very much appreciated by many of my own former colleagues still at Sydney. In 2024, she was identified as the new Vice-Chancellor of the University of Melbourne, and returned to her alma mater to help guide that outstanding institution.

Again, Emma's time at Melbourne was appreciated by all those I have spoken to there. Melbourne has always been a leading Australian university, but morale had faltered during the many COVID lockdowns in Melbourne. Furthermore, student politics and protests, especially around the awful crisis in the Middle East, had become more and more visible and divisive.

It is hard to imagine anyone better equipped to restore positivity in Melbourne. As a former student leader, it was hard to frame her as an out-of-touch or self-interested manager. As a star researcher, she represented hope for those committed to building Australia's capacity in research and belief in science. As an inspiring, award-winning university teacher and science communicator, it was clear that teaching would be highly valued on her watch. And finally, as a superb mentor and supporter of women and disadvantaged groups, it was also clear that Melbourne would continue with its efforts to throw open its doors and provide opportunities for students of potential wherever it could. Morale at Melbourne soared.

At the excellent commemorative event held at Wilson Hall at the University of Melbourne, where Emma (and indeed I) had sat tens, if not hundreds of exams long ago, Emma's former students, colleagues, mentors, and even the Education Minister Jason Clare, spoke of her contributions and her capabilities, and recounted heart-warming and heart-wrenching stories of her life and achievements. Her husband Sam Maresh, an expert in renewable energy systems, spoke superbly, at this and other events held in Sydney. The then-Chair of the Australian Tertiary Education Commission and former NSW Chief Scientist and Engineer, Mary O'Kane, also spoke in Melbourne and later in Sydney. At the events were: many of Emma's mentors; Peter Steinberg, who had been inaugural Director of the Sydney Institute of Marine Science; Rob Brooks, who had set the standards for supportive lab cultures in his Evolution and Ecology Research Centre; her students, most notably Katherine Dafforn; and her colleague Kylie Walker from the Science and Technology Australia, who had helped found Superstars of STEM. Her children Amelia and Antonin and their friends also attended several of the events. There were others, of course, who had worked with Emma, or had been led by Emma, or had learnt from Emma.

Emma was a star and she was recognised as a star, being welcomed as a Fellow of the Australian Academy of Science, the Academy Technological Sciences and Engineering, the Royal Society of New South Wales, and she was appointed as an Officer of the Order of Australia in 2018.

She was taken away from us too soon, and we lost future contributions she might have made when she died from complications of cancer.

But her story was not taken away. It is a story of dedication to learning, of seizing of opportunities, of focus and hard work, of using leadership to give back, of inspiring through purpose, and of growing confidence through application and success. It is an inspiring story. It shows what can be done and what should be done. It shows the importance of teaching and the importance of science, and the importance of excellence. It shows also the importance of humanity and of putting support for others first. The achievements of any single day are often invisible — even the achievements of a lifetime can be hard to identify — but good work does have good impacts. Australia is a good place to live, and it has an environment which is still, in many places, like the Garden of Eden, and is the envy of the world. Emma saw all that from an early age and worked to try to ensure that in this garden the opportunities that it offers would be available to all and forever. She never stopped talking about saving the planet. We should also continue to learn and work to deliver what we can.

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, p. 161. ISSN 0035-9173/26/010161-01

Sir Anthony Mason (1925–2026) AC KBE DistFRSN FASSA HonFAIB KC

Lindsay Botten FRSN

Email: lindsay.botten@outlook.com

Distinguished Fellow Sir Anthony Mason AC KBE DistFRSN FASSA HonFAIB KC died on 17 March 2026, just short of his 101st birthday.

Sir Anthony Mason, a former Chief Justice of the High Court of Australia (1987–1995), was a towering figure in Australian law and one of the nation's most influential jurists, whose influence reshaped the nation's constitutional and common law. Renowned for his intellectual rigour, principled independence and dry wit, his judicial legacy profoundly reshaped Australian law.

Born in Sydney on 21 April 1925, Anthony Frank Mason grew up during the Great Depression and the political turbulence of the Jack Lang era. After leaving school, he enlisted in the Royal Australian Air Force in 1944, training as a navigator in Australia and Canada before the war ended. Returning home, Mason studied arts and law at the University of Sydney, graduating with first-class honours in both. He was called to the bar in 1951 and rapidly established himself as an outstanding barrister. In 1964, he was appointed Solicitor-General of the Commonwealth, before joining the New South Wales Supreme Court in 1969 and the High Court of Australia in 1972.

Initially regarded as a conservative judge, Mason's jurisprudence evolved markedly. As Chief Justice from 1987, he championed principled legal development over rigid adherence to precedent. His leadership shaped landmark decisions, including Cole v Whitfield, the Franklin Dam case, Australian Capital Television and Mabo, as well as the development of implied constitutional freedoms and national powers.

After retiring in 1995, Mason continued extensive public service in Australia and abroad. Sir Anthony Mason leaves a profound and enduring legacy in Australian jurisprudence. He is mourned by his family, colleagues and the legal profession he served with distinction, and the Royal Society of NSW, of which he had been a Distinguished Fellow since 2018.

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, p. 162. ISSN 0035-9173/26/010162-01

Keith Douglas Suter (1948–2025) AM FRSN

Lindsay Botten FRSN

Email: lindsay.botten@outlook.com

Dr Keith Douglas Suter AM FRSN, died on 10 December 2025 after a long illness. Dr Suter was a distinguished foreign affairs commentator, human rights advocate, and futurist whose career spanned more than five decades.

Born in London to modest circumstances, he left school at 15 to work at the Ministry of Defence, where mentors inspired his lifelong commitment to peace and justice. He later studied political science at the University of Sussex, becoming the first in his family to attend university. His activism began early, pressing British leaders on apartheid and UN obligations, while honing his oratory at Speakers' Corner in Hyde Park.

Arriving in Sydney in 1973 on a scholarship, Suter completed three doctorates and published 12 books, including the bestseller *Global Notebook: 50 Things You Want to Know About World Issues … but Were Afraid to Ask*. He worked for 32 years at Wesley Mission, championing causes from Indigenous justice to nuclear disarmament. He founded the Trinity Peace Research Institute and became a leading figure in scenario planning and futurism. His voluntary service was extensive, spanning decades with the Club of Rome, the Anti-Slavery Society, the International Commission of Jurists, the Red Cross, and the UN Association of Australia. He also participated in the Foreign Minister's National Consultative Committee on International Security for nearly 25 years.

Known for his wit and clarity, Suter was a familiar media presence, appearing on the *Sunrise* program of the Seven Network, delivering lectures, and writing widely. He was appointed a Member of the Order of Australia in 2019 and received numerous peace and service awards. Until weeks before his death, he remained active in commentary and preaching. He is survived by his partner, Jane Phelan, and brother Clive. His legacy endures as a communicator, educator, and peacemaker.

Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, pp. 163–168. ISSN 0035-9173/26/010163-06

Royal Society of New South Wales

Awards 2026

The Royal Society of New South Wales has long recognised distinguished achievements in various fields of knowledge through its Awards. Some are amongst the oldest in Australia while others are more recent. From its Act of Incorporation in 1881, the Society's mission has been to encourage "studies and investigations in Science, Art, Literature and Philosophy."

Nominations for all available Awards open on 1 July each year and close on 30 September. Awardees are announced by the end of that calendar year with formal presentations of their Awards in the following year. Almost all nominations require a nominator and a seconder. All RSNSW Awards are assessed relative to opportunity.

See the Awards page for all links, at https://royalsoc.org.au/awards.

Career Excellence Medals

Nominations for Career Excellence Medals are called for annually in the four categories described below. Please note that the call for nominations for all awards opens on 1 July and closes on 30 September of each year.

RSNSW Aboriginal and/or Torres Strait Islander Scholars Medal

Awarded for sustained, meritorious contributions to knowledge and society made by scholars identifying as Australian Aboriginal or Torres Strait Islander and conducted mainly in NSW. Recipients may be resident in Australia or elsewhere.

The application procedure for this Medal is described on the nomination form. Each nomination must comply with the conditions of award and will consist of a completed nomination form together with supporting documentation as specified on the form.

Note: When appropriate, this Medal recognises teams as well as individuals. Nominators are welcome to consult the Society for guidance before making a team nomination, noting that only one physical medal is presented.

RSNSW James Cook Medal — for lifetime career contributions

Awarded for the most meritorious lifetime contributions to knowledge and society in Australia or its territories made by an individual and conducted mainly in NSW. The recipient may be resident in Australia or elsewhere.

The James Cook Medal was established by Council in 1943 following a donation made by Henry Ferdinand Halloran to celebrate his 50 years as a member of the Society and it has been awarded periodically since 1947.

The application procedure for this Medal is described on the nomination form. Each nomination must comply with the conditions of award and will consist of a completed nomination form together with supporting documentation as specified on the form.

RSNSW Edgeworth David Medal — for mid-career researchers

Awarded for the most meritorious contributions to knowledge and society in Australia or its territories, conducted mainly in NSW by an individual who is from 5–15 years post-PhD or equivalent on 1 January of the year of the award, together with signs of leadership. The recipient may be resident in Australia or elsewhere. All assessments consider interruptions and performance relative to opportunity.

The Edgeworth David Medal was established by Council in 1943 in honour of Sir TW Edgeworth David FRS, who compiled the first comprehensive record of the geology of Australia, and following a donation made by Henry Ferdinand Halloran to celebrate his 50 years as a member of the Society. It has been periodically awarded since 1948.

The application procedure for this Medal is described on the nomination form. Each nomination must comply with the conditions of award and will consist of a completed nomination form together with supporting documentation as specified on the form.

RSNSW Ida Browne Early Career Medal

Awarded for the most meritorious contributions to knowledge and society in Australia or its territories made mainly in NSW, together with signs of emerging leadership. The recipient may be resident in Australia or elsewhere. Applicants are expected to have been awarded a doctorate or equivalent postgraduate degree. In addition, since this is an early academic career award, applicants are expected to be in the first few years of their career. To be eligible they must have published their first research paper no more than 6 years ago (that is the date that the publication was formally accepted, as shown on the formal publication — not electronic pre-print, should be no more than 5 years prior to the close of applications, normally 30 September. So, for 2025 the publication acceptance date of the first paper should be no earlier than 30 September 2019). All assessments consider career interruptions and performance relative to opportunity.

The Ida Browne Medal was established by Council in 2023 in honour of Ida Browne DSc, palaeontologist and first woman President of the Royal Society of NSW, serving from 1953–1954.

The application procedure for this Medal is described on the nomination form. Each nomination must comply with the conditions of award and will consist of a completed nomination form together with supporting documentation as specified on the form.

Discipline Awards and Lectureships

Applications for the Society's discipline awards and lectureships will be called for annually. These Awards are made on a three-yearly cycle, Years A, B and C, as indicated in the listing of the Awards, with Year A Awards in 2026, Year B in 2027, and Year C in 2028. The discipline awards for Year A (offered in 2026, 2029, 2032, …) are as follows.

RSNSW Clarke Medal and Lectureship in the Earth Sciences

Awarded for distinguished research in any area of the sciences affecting the planet, excluding Medicine and Veterinary Science, and Agricultural and Environmental Science, conducted mainly in NSW. Recipients may be resident in Australia or elsewhere.

The Royal Society of NSW Clarke Medal honours Rev. William Branwhite Clarke, a geologist, and a significant figure in the history of the Royal Society of NSW, who served for a decade as a highly influential Senior Vice-President of the Society (in the early years, the Governor of NSW was the President of the Society). The Medal was first awarded in 1878 "*for meritorious contributions to Geology, Mineralogy, and Natural History of Australasia*." The first Clarke Memorial Lecture was delivered in 1903. Since 2018, the Clarke Medal and Lectureship have been conjoined.

The application procedure for this award is described on the nomination form. Each nomination must comply with the conditions of award and will consist of a completed nomination form together with supporting documentation as specified on the form. Completed nominations (ideally compiled into a single PDF file) should be sent to the email address listed on the nomination form.

RSNSW Walter Burfitt Award in Medical and Veterinary Sciences and Technologies

Awarded for distinguished research in any area of the Medical and Veterinary Sciences and Technologies, conducted mainly in NSW. Recipients may be resident in Australia or elsewhere.

The Walter Burfitt Award honours the life and work of Walter F Burfitt BA MB ChM BSc, an eminent Sydney surgeon in the early twentieth century. It was established as a prize with generous support from Dr Burfitt and his wife, and was first awarded in 1929. In 2004, funding for the prize was augmented by Dr Burfitt's granddaughter, Dr Anne Thoeming. In 2023, Council designated it the Royal Society of NSW Walter Burfitt Award.

The application procedure for this award is described on the nomination form. Each nomination must comply with the conditions of award and will consist of a completed nomination form together with supporting documentation as specified on the form.

RSNSW Award in the Social and Behavioural Sciences

Awarded for distinguished research in any area of the Social and Behavioural Sciences including Economics, Management, Psychology (with the exclusion of clinical psychology, which is included in the RSNSW Walter Burfitt Award in Medical and Veterinary Sciences and Technologies) and related disciplines, conducted mainly in NSW. Recipients may be resident in Australia or elsewhere.

Council established the Royal Society of NSW Social and Behavioural Sciences Award in 2023 to reflect the full scope of the Society's founding values.

The application procedure for this award is described on the nomination form. Each nomination must comply with the conditions of award and will consist of a completed nomination form together with supporting documentation as specified on the form.

RSNSW Award in the History and Philosophy of Science

Awarded for distinguished research in the History and Philosophy of Science conducted mainly in NSW. Recipients may be resident in Australia or elsewhere.

The Royal Society of NSW History and Philosophy of Science Award was established by Council in 2013 to reflect the founding values of the Society and was first awarded in 2014.

The application procedure for this award is described on the nomination form. Each nomination must comply with the conditions of award and will consist of a completed nomination form together with supporting documentation as specified on the form.

RSNSW Award in Computer Science or Information Technology

Awarded for distinguished research in Computer Science and Information Technology, including recent subdisciplines such as Artificial Intelligence, Data Science, Cyber Security and existing areas of computer science and information technology more generally. Recipients may be resident in Australia or elsewhere.

The Royal Society of NSW Award in Computer Science or Information Technology was established by Council in 2025 to reflect the founding values of the Society and will be first awarded in 2026.

The application procedure for this award is described on the nomination form. Each nomination must comply with the conditions of award and will consist of a completed nomination form together with supporting documentation as specified on the form.

Scholarships, Early Career, and Student Awards

RSNSW Bicentennial Postgraduate Scholarships (3)

Each year, the Society awards three Bicentennial Postgraduate Scholarships. The value of the scholarships is determined annually by Council, and includes a complimentary year of Associate Membership of the Society. The awards will be made primarily on the basis of an academic paper that is published by a higher-degree research student, enrolled at a university in NSW or the ACT. The paper must have been published in final form (electronic or paper form) after 1 July in the preceding year. The student must have been enrolled as a higher-degree student at the time of the initial submission to the journal in which the paper appeared. The student need not be the only or the first author on the paper, but, if there are multiple authors, their role in the work should be carefully explained. These details and an outline of the paper and its significance should be summarised in no more than one page. The curriculum vitæ of the student will also be considered, as will recommendations by the nominators, one of whom should be the supervisor.

Winners will be expected to deliver a short presentation of their work at a general meeting of the Society in February or later of the year following that in which the award was made, and also submit a paper to the *Journal & Proceedings*.

Scholarships were first awarded by the Royal Society of NSW in 1999 and in 2023 were redesignated by Council to commemorate the Society's Bicentenary.

The application procedure for this award is described on the nomination form. Each nomination must comply with the conditions of award and will consist of a completed nomination form together with supporting documentation as specified on the form.

RSNSW Bicentennial Early Career Research and Service Citation (3)

Three citations plus a complimentary year of Associate Membership of the Society are awarded each year to recognise outstanding contributions to research and service to the academic and wider community. Applicants must on 1 January of the year of nomination

be no more than 5 years after the award of their PhD or equivalent by a university or other research institution in NSW or the ACT.

Winners will be expected to deliver a short presentation of their work at a general meeting of the Society in February or later of the year following that in which the award was made, and also submit a paper to the *Journal & Proceedings*.

Council established these Early Career Citations in 2023 to commemorate the Society's Bicentenary.

The application procedure for this award is described on the nomination form. Each nomination must comply with the conditions of award and will consist of a completed nomination form together with supporting documentation as specified on the form.

RSNSW Jak Kelly Postgraduate Award

Awarded for excellence in postgraduate research in physics annually. The winner is selected from presenters at each year's Australian Institute of Physics, NSW Branch Postgraduate Awards, as advised to the Awards Committee of the Royal Society of NSW.

The Jak Kelly Award honours Jak Kelly (1928–2012), Professor and head of Physics at the University of NSW (1985–1989), Honorary Professor at The University of Sydney (2004), and President of the Royal Society of NSW (2005–2006). It was first awarded in 2010.

There is no nomination form for this award.

Internal and Discretionary Awards

Notes relating to Internal and Discretionary Awards:

- For Internal Awards, the nominator and seconder must be either a current Member or a current Fellow of the Royal Society of NSW.
- Selection of these Awards is made by the Council of the RSNSW, excepting for the Archibald Ollé Award.

RSNSW President's Award

Awarded at the discretion of the President and Council of the RSNSW to an individual whose distinguished work in any area has made an outstanding and eminent contribution to the State and people of NSW. The recipient may be resident in Australia or elsewhere.

Council established the Royal Society of NSW President's Award in 2023 to reflect the full scope of the Society's founding values. There is no nomination form for this award.

RSNSW Citations (3)

The Royal Society of New South Wales Citations recognise an individual who has made significant contributions to the Society, but who has not been recognised in any other way.

The Royal Society of NSW Citation was first awarded in 2019. Council may make up to three Citations in any year at its discretion.

The application procedure for this award is described on the nomination form. Each nomination must comply with the conditions of award and will consist of a completed nomination form together with supporting documentation as specified on the form.

RSNSW Medal

The Royal Society of New South Wales Medal recognises an individual who has made meritorious contributions to the advancement of knowledge in any field and also to the Society's administration, organisation, and endeavours.

The Royal Society of NSW Medal was first awarded in 1884, revived in 1943, and has been awarded periodically thereafter. Council may award the Medal in any year at its discretion.

The application procedure for this award is described on the nomination form. Each nomination must comply with the conditions of award and will consist of a completed nomination form together with supporting documentation as specified on the form.

RSNSW Archibald Ollé Award

Awarded to the author/s of the best paper submitted to the Society's *Journal & Proceedings* in any year in which the Award is made.

The Archibald Ollé Prize was first awarded in 1956, established by a bequest from Mrs AD Ollé. The award of the Prize (currently $500) is determined by the Editor of the Society's *Journal*, in consultation with the Editorial Board. There is no nomination form for this award.

Knowledge for a Nation:
Origins of the Royal Society of New South Wales

Anne Coote

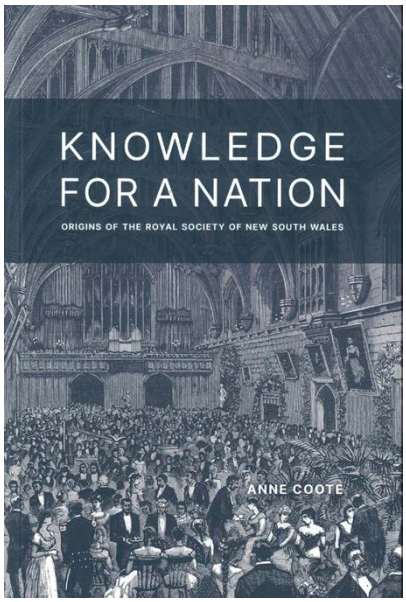

The Governor: "this wonderful history"
The State Librarian: "a crystal-clear and beautifully constructed exposition on the intellectual milieu, international parallels and the political machinations behind the creation and administration of the Royal Society. It helps inform not only the history of science in NSW, but more broadly the larger intellectual climate of Australia."

"a book which is genuinely engaging" "the storytelling is compelling"

"takes the story outside the constrictions of Royal Society personalities, into the much more significant world of colonial society and politics"

"tells the early history of a learned society still active in the intellectual culture of 21st century Australia. Dr Coote's *Knowledge for a Nation* beautifully captures this story."

The Society President: "a very readable account that is hard to put down. By using the characters involved in the Society to tell its story, she has produced a rollicking recital of their lives and times, of the Society's fortunes and misfortunes, ebbs and flows, highs and lows."

Peter Shergold: "Pick up this wonderful and engaging history by Dr Anne Coote, and you will find yourself in a different place. You will be transported to another time, listening in on vigorous debates about science, technology, medicine, philosophy and society. More importantly, you will gain a sense of the world which shaped their views and their approach to intellectual discovery."

"because of Coote's wonderful attention to just the right detail, we are allowed to comprehend the world that existed" 140 years ago

"Coote centres her history of the scientific establishment on the flesh, blood and temperament of the participants, not just on the ideas that they espoused. We meet real people, some snooty, some engaging, some dull, some respectable, some downright unethical but, nearly all of them, insatiably curious or at least feeling the need to behave so in front of others."

"between the covers of this absorbing book lies ... more than a travel guide to a foreign country... History is not just a quaint place to visit. It has made the world in which we live today and beckons to the future which awaits."

"To potential readers — Royal Society members and beyond — I thoroughly recommend this book. **Buy it. Read it. Think about it."**

Knowledge for a Nation: Origins of the Royal Society of New South Wales by Anne Coote

ISBN: 9780645859409

Publisher: The Royal Society of NSW, Sydney

Publication date: 1 October 2024

The book is available from the RSNSW Online Shop at $50 for RSNSW members (https://members.royalsoc.org.au/rsnsw- shop/), or from John Reed Books ($59.95), https://johnreedbooks.com.au/p/knowledge-for-a-nation?barcode=9780645859409

The Royal Society of New South Wales

Information for authors

The *Journal & Proceedings* welcomes manuscripts for review in the many areas of interest to the Royal Society: science, engineering, social science, politics, arts, philosophy, and the humanities. Papers presenting aspects of the historical record of research carried out by Australians or within Australia are particularly welcome. Papers (other than those specially invited by the Editorial Board) will only be considered if the content is either substantially new material that has not been published previously, or is a review of a major research programme.

Letters to the Editor, Discourses, Short Notes and Abstracts of Australian PhD theses may also be submitted for publication. We welcome sets of papers where disagreements among authors are ventilated and argued. Please contact the Editor if you would like to discuss a possible article for inclusion in the *Journal*.

In the case of papers presenting new research, the author must certify that the material has not been submitted concurrently elsewhere nor is likely to be published elsewhere in substantially the same form. In the case of papers reviewing a major research programme, the author must certify that the material has not been published substantially in the same form elsewhere and that permission for the Society to publish has been granted by all copyright holders.

Details of submission guidelines can be found in the online Style Guide for Authors at: https://www.royalsoc.org.au/society-publications/the-royal-society-of-nsw-journal/information-for-authors/.

Note that references should be formatted using the Sage-Harvard reference style. Include the DOI or ISBN of each reference when available. Do not attempt to copy the page layout of the *Journal*.

Manuscripts are only accepted in digital format (generally Microsoft Word) and should be e-mailed to: editor@royalsoc.org.au. In the unlikely event that a file is too large to email, it should be placed on digital media and posted to:

The Editor,
The Royal Society of New South Wales,
PO Box 576,
Crows Nest, NSW 1585

Manuscripts will be reviewed by the Editor, in consultation with the Editorial Board, to decide whether the paper will be considered for publication in the *Journal*. Manuscripts are subjected to peer review by at least one independent reviewer. In the event of initial rejection, manuscripts may be sent to other reviewers.

The Society does not require authors to transfer the copyright of their manuscript to the Society, but authors are required to grant the Society an unrestricted licence to reproduce in any form manuscripts accepted for publication in the *Journal & Proceedings*. Enquiries relating to copyright or reproduction of an article should be directed to the Editor.

Volume 159 Part 1 2026 Numbers 501 & 502

CONTENTS

The Royal Society of New South Wales
P.O. Box 576, Crows Nest, NSW 1585, Australia

ISSN 2653-1305

9 772653 130007 01 >

info@royalsoc.org.au (general)
editor@royalsoc.org.au (editorial)
www.royalsoc.org.au
www.facebook.com/royalsoc
www.youtube.com/@RoyalSocNSW

Published July 2026