

Fairness, Accountability, Transparency and Ethics for a more equitable future



FOREWORD

A steady growth in the general public's concerns about security can be seen around the globe and in Australia, as the Unisys Security Index has reported once again in 2019. These results further reveal aspects of anxiety specifically around data privacy, digital security, and emerging technologies. For example, Australians fear that their smart devices are listening to their conversations, with 41% saying they received related social media posts and advertisements after discussing a topic around their devices.

Fairness, Accountability, Transparency and Ethics (FATE) are elements that need to be considered when emerging technologies are developed and widely used, impacting our everyday life. While this consideration is not unique to Artificial Intelligence (AI) and Machine Learning (ML), it is particularly important to this domain because virtually all aspects of AI/ML require governance on FATE elements, including data, algorithms, training process, deployment, and so on. However, compared with the incredibly fast progression of AI technology over the past decades, there is a limited awareness from the broader audience and stakeholders, about the combination and consideration of FATE elements. Therefore, it is increasingly important for us to invest in a broader, more robust discussion, actively educating, participating and contributing to shaping the future of AI, as well as its impacts on FATE and how we can mitigate relevant risks.

The complexities with respect to FATE are beyond technology alone, and there is a need for wider collaboration between experienced parties from different industry sectors and different regions. Unisys has been helping many governments around the world, solving mission-critical problems like intelligent targeting for border protection, child/family protection, social security, welfare and biometric solutions. The practical hands-on knowledge brings great responsibility for Unisys to take a more active role in facilitating the discussions and shaping the future of AI, through deliberation of tough FATE questions and balancing the governance side of the equation.

Unisys continues to invest in raising the awareness of the global citizen, leading the industry, facilitating roundtable discussions, participating in related committees, with active memberships in leading organisations like the Biometrics Institute. In November 2019, Unisys proudly sponsored the Women Ambassadors Luncheon (Canberra) in partnership with C-Suite Exchange, facilitating a robust panel discussion surrounding FATE in Al.

The following report is a synopsis from this event, detailing the critical interest and support for this topic and the key takeaways, highlighting the wider landscape that is yet to be explored.

Dr. Xin Xu Practice Lead of AI/ML & Data Science Unisys Asia Pacific







Fairness, Accountability, Transparency and Ethics for a more equitable future







Catrina Denvir
Senior Lecturer,
Monash Business
School and Research
Affiliate, Monash Data
Futures Institute,
Monash University

Kate Marshall
Partner, Head of KPMG Law





Jeannie Marie Paterson Professor of Law and co-director of the Centre for AI and Digital Ethics, University of Melbourne

Angela Kim Education Ambassador, WAI Australia



Carina Kemp
Director eResearch, AARNet

Gandhi Sivakumar

Executive Architect and Vice President for AI, IBM AoT, IBM Australia



Fairness, Accountability, Transparency and Ethics for a more equitable future



Overview

Artificial intelligence (AI) is a theme never far from the news headlines. In a typical week the interested reader can learn about its take up by corporations, academia or government.

Examples from recent media coverage prove the breadth of possibilities Al offers. Reporting on the corporate world, Techcrunch.com noted Microsoft has been piloting Al programs in India, where it's trialling the use of machine learning to analyse pap smear results and detect cervical cancer. In the USA, Forbes was one of many reporting on coffee behemoth Starbucks's recent announcement about the suite of Al applications involved in its Deep Brew program; Al will support Starbucks operations on tasks as diverse as espresso machine maintenance, store inventory and order customisation.

During the same period, Science Daily reported on developments in academia, where researchers are using AI to predict lightning strikes. The New Yorker asked whether a machine would be able to write its stories in future, part of its report on Open AI's next language model, GPT-2 – a tool that is generating coherent written text and has potential as a curator of new ideas.

Governments too, are speaking more about AI. Russia's Vladimir Putin recently promoted the potential for AI to help cut through bureaucratic red tape. Back in Australia, Business Insider was one of many mastheads covering the Australian Government's recent announcement of its AI ethics framework; eight ethical principles report that are the government's nod towards ensuring AI functions for the good of society and the environment.

As even these few examples show, AI has the potential to offer society myriad benefits. But an AI future also brings up a raft of complex questions. In a world where AI is deeply embedded into our work and private lives, how do we ensure this future is equitable, and that fairness, accountability, transparency and ethics play a central role?





Fairness, Accountability, Transparency and Ethics for a more equitable future

To understand the future, we must look to the past

Most people thinking about AI look purely to the future, but as Professor Genevieve Bell, Director of the 3A Institute points out, we also need to acknowledge its roots.

Humans have long imagined a world where we could build the perfect human, Bell notes, although many of us don't realise that in the last century this conversation was popularised via the creative industries as a response to World War I.

The Czech playwright Carol Kafek's 1920 play R.U.R. (Rossum's Universal Robots) played a key role. The story of "creatures that only felt efficiency, and that would replace the tasks humans deem dangerous and tedious" ended up on Broadway, the West End, Tokyo and Sydney.

It also became the first piece of science fiction played on the radio and on television, and that most US students encountered in high school. But as Bell notes, the play's most lasting effect was not its storyline, but an invented word it contained within: robot.

The term, derived from the Czech word "Robota" translated to serf or drudge. Today, we don't think of it as a word invented in art, but regardless, the concept of robots has remained a seductive idea scientists have found difficult to let go of.

As early as the 1930s scientists across the globe spent their time trying to design robots, and by 1956, when a conference in Dartmouth coined the term artificial intelligence, the two became intertwined.

Al is a story with multiple beginnings. Bell points out that this means it's not possible to have a conversation about Al without also having a parallel conversation about where it came from, what's encoded in it, and its social and cultural roots as a reaction to a particular moment in time.

How do we define AI?

The definition of AI is far from neat and tidy. The panellists all explain it differently, depending on their audience.

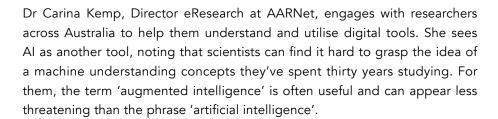
Gandhi Sivakumar, Executive Architect (Data & Al Expert Labs) and Vice President for Al, IBM AoT, IBM Australia, says it's important to make it easy for people to understand. She tries to get people to think of Al by explaining that when humans learn something, they need to be trained. Al systems do too, however they have the advantage of collective intelligence to ultimately support "cognition amplification".







Fairness, Accountability, Transparency and Ethics for a more equitable future



Explaining AI to lawyers has also created some challenges for Dr Catrina Denvir, Senior Lecturer in the Monash Business School at Monash University. She's tried to get it down to an art form, explaining the distinction between symbolic AI which formalises a series of rules, and sub-symbolic (data-driven) methods which apply machine learning techniques to large datasets to reveal relationships between data-points. While both of these approaches help the legal profession to understand how AI could work, Denvir notes that this type of discussion is just referring to the "superficial" processes known as Shallow AI.

Al is just an algorithm, and there is a lot of hype about Al can deliver on complex problems. Some of it is true. Most of it is not. This is because the superficial use of conventional "Big Data" methods are wholly inadequate and often dangerously misleading as they can create a false sense of security when dealing with complex problems.

To explain this statement, it is useful to think of Al algorithms as tackling two types of problems; prediction (sometimes called shallow AI) and inference (sometimes referred as deep Al). The distinction between these two types of Al is important, particularly in the context of an ethical future. As Professor Sally Cripps, Director of the Centre for Transactional Data Science at the University of Sydney points out, Shallow AI is what we mostly have now; relatively simple models for large volumes of data to help decision making in a context where decisions

are repeatable and reversible. Deep Al is characterised by complex models with relatively small amounts of data, to aid decision making in highly ambiguous environments. For example, "What led to the war?" and "Was there a point we could have intervened to stop it?" and get actionable answers?

The bushfires currently raging across New South Wales provide another future scenario. What if we could not only ask, "What brought them about?" but "What could we have done to reduce the risk?" and "How do we protect Australia for future generations?"

These are complex issues requiring multiple sources of information, and it's Deep AI that will enable us to answer a whole different set of problems beyond the simple predictions most people think of AI in terms of at the moment.







Fairness, Accountability, Transparency and Ethics for a more equitable future

What opportunities lie ahead?

Not only will the opportunities for using Deep AI provide researchers with answers, but also they'll help support the researchers and CEOs of the future to come up with their own solutions.

Angela Kim, Women in AI (WAI) Education Ambassador for Australia says education for young people on AI is one space where government and corporates are currently cooperating well. She's particularly involved in the efforts to provide environments for girls from low socioeconomic environments to learn about coding and start thinking about what technology means to them as global citizens.

Ethics and regulation in a global market

Al may be global, ubiquitous and doesn't care about geographic boundaries, but governments, corporations and societies need to realise that regardless, Al is developed within specific cultural contexts.

This means, as Bell reminds the audience, that when we see a data set we need to think: "Whose world does that data reflect?" pointing out that if we don't know the answer we can train a system on that data set and think it's stable but then, upon transferring it to another environment, bring problems with us.

This is important when considering if regulatory environments (and notions of a regulatory sandbox) are the right space to develop Al. Kemp notes that in research they are working to international standards so that conversations about specific pieces of data translate internationally, but translating that approach to Al is tricky.

So will one set of norms and ethics prevail? While the idea of "winners and losers" in an "Al race" is problematic according to the panel, the issue of values is real. As Professor Jeannie Marie Paterson, of the Melbourne Law School (University of Melbourne) notes, citing a recent ACCC action against Google (about consumers being misled about privacy statements), part of what we need to do in Australia is to make decisions about the values we want to regulate for. Law is shaped by those important decisions about fundamental values, and the conversation about values needs to be pitched widely and be genuinely inclusive.

It also seems the "who" should never be far from the equation. Denvir believes a distinction should be made between ethicality in AI systems developed by governments versus those developed by corporations. In part, this is because government led programs will primarily be used here in Australia, making it easier to define the norms and values it should adhere to. In a global, corporate environment, this is obviously more challenging.







Fairness, Accountability, Transparency and Ethics for a more equitable future

When considering regulatory structures around AI, while some decisions about this are important, Kate Marshall, a Partner and Head of Law at KPMG says she often hears resistance to AI posed in terms of: "But we don't know what the liability environment looks like". She suggests that while we will require some new legislation around AI, including aspects of it we don't yet understand, there are areas where we'll have to be more pragmatic, acting now and reacting afterwards. "It's a challenge but should not stop all progression," Marshall says. She sees an important role for Data61 and the government in supporting Australians to manage those fears.

Paterson added that we have a robust legal regime in Australia and we need to think in a nuanced way about how that regime may work with new technologies. This adaptive model may work better than simply rushing to enact more law that may be "out of date' by the time it is enacted, given the pace of technological change. We need to bring design values to ethical AI and also to the law that will apply to demand fair and safe outcomes from AI.

Marshall wonders whether we seek to hold AI to account at a different level. She notes autonomous vehicles as an example, pointing out that society is unwilling to accept some risks even if this is a significant improvement on the current state, while with the development of other types of vehicles throughout history we have accepted this reality and tolerated this level of risk.

As an interesting aside on the autonomous vehicle issue, Cripps points out that the recent pedestrian death by an autonomous vehicle during a testing phase in the US was an example of the Shallow AI; the algorithm was trained only to make limited predictions rather than act when uncertainty was present (the uncertainty in this case was a pedestrian who was jaywalking, and the vehicle was only trained to consider pedestrians on crossings). It's one reason why Cripps says a fundamental shift away from Shallow AI is vital.

Privacy issues and FATE in Al

Are our privacy laws fit for the future? Marshall doesn't think so - some reform is needed and strong privacy laws need to provide that framework and harmonising international laws would help Al's maturity.

Cripps says we need a future of democratised data. She worries that companies like Google, Facebook and Apple own the world's biggest asset. She suggests that if data on an individual was released to that individual, that action could represent the start of data democratisation and reiterates that this is a serious challenge, particularly for vulnerable populations.

...we have a robust legal regime in Australia and we need to think in a nuanced way about how that regime may work with new technologies.





Fairness, Accountability, Transparency and Ethics for a more equitable future

Kim agrees, adding that this return of data should be done in a way that even those without technical understanding can understand.

Denvir stipulates that integrated education is vital: Every student in every discipline needs to understand data security and privacy, what automated and algorithmic decision making actually entails, and what risks surround the use of AI in a range of different context.

Bell takes this idea one step further, suggesting that there needs to be a shift in the discourse that "more data equals more truth". After all, more data won't be what turns us into "better people". However, maybe it will do what Cripps suggests, and give us a way to deeply understand what we don't know. "If we can measure uncertainty we can be really critical about what we are using data for," Cripps says.

Measuring Al's success

Many panellists mentioned their hope that a future of AI can be a societal turning point: something that delivers benefits that go beyond the solutions offered by the technologies themselves. Instead, they suggest, an equitable AI future could – and should – herald an opportunity to question broader systems like capitalism, democracy and economic progress.

So is productivity going to be the only measure of the success of AI? Is that the way we are going to continue to define its achievements? What if instead, suggests Bell, we use a "whole of society" indicator that includes metrics like civic engagement, trust, and the role of AI in the creative sector, rather than simply sticking with the current metrics of productivity and efficiency?

The environment too, needs to be considered. All systems are energy intensive, with data requirements already using ten per cent of the world's energy just to power where we are now. Imagine the energy needs of the All systems of tomorrow; just one reason this is a conversation governments need to be having with businesses. After all, as Paterson points out, if we don't do something about the environment then All simply won't be relevant.

In terms of timing, Bell reiterates the need to act quickly and ask these types of vital questions now: in ten years time, Al won't simply be in our phones, but in the walls of the buildings we meet in.

It's one more reason supporting Marshall's greatest hope for AI. When we look to a future, she wants to see a world that doesn't just work for "most" people. "AI may, in some instances, allow us to look at things for all people."











Unisys Canberra
70 Kent St
Deakin
CANBERRA ACT 2600
www.unisys.com.au

Dr. Xin XuPractice Lead of AI/ML & Data Science
Xin.Xu@unisys.com



Suite 3, Level 27
Governor Macquarie Tower
1 Farrer Place, Sydney NSW 2000
AUSTRALIA

www.c-suiteexchange.com connect@c-suiteexchange.com (61) 2 9401 0680

This issues paper is the third in a series of C-Suite Exchange
Women Ambassadors Luncheons on FATE in AI, that collates learnings and
insights from influencers in government, academia and industry.

Authored by Sue White, The Comms Collective

© Copyright C-Suite Exchange Pty Ltd. 2019 All rights reserved.