Editorial: Artificial Intelligence: A New Transactions

I. Introduction

COGITO Ergo Sum, 'I think, therefore I am' (René Descartes, 1596–1650). Cognition, thinking, intelligence and existence are inseparable concepts in many philosophical inquiries. David Hume (1711–1776) brought computations to the discussion, deeply embedding a widely believed premise in Cognitive Science that cognition is computation. George Boole (1815–1864) challenged human presumptions of complexity using three binary operators: 'and,' 'or' and 'not.' Boole demonstrated that these are sufficient for universal computations.

By deduction, and by taking the perspectives of Descartes, Hume and Boole as axioms, we dare to conclude that 'and,' 'or' and 'not' are the core enablers of existence. Most humans would be extremely sceptical of the validity of this conclusion, but not perhaps Herbert Simon (1916–2001), who believed that behavioural complexity is attributable to the environment, and not to the agent. That is, we are simple machines. Our, possibly self-made, illusion that we are complex beings is just that-an illusion-because all complexity is caused by the environment in which we are situated. The units of our brain are neurons made of synapses and axons. The complexity in the organisation of the brain is determined by what the environment allowed us to perform over the course of our evolution.

Let us suspend our scepticism for a bit longer and allow our free-of-scepticism mind to develop this line of thinking. After all, regardless of whether our human brain is simple or complex, for centuries biological brains monopolised computations and intelligence as the only computational devices we know of that were capable of exhibiting intelligent behaviours.

Unsurprisingly, this monopoly ceased to exist when Charles Babbage (1791–1871) invented the first mechanical computational device, the computer. When Alan Turing (1912–1954) designed the Turing machine as a general-purpose computer, it became clear that biological brains would never again monopolise computations and intelligence. Another machine, called the computer, could now perform any form of computation, work with zeroes and ones and implement 'and,' 'or' and 'not' operations. Since computers can compute and are universal computational machines, they could form cognition of their own to think. With the ability to think, they can assert their existence. Thus, the two fundamentally challenging questions remaining were as follows: Can computers actually advance their ability to think to a level of intelligence comparable to that of humans? If yes, and with the early logic discussed above, they can assert their existence. This raised the second question: Can computers coexist with humans? Interestingly, within six years of Turing's death, answers to both questions started to emerge.

Can computers develop intelligence? The question was answered in the summer of 1956 during the famous Dartmouth Workshop led by four scientists: John McCarthy, Marvin Minsky, Claude Shannon and Nathaniel Rochester. In their proposal to obtain funding for the workshop, they stated 'The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it' [item 1) in the Appendix]. They also stated, 'The major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have' [item 1) in the Appendix].

We will revisit these statements, but given that the workshop gave birth to the new field now termed 'Artificial Intelligence' (AI), it became pertinent to ask whether such a new cognitive agent with computational abilities leading to cognition, which, in turn, is leading to existence, can in fact exist at all, let alone coexist with humans. Within four years of the Dartmouth Workshop, the ambition to leverage this revolutionary technology was taken to a new level by a project manager from the Defense Advanced Research Projects Agency (DARPA).

In 1960, Licklider [item 2) in the Appendix] published his famous paper, 'Man-Computer Symbiosis.' Notably, he used the word 'man' since at that time it meant 'males and females,' whereas since the 1970s, it has been used to refer to 'males' alone and has been replaced with the gender-neutral word 'human.' Licklider predicted that by 1980, AI will have the ability to think and solve problems and that by 1985, humans and AI machines will coexist, working in harmony together.

However, two different schools have contributed to the evolution of AI. The symbolic school, with its emphasis on symbolic representations, relied mostly on deductive reasoning following Aristotelian logic. Aristotle's (385 BC–322 BC) famous modus ponens sat at the core of deductive reasoning. Later, abductive logic, as introduced by Peirce in 1887 [item 3) in the Appendix], was added to the symbolic school. A non-symbolic school that focused on induction emerged with Rosenblatt's [item 4) in the Appendix] work in 1958 on the Perceptron. The non-symbolic school continued to grow in ambition with its work on emergence and adaptation, which led to the field of evolutionary computation [item 5) in the Appendix], and with its work on computing with words, which led to the field of fuzzy sets [item 6) in the Appendix] and fuzzy logic [item 7) in the Appendix]. Research areas at the interface of the two schools started to

Digital Object Identifier 10.1109/TAI.2020.3035827

emerge, including inductive logic programming, constraint logic programming and approximate reasoning.

The fulfillment of Licklider's vision seemed imminent, given the emerging research areas. Symbolism mimics our ability to reason, whereas connectionism mimics the architecture of the machine sitting underneath our sense-making, action-production and thinking-the brain. All indicators raised the ambition and expectations that AI would deliver.

Following Minsky and Papert' s [item 8) in the Appendix] technical scepticism regarding Rosenblatt's claims about the Perceptron, the Science Research Council started an inquiry into AI, which concluded with the Lighthill Report in 1973. The report criticised the failure of AI to deliver on the promises that researchers in the field had made. This was the first AI winter. Later, as researchers started to recover-making new grand promises about expert systems-a second AI winter was witnessed in 1993.

The evolution of AI has not matched Licklider's vision. The year 1980 did not record any serious implementation of AI that could compete with humans. The year 1985 did not witness human-computer symbiosis. The year 1993 only saw an AI winter. Nevertheless, hope started to rise again with the IBM Deep Blue program winning against Kasparov in 1997.

Deep Blue is a milestone in the history of AI. However, it was the perfect example to showcase the worries raised by the Dartmouth Workshop organisers, namely, that machine capacity is not the bottleneck, but our ability to write programs is. If it is indeed a problem, then why must we write programs ourselves? Why should we not make computers write programs for themselves? Indeed, this was the solution to overcome the bottleneck that the organisers of the Dartmouth Workshop identified, and it took 60 years to find this answer using deep learning, thus progressing AI towards a new era.

In 2016, the AlphaGo deep learning and Monte-Carlo Tree Search based program defeated a Go world champion. Cracking Go was different from winning at chess. Go is much more complex than chess. AlphaGo was not rule-based–it **learned** from previous human games and made unconventional moves during the challenge. Thus, the bottleneck recognised by the Dartmouth Workshop organisers about human ability to write programs was resolved; computers had begun to learn, write their own programs, and were able to compete against human world champions.

This time, however, it did not take long to see AI penetrating every industry. Schools and universities started using AI for learning analytics. The healthcare sector adopted various forms of AI, ranging from AI systems that assist medical practitioners in diagnosis to those that allocate beds in hospitals. The financial sector, an early adopter of AI technologies, has used it for various purposes, ranging from share market prediction to fraud detection and credit risk scoring. Thus, AI became ubiquitous and started to be widely used.

Moreover, AI is developing towards fulfilling Licklider's vision. Perhaps AI has matured 46 years later than Licklider had anticipated. However, his second prediction-that within five years of AI rising to the challenge, human-AI symbiosis will be feasible-seems more likely to occur. Just four years have passed since AlphaGo defeated Lee Sedol. In 2019, Lee Sedol retired because he did not consider it possible to ever beat AI at GO. Kasparov and Sedol are probably the most famous people whose career was impacted by AI. Social implications, such as unemployment, ethics, accountability, reliability, explainability, transparency and trust, are but some of the contemporary topics in AI. Nevertheless, the future possibly lies in what Licklider envisioned: human-AI teaming, AI augmenting humans and assisting them in their daily jobs. This is the area in which AI can deliver the most benefits to society, at least in the near term.

In this century, AI has evolved into a different form. Most contemporary research no longer attempts to replicate human intelligence but to discover new forms of intelligence to augment human intelligence. Most researchers have become conscious of the social and ethical implications of AI and are striving to develop beneficial, responsible AI. In particular, since the ongoing COVID-19 pandemic has transformed the work environment to teleworking in an amazingly short timeframe, workplace digitisation has grown in an unprecedented manner, which has increased the opportunities to automate and use AI. Organisations have started to realise that the nature of the work environment has changed, and as a result, are increasingly reskilling their workforce to use AI.

In view of these developments, the establishment of the IEEE TRANSACTIONS ON ARTIFICIAL INTELLIGENCE (IEEE TAI) in 2020 is indeed timely. The landscape of AI research is evolving rapidly in a wide variety of areas, including (deep) neural networks, fuzzy systems for explainable AI, applications involving evolutionary computation techniques for autonomous and adaptive machine learning and optimisation, swarm intelligence in distributed AI, natural language processing, constraints, multi-agent systems, reasoning, human-AI teaming, social integration, and societal and ethical implications. The creation of IEEE TAI has involved a significant amount of effort from many people. However, since we have limited space, we can list only a few names. Prof. Nikhil Pal was the President of the IEEE Computational Intelligence Society (CIS) when we proposed launching IEEE TAI. Nik championed the effort and remarkably navigated the proposal past many hurdles. Prof. James Keller is the Vice-President for Publications. Jim has unprecedented knowledge on publications, which was vital to the success of the proposal. He worked tirelessly to bring IEEE TAI into existence. Prof Bernadette Bouchon-Meunier was the IEEE CIS President when the IEEE TAI proposal reached the most critical stage of negotiation. Bernadette's wisdom and experience are superb, and her belief in the importance of IEEE TAI is exemplary. Prof. Pablo Estevez, the IEEE CIS Vice-President, Finance, has backed the process since inception. IEEE TAI has four financial sponsors and one technical sponsor. The presidents of the IEEE sponsoring societies and their vice-presidents for publications were so professional, friendly and positive that it made the creation of IEEE TAI a real joy. In particular, acknowledgement and appreciation go to the four presidents and four vice-presidents for Publications: Leila De Floriani and Lombardi Fabrizio (IEEE-CS), Rudas Imre and Enrique Herrera Viedma (IEEE-SMC), Ahmed Tewfik and Sergios Theodoridis

(IEEE-SPS), and Seth Hutchinson and Aude Billard (IEEE-RAS).

IEEE TAI has a management committee of eminent scientists from all sponsoring societies. A big thank you to the management committee for using their experience and wisdom to support the journal. This acknowledgement is incomplete without a special thank you to all the IEEE staff, who have worked hard to launch the journal; to the IEEE TAI Associate Editors and reviewers, who have been working tirelessly to support the journal; and to current and future authors, who are sending their ideas to IEEE TAI for the world to witness their contributions to the field of AI. The success of the journal is in the hands of the community, and together, we can showcase AI to the world to ensure that from this moment on, AI will grow in a continuous and endless spring, with no more winters to come.

HUSSEIN ABBASS, *Founding Editor-in-Chief*
University of New South Wales
Canberra, ACT 2600, Australia

APPENDIX
RELATED WORK

1] J. McCarthy, M. L. Minsky, N. Rochester, and C. E. Shannon, "A proposal for the dartmouth summer research project on artificial intelligence, August 31, 1955," *AI Mag.*, vol. 27, no. 4, pp. 12–12, 2006.
2] J. C. Licklider, "Man-computer symbiosis," *IRE Trans. Hum. Factors Electron.*, no. 1, pp. 4–11, 1960.
3] C. Peirce, "Logical machines," *Amer. J. Psychol.*, vol. 1, no. 1, pp. 165–170, 1887.
4] F. Rosenblatt, "The perceptron: A probabilistic model for information storage and organization in the brain," *Psychol. Rev.*, vol. 65, no. 6, p. 386, 1958.
5] D. B. Fogel, *Evolutionary computation: The fossil record*. Wiley-IEEE, 1998.
6] L. A. Zadeh, "Fuzzy sets," *Inf. Control*, vol. 8, no. 3, pp. 338–353, 1965.
7] L. A. Zadeh, "Fuzzy logic," *Computer*, vol. 21, no. 4, pp. 83–93, 1988.
8] M. Minsky and S. Papert, *Perceptrons: an introduction to computational Geometry*. MIT, 1969.

Hussein Abbass (Fellow, IEEE) received two bachelor's degrees from Cairo University in 1990 and 1995, the Postgraduate Diploma in operations research and the M.Sc. degree in constraint logic programming from Cairo University, in 1992 and 1995, respectively, the M.Sc. degree in artificial intelligence from the University of Edinburgh, U.K., in 1997, the Ph.D. degree in computer science from the Queensland University of Technology, Brisbane, Australia, in 2000, and a Diploma from the Australian Institute of Company Directors in 2010.

He commenced his career in the IT industry before his transition into academia when he joined Cairo University in 1994. In 2000, he took up a Lecturer position with the University of New South Wales Canberra, Australia. He was promoted to Senior Lecturer (2002), Associate Professor (2005), and Full Professor (2007). He is currently a Full Professor with the School of Engineering and Information Technology, University of New South Wales Canberra, Australia. He is a member of the American Association for Artificial Intelligence, the Human Factors and Ergonomics Society, and a number of IEEE societies including the IEEE Computational Intelligence Society, the IEEE Computer Society, the IEEE Robotics and Automation Society, and the IEEE Systems, Man, and Cybernetics Society. He is a member (2019–2020) of the Artificial Intelligence Ethics Committee of the Australian Computer Society.

Prof. Abbass is the Founding Editor-in-Chief of the IEEE TRANSACTIONS ON ARTIFICIAL INTELLIGENCE. He is an Associate Editor for the IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS, IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION, IEEE TRANSACTIONS ON COGNITIVE AND DEVELOPMENTAL SYSTEMS, IEEE TRANSACTIONS ON CYBERNETICS, *ACM Computing Surveys*, and four other journals. He was the Vice-President for Technical Activities (2016–2019) for the IEEE Computational Intelligence Society, and the National President for the Australian Operations Research Society (2016–2019).

Prof. Abbass is a Fellow of the Australian Computer Society (FACS), Operational Research Society (FORS), Institute of Managers and Leaders Australia and New Zealand (FIML), and the IEEE. He has authored or coauthored more than 300 papers in journals and conferences, and four authored books. His current research contributes to trusted human-swarm teaming with an aim to design next generation trusted and distributed artificial intelligence systems that seamlessly integrate humans and machines.