EDITORIAL

Past the tipping point?

John MacIntyre[1] · Larry Medsker[2] · Rachel Moriarty[3]

We are delighted to welcome readers to our new journal, *AI and Ethics*. Launching a new journal is a significant undertaking, and this one has been in gestation for at least 5 years. It is both rewarding and satisfying to see the first peer reviewed original work published alongside the thought leadership pieces contributed by our outstanding Editorial Board members. We are grateful to all authors for their excellent contributions, and to the reviewers who have provided thoughtful and useful feedback to our authors.

Our journey to this point has been one of learning and enlightenment. Both John and Larry are practitioners of applied artificial intelligence, developing solutions to problems and focussing on the technical aspects of our various developments.

Larry's career started with research in physics and computer science, with increasing interest in applied and curricular work in AI and data science. His Artificial Neural Network research included a focus on hybrid intelligent systems. While the initial decades concerned technological aspects, during the past two decades his evolving interests and concerns focused on the nature of humans and machines and the impacts of AI on individuals and society. A professional development has been his role as ACM SIGAI Public Policy Office, which naturally includes involvement in AI ethics. The advent of the new journal *AI and Ethics* is an obvious and welcome next step.

John came late into academia, having started his career in industry, spanning both engineering and computer science; this led seamlessly to his doctorate in applying neural networks to engineering problems. He established a research centre—the Centre for Adaptive Systems—in the mid-1990s, which became recognised by the UK Government as a 'Centre of Excellence' in applied artificial intelligence, developing applications in engineering, finance, medicine, security, and a range of other areas. He served on the committee of the UK's Neural Computing Applications Forum, and has been Editor-in-Chief of *Neural Computing & Applications*, published by Springer, since 1996. This career-long focus on applied AI has given him the vantage point to see how our field has developed both technically, and in terms of the scope and scale of where AI is being applied—and the moral and ethical issues which are emerging as a result.

As Senior Editor at Springer Nature, Rachel has been working on Computer Science journals for over a decade, mainly focusing on AI, Human Computer Interaction, Graphics and Virtual Reality. Over the years of managing portfolios ranging from 15 to 30 journals, as a Publishing Editor she forms the link between the external communities a journal serves, and the dedicated internal teams within Springer Nature. Through working with John and Larry on *Neural Computing and Applications*, another Springer Nature journal, it became evident in our discussions that the time was right for a new journal on AI and Ethics and our successful relationship established on NCA could be continued into this new launch.

In the three decades in which John and Larry have been active in applied artificial intelligence, so much has changed. Computing power has increased exponentially, and alongside it, the availability of, and ability to manipulate, very large bodies of data. The development of the Internet has, of course, also radically changed the world in that period.

In the 1990s, there was a sense that neural networks had developed to the point that they were a 'clever technique looking for the right problem'. Geoff Hinton's demonstration of the usefulness of the back-propagation algorithm for multi-layered perceptron networks was a watershed, allowing neural nets to be applied to serious non-linear problems and leading to the currently widespread use of deep learning. This led to a re-awakening of research and development in artificial intelligence, which had stagnated somewhat from the period during the 1970s and 1980s when rule-based expert

✉ John MacIntyre
John.Macintyre@sunderland.ac.uk

1 University of Sunderland, Sunderland, UK

2 The George Washington University, Washington, DC, USA

3 Springer Nature, London, UK

systems and predicate logic had promised so much, but flattered to deceive. Even so, progress was slow, partly due to the constraints of computational power, training data, and the tools and techniques of artificial intelligence at the time. So-called 'real world' applications were scarce, and commercial investment in the field was limited.

Now, artificial intelligence is booming, with commercial products, development tools, training and consultancy relentlessly marketed through multiple channels. Product reviews, opinion articles, and analysis pieces are now published daily all over the world; the AI bandwagon is really rolling, and gathering speed.

The pace and breadth of the expansion of AI is breathtaking. There is hardly an application domain, problem type, or market sector where AI is not already in use. Products are being released at an amazing rate, commercial investment and patents are increasing, and investment in R&D is escalating incredibly.

We have sight of some of this work through our collaboration on *Neural Computing & Applications*. This journal publishes high-quality work in applied artificial intelligence, focussing on innovative ways in which AI is developed and used to solve real problems. *Neural Computing & Applications* has seen substantial growth in recent years, in terms of the number of submissions received, but also in the range and type of applications and techniques which are being developed all over the world. This is but one indication that AI is "of its time".

This significant growth brings with it both benefits and risks. Clearly, the capability of AI tools and techniques has increased massively since both John's and Larry's early work in the 1990s, and there are many wonderful examples of how these have been applied to create elegant and effective solutions which bring many benefits to society. Artificial intelligence is generating economic benefit to many companies, as well as improvements in agriculture, manufacturing, health, finance, and so many other fields. It is not an exaggeration to say that AI is the underpinning technology (or group of technologies) powering a new industrial revolution, with profound consequences for the future in an increasingly integrated and internet-enabled world. We celebrate the achievements of AI, and congratulate those who have developed AI applications bringing these benefits.

Disruptive technologies are rarely without some negative consequences and risks. The development of electricity transformed almost every aspect of human life, but also had its casualties, both economically and culturally. The advent of the printing press transformed education across the globe and powered the scientific revolution of the 16th and 17th centuries, but was also responsible for the elimination of the livelihoods of many craftsmen whose artisan skills were replaced by mass production. So it has been throughout history—technology comes with consequences. Artificial intelligence is no exception, but may be exceptional in terms of the unprecedented short time frame of change and the extent of the potential negative consequences if we get it wrong.

Powerful voices have expressed concerns about the possible impact of AI on the future of humanity. Elon Musk has called AI the "greatest existential threat" to our species, and yet is actively developing AI technology. Both Bill Gates and Stephen Hawking have expressed serious misgivings over how AI might be used, and the direction of travel of developments. These notable individuals are high-profile and visible, but many others, including some within the AI community itself, also express concerns, just as there are those who suggest that we need not be worried. As with other fields, many viewpoints and perspectives exist, and the wonderful documentary film released in 2019 entitled "We Need to Talk About A.I.", made by the New Zealand-based directly Leanne Pooley, captures these very well—it is well worth 90 min of your time.

At the crux of the questions raised is the fact that AI has developed such powerful capabilities, and applications are so widespread and ubiquitous, we cannot escape its impacts. As AI continues to develop, it is possible that our field has reached, or even passed, a tipping point where we must stop to ask ourselves a fundamental question—does the fact that we CAN do something, mean that we SHOULD?

This is not an abstract or esoteric conundrum, posed as a philosophical debating point. It is real, and it is with us now, not only for the future.

Other areas of science and technology have wrestled with this question, at various times and in different ways. Genetic science had developed technical capabilities in human embryology and gene editing, which in turn led to the global debate on if, when, and how these techniques should be applied, and how they should be regulated and governed. That debate continues, but much progress has been made (including significant elements of self-regulation by the scientific community itself).

We believe that our field, artificial intelligence, needs to engage in a similar level of global debate, and urgently. The moral, social, and economic questions which are increasingly emerging as AI becomes more powerful and ever-present are profound. They touch on the future of human work, privacy, freedom, fairness, and perhaps our very existence. While that may sound dramatic, it is true.

As scientists and academics we have both been proponents of the development of artificial intelligence, and we continue to believe in the opportunity to develop AI for the good of humanity. At the same time, we have recognised, through our own journeys of learning and enlightenment, that there are also significant risks which should be the subject of open debate.

Our ambition is that *AI and Ethics* becomes a vehicle for that debate, facilitating it through publication of high-quality material which adds to the body of knowledge. We wish for it to be useful to a wide range of audiences—to scientists and academics, researchers and developers (including industrial and commercial developers), users, those developing regulatory and governance frameworks, lawmakers, and ultimately, the public. If we are successful we will have made a positive contribution to the future of the development of ethical artificial intelligence which brings the many benefits to society which we know are possible, but limits negative consequences and risks as far as possible.

As a journal focusing on "ethical AI", we must be held to a higher standard by our community—and we welcome that. One of these standards is the composition of our Editorial Board. From the start, we sought to welcome voices from the AI and ethics field that are diverse in gender, race, ethnicity, ability/disability, expertise, field of focus, age and geography. We are very proud of where we are with our Editorial Board, and we expect it to naturally change over the coming years in response to the needs of the journal and the community.

We are also grateful to colleagues who have agreed to join our growing Review Board, providing the essential component of robust and independent peer review which is the underpinning plank of our commitment to maintaining high publishing standards. We will continue to add to our list of AI and Ethics experts until our Editorial and Review board reflects the many facets of our community. We are deeply grateful to the outstanding individuals who have agreed to join us and we are humbled by your kindness and support.

We would like to thank Springer Nature for their foresight and courage in supporting us to launch the journal, and also the internal colleagues, and specifically Annette Hinze, who will be an essential part of the Editorial Team going forward.

Finally, we would like to thank all authors and reviewers—past, current, and future—for your vital contributions. Without you, there would be no AI and Ethics, no work which can be published and used by others to help them on their journey to developing ethical artificial intelligence.

<div align="right">

John MacIntyre, Larry Medsker and Rachel Moriarty
Co-Editors-in-Chief and Publishing Editor
AI and Ethics
17 September 2020

</div>