



# Actionable ethics

Karamjit S. Gill<sup>1</sup>

Accepted: 18 January 2022 / Published online: 27 January 2022  
© The Author(s), under exclusive licence to Springer-Verlag London Ltd., part of Springer Nature 2022

Whilst grappling with the impact of predictive algorithms and making sense of the data driven society on the one hand, and envisioning the common-good potential of augmented artificial intelligence (AI) systems on the other, we face social challenges of governance, ethics, accountability, and intervention arising from the accelerated integration of powerful artificial intelligence systems into core social institutions. With the exponential rise of big data flows in networked communications and their manipulating algorithms, the gaps in translation from prediction to actuality are now too vast to grasp and address, rendering us unable to engage with difference through the shadows of machine thinking. Augmentation and automation place the human in the predicament of accepting the calculation of the machine without judgement (Cooley 2020).

In a similar vein, Nowotny (2021) proposes that there is a tacit assumption and misplaced confidence that ethical AI would ultimately take care of the unresolved ethical, transparency, and accountability conflicts when we are able to develop computational tools ‘to assess the performance and output quality of Deep Learning algorithms and to optimise their training’. The danger, she says, is that we end up trusting the automatic pilot whilst flying blindly in the fog, becoming part of a fine-tuned and inter-connected predictive system, thereby diminishing our motivation and ability to stretch the boundaries of imagination. Although there is considerable recognition of the wide spread prejudice, discriminatory practises, and biases that reside in the development and use of algorithmic systems and tools such as that of facial recognition and racial profiling, there is belief amongst many of the AI community that solutions to many of the problems besetting society will be found through the ethical machine by substituting the ‘irrational beliefs’, diverse and specific with a single universal logic of ethics.

Human future, from this perspective, is seen in terms of human and AI machine co-evolution.

But what drives this idea of the ethical machine? First, the desire to seek objectified solutions without prejudice in the scientific tradition; Second, belief in calculation as measurement of objectification; Third, confusion in the idea that data are objectivity and not calculation; Fourth, the idea of machine ethics as an extension of human ethics, ultimately becoming fully aligned with the machine’s operations—just as the machine was seen as an extension of human body, now machine intelligence is seen as an extension of human intelligence. All these drivers are contributing to a belief in the pure, universal, and unifying logic of the machine. Nowotny (ibid) reminds us that in delegating more and more human tasks to AIs, human responsibility is being diluted, raising concerns of a fundamental incompatibility between the logic of algorithms and that of human institutions. We echo Cooley’s concerns of ‘socially irresponsible’ science and ask whether we can transcend the instrumental reason of machine thinking to mould technological futures for common good, rather than turning them into a single storey of ‘singularity’. Can we re-appropriate the idea of causality that has been taken by ‘science’ and reframe it in the making of everyday judgments and decisions? How can we harness collective intelligence as a transformational tool for addressing complex social problems? Cooley notes that it is true that the drive for scientific knowledge has provided the material basis for a fuller and more dignified existence for the community as a whole. It must not, however, be a blind and unthinking drive forward, shirking our social and political responsibility to analyse its effects upon society. Any meaningful analysis of scientific abuse must probe the very nature of the scientific process itself, and the objective role of science within the ideological framework of a given society. As such, it ceases to be merely a ‘problem of science’ and takes on a political dimension. It extends beyond the idea of important, but limited, introverted soul-searching of the scientific community, and recognises the need for wider public involvement. The challenge is to create a strategic framework that facilitates this change, in response

✉ Karamjit S. Gill  
editoraisoc@yahoo.co.uk

<sup>1</sup> University of Brighton, Brighton, UK

to the technologies of computerization and automation, for example in dealing with the disruption of social, economic, and cultural life, especially when life becomes synchronised with the computerised environment. Although humans with their skill and ingenuity were able to create technological change from the early stages to the advance of artificial intelligence, the society which has given birth to them tends to fail to keep pace.

Those who are engaged in the pursuit of machine ethics and governance are reminded that actionable ethics is also about the pursuit of inclusive participation and openness towards knowledge of the past, complexities of the present, and uncertainties of the future. In the end, it is not important how the AI machine can be aligned with human values or visualising as how human values are fully aligned with the AI machine, converging to the post-human world, what it is important to know is that human values are diverse, social, cultural, and contextual, and they do not fit into the logic of the AI machine. It may be time to rethink, if not ‘debunk’, the mythical notions of the ethical machine, post-humanism, and singularity that promote the ideas that machine learning can create ethics. In many way ways, our authors in this volume are moving away from these mythical notions in their examination of the shifting human–machine relations. AI&Society welcomes contributions on the debate on human machine co-evolution and actionable ethics.

In this volume, our authors continue the debate on actionable ethics from a multiple perspective, ranging from Framing AI systems in healthcare sector; Social machine as a tool for shaping interactions between individuals and algorithms; Algorithmic accountability, transparency, and intentional biases; Algorithmic augmentation of democratic processes, Discrimination in the age of artificial intelligence; Ethics and biometric facial recognition technology; Algorithmic and human decision-making, and standards of transparency; Explainable artificial intelligence and its intrinsically value and desirability; How do people judge the credibility of algorithmic sources; Shifting relations of human autonomy and technological automation; Ethical challenges and organisational responses on responsibilities of policymakers, professional bodies, and regulators; Data objects for knowing—Data Science a technology-driven science; Endowing artificial intelligence with legal subjectivity; In search of the moral status of AI; Actionable Ethics for governance; Sensorimotor debilities in digital cultures; Social acceptance of robots; Child–robot relationship formation; AI machine and the art of education; The challenge of defining cross-cultural fairness assessments of texts; Multifaceted nature of the transformation; Impact of AI on human behaviour and emotions in a multicultural educational context; AI for seeing creativity assessment of culinary products as art; The making of AI futures in German context; The limit of human anthropocentric tendencies of control; Utilitarianisms and

machine ethics; Dystopian conception of post-humanism vs Africanist civilizational humanism.

Mercedes Bunz and Marco Braghieri, in ‘The AI Doctor will see you now’ (this volume), study the qualitative outlook into the relationship between news media production and the depiction of Artificial Intelligence in the healthcare sector. They discuss the framing of AI systems and their agency replacing and outperforming the human medical expert, with the consequence that this framing of ‘outperforming’ might place AI systems above critique and concern including the Hippocratic oath. AI systems are already becoming more and more common in healthcare sector, being implemented for delivering tasks within healthcare, and at times supplanting and delivering institutional decisions. The authors argue that by suggesting decisions such as diagnosing an illness or suggesting a personalised treatment, these technologies are taking on societal functions and take part in shaping our societies. There is thus a need of space for reflection and an open debate on their ‘traits’ and operations that encourages an understanding of their technical capacities and implication of AI systems.

Orestis Papakyriakopoulos, in ‘Political machines’ (this volume), explores the way *Social Machines*, per se political machines, serve as a framework for understanding and interpreting interactions in socio-algorithmic ecosystems. The author posits that a cybernetic perspective of *Social Machines* allows for the investigation of the interplay of political processes, and the proposed framework can be used to categorise dimensions of influence that shape interactions between individuals and algorithms. It is argued that since socio-algorithmic ecosystems constantly face important ethical and political challenges, the framework can be used to semantically plan how potential interventions, either in the regulation, or in the political machine itself, might change systems’ dynamics. The author suggests that researchers should not only describe how political machines function, but also define principles, frameworks, and constraints that can lead to the creation of socio-algorithmic ecosystems that serve the public interest.

Antonin Descampe et al., in ‘Automated News Recommendation... Algorithmic Accountability’ (this volume), discuss the question algorithmic accountability and suggest that it essentially boils down to a question of transparency, and thus of technical limitation to deal with intentional biases. The authors, in their work on an automated news recommendation system, illustrate that adversarial behaviours resulting from biases may make the algorithm deviate from its publicised intent embedded into the algorithm. They suggest that robustness against adversarial behaviours should be taken into account in the definition of algorithmic accountability, to better capture the risks inherent to algorithmic decision-making. In reference to computational journalism, the authors note that the acceptance of a level of

risk may also depend on the values that may be threatened by the biases potentially induced by adversarial machine learning tools. Although, in general, the adoption of precautionary measures may be the most effective way to cope with manufactured risks, the situation of a newsroom significantly differs in the sense that journalists may act as an interface (‘human in the loop’) between automated news production (and compilation, ...) and their readers. It is noted that as long as robust algorithmic accountability cannot be ensured by technical means, a more user-centric (responsibility-based) view of algorithmic accountability therefore appears as a necessary complement to the design-centric approach, to compensate its limitations.

Paul Burgess, in ‘Algorithmic Augmentation of Democracy’ (this volume), explores the ways AI and other forms of technology could be used to augment the representative democratic process. The augmentations range from voting online to the wholesale replacement of the legislature’s human representatives with algorithms. It is posited that AI and other forms of technology, if considered and applied in the right way, can move society forward, so long as technological innovations can provide a positive enhancement of core ideas—like democracy and the Rule of Law. Some innovations can facilitate augmented democracy, and augmenting democracy even in its most extreme forms can be beneficial. These forms of technological innovation, and others like them, regardless of how extreme they may appear, should remain open to interrogation.

Bert Heinrichs, in ‘Discrimination in the Age of Artificial Intelligence’ (this volume), examines whether the use of artificial intelligence (AI) and automated decision-making (ADM) aggravates issues of discrimination, and argues that the use of AI/ADM can, in fact, increase issues of discrimination, but in a different way than most critics assume. It is due to its epistemic opacity that AI/ADM threaten to undermine our moral deliberation, which is essential for reaching a common understanding of what should count as discrimination. As a consequence, it turns out that algorithms may actually help to detect hidden forms of discrimination. Against this background, research initiatives for explainable AI are especially important from an ethical point of view.

Marcus Smith and Seumas Miller, in ‘The Ethical Application of Biometric Facial Recognition Technology’ (this volume), examine the rise of biometric facial recognition, current applications, and legal developments, and conduct an ethical analysis of the issues that arise. It is recognised that biometric facial recognition technology gives rise to security concerns, such as the possibility of identity theft by a sophisticated malevolent actor, even as they resolve old privacy and confidentiality concerns, such as by reducing unauthorised access to private information and thereby strengthening privacy protection. However, the authors conclude that the problems in this area cannot be framed in

terms of a simple weighing of, let alone trade-off between, individual privacy rights versus the community’s interest in security. In view of the expanding use of biometric facial recognition for security and public safety, the paper outlines relevant ethical principles and identifies a number of actual or potential problems that arise in relation to this rapidly developing form of information technology.

Mario Günther and Atoosa Kasirzadeh, in ‘Algorithmic and Human Decision Making’ (this volume), pose a question of whether decision-making algorithms be held to higher standards of transparency than human beings. The paper argues that the answer depends upon what we demand from explainable algorithms, how we govern them via regulatory proposals, and how explainable algorithms may help resolve the social problems associated with decision-making supported by artificial intelligence. In recognition of the debate on the same or double standards of transparency for humans and algorithms, the authors put forth two arguments for how and when a double standard is justified, the first is the need to take design explanations into account with respect to algorithmic decision-making and the second is that the intentional stance does not apply to proper black box algorithms. The paper suggests that the next steps of research are a systematic exploration of the classes of algorithmic decision-making scenarios that require a higher standard of transparency, and articulation of how the algorithmic governance and regulatory proposals would look like in cases of the double standard of transparency.

Nathan Colaner, in ‘Is Explainable Artificial Intelligence Intrinsically Valuable?’ (this volume), suggests that in addition to asking for value and desirability of explainable artificial intelligence (“XAI”), we need to ask: How do we develop technical strategies to achieve XAI? And what kind of explanation is worth having in the first place? Although XAI is desirable to attain some other value, such as fairness, trust, accountability, or governance, the author argues that it is also crucial to consider the legal and ethical values that may be undermined by use of such models, such as trust and fairness. Rather than depending on the notoriously elaborate and evasive concept of freedom and/or exploring precisely what autonomy entails and where its limits lie, the author suggests that we can focus on exploring the availability of real opportunities and capacities to resist coercion, and this focus can be used not only in evaluating past and prevailing moral relations, but also in the moral dimension of decision-making for the future.

Donghee Shin, in ‘How Do People Judge the Credibility of Algorithmic Sources?’ (this volume), discusses the need for the creation of understandable/explainable AI for establishing trust and credibility by engaging human agency into the AI. In his study on chatbots, the author examines users’ belief in credibility of chatbot’s information, and notes that users’ algorithmic literacy as well as users’ trust play a

pivotal role in how users form perceptions of the credibility of chatbot messages and recommendations. This insight on credibility, it is suggested, provides a better foundation for algorithm design and a stronger basis for the design of sensemaking chatbot journalism.

Simona Chiodo, in ‘Human autonomy, technological automation (and reverse)’ (this volume), sheds light on the question of what is essentially being human and what is being autonomous. He says that on one hand, we have the notion of (human) ‘autonomy’, meaning that there is a ‘law’ that is ‘self-given’, and, on the other hand, we have the notion of (technological) ‘automation’, meaning that there is something ‘offhand’ that is ‘self-given’. Yet, we are experiencing a kind of twofold shift: on one hand, the shift from defining technologies in terms of automation to defining technologies in terms of autonomy and, on the other hand, the shift from defining humans in terms of autonomy to defining humans in terms of automation. This shift raises a concern that we may be using technologies, and in particular algorithmic technologies, as scapegoats that bear responsibility for making decisions for us. Or, we may argue that we create a kind of technological divine that, by being always with us through its immanent omnipresence, omniscience, omnipotence, and inscrutability, that can always be our technological scapegoat freeing us from the most unbearable burden of individual responsibility resulting from individual autonomy.

Bernd Carsten Stahl et al., in ‘Organisational Responses to the Ethical Issues of Artificial Intelligence’ (this volume), note that whilst there is an abundance of conceptual work on AI ethics, empirical insights are rare and often anecdotal. It is also generally unclear how organisations that make use of AI understand and address these ethical issues in practise. They posit that although organisations are highly aware of the AI ethics debate and keen to engage with ethical issues proactively, many of the ethical issues are seen to be either beyond the organisations’ expertise or lie outside their remit. This raises a question of what exactly lies within the remit of organisations and which issues and measures are the responsibility of policymakers, professional bodies, and regulators that needs also to be addressed. The authors argue for a broader framework that engages a diversity of stakeholders to ensure a more comprehensive coverage of AI ethics.

Fred Fonseca, in ‘Data Objects for Knowing’ (this volume), commenting on the claim that Data Science does not need theories, argues that technology-driven science created a need for more technology-driven science, culminating in data science. Data science is thus an experimental science, which uses data objects in experiments, and these objects are called ‘data-objects-for-knowing’. The author concludes that data science is a science to study artificially created phenomena—a science to study the data manipulated by the equations and operations of AI. It is highlighted that Data

Science is really the science of data and data are its phenomena. The data are experimented with regardless of the theories that generated it. In this sense, data Science disregards the connections between data and the real world that were carefully built by the theories from other sciences. In the experiments of data science, data are the world itself. The knowledge created by data science is purposely disconnected from any theory from other sciences; it is a knowledge for the sake of itself. In other words, the purpose of data in Data Science is to understand not *what* the data mean in the context of the world, but rather what the data say about *itself*.

Sylwia Wojtczak, in ‘Endowing Artificial Intelligence with legal subjectivity’ (this volume), commenting upon the possible participation or presence in social life, argues that despite the potential dangers associated with endowing AI with some kind of subjectivity, such a course is inescapable, and should be considered sooner rather than later. However, the author notes that social recognition is neither necessary nor enough for legal personhood, but it means that the lack of social recognition is a crucial obstacle for untypical legal persons. The paper offers possible options to deal with the question of legal personhood—connect legal subjectivity with some financial autonomy of the entity; legal rules demand safety and explainability by default; bundle Theory of Subjectivity—adjust the scope of subjectivity to practical needs, by only assigning the AI competences, claim-rights or duties that are acceptable, useful, and safe. The author concludes that legal science should avoid falling into ideological boost or simply the ‘guarding’ tradition.

Martin Gibert and Dominic Martin, in ‘In Search of the Moral Status of AI’ (this volume), discuss different arguments for granting moral status to an AI system: the idea of indirect duties, the relational argument, the argument from intelligence, the argument from life and information, and the argument from sentience. The paper argues that sentience, by contrast to other arguments, provides a strong argument for the moral status of an AI system, pointing out, however, that no AI system is sentient given the current level of technological development. Drawing upon the use of the argument from sentience to defend moral status of AI systems in science fiction storeys, it is posited that that the sentient argument captures widely held judgments. It thus avoids: (i) the constraint of the idea of indirect duties to consider an AI system for its own sake; (ii) the issues of consistency or objectivity of the relational argument, and (iii) the criteria of just being intelligent in place of sentience also being a relevant criterion of moral status for human beings; (iv) Although one can grant a moral status to an AI system on the basis that this system is a living or information entity, the argument may grant only a minimal moral status, one that is easily overridable. The paper asserts that the argument from sentience, however, could lead us to grant moral status to an AI system at the point where it would be considered

equal to a human in Sparrow's triage test. The argument from sentience is thus the strongest for defending the moral status of an AI.

Luciano Floridi et al. in 'The Ethics of Algorithms' (this volume), undertake a review of ethics of algorithms with a view to contribute to the debate on the identification and analysis of ethical implications of algorithms, provide an updated analysis of epistemic and normative concerns, and offer actionable guidance for the governance of the design, development, and deployment of algorithms. The authors conclude that ethical analyses are necessary to mitigate the risks whilst harnessing the potential for good of these technologies.

Simon Penny, in 'Sensorimotor debilities in digital cultures' (this volume), reflects on the qualities of living and learning in digital cultures, the design of digital technologies, and the philosophical history that has informed that design. Here, 'sensorimotor debility' is described as a condition, especially amongst users of digital technology, arising from cognitive, neurological and physiological effects. The paper argues that the longstanding Enlightenment-humanist privileging of reason and of abstraction, combined with the emergence of a technology of abstract symbol manipulation, and neoliberal educational agendas that slash 'soft' or 'applied' aspects of learning are actually valorizing abstract symbol manipulation and are thus creating a perfect storm for sensorimotor competence. The author proposes that leveraging postcognitivist, embodied, enactive and distributed approaches to cognition to analyse human computer interaction can provide new insights into growing social and public-health concerns around emerging computer-use issues. In reasserting the holism of the cognizing organism, such a leverage destabilises axiomatic assumptions about the separability of mind and body, and thus of intelligence and skill.

Tatsuya Nomura and Motoharu Tanaka, in 'Experiences, Knowledge of Functions, and Social Acceptance of Robots' (this volume), undertook an exploratory case study focusing on Japan. The study suggests that although Japanese society has become aware of some types of robots, social acceptance of robots is still not widespread. The authors note that although the results of the survey revealed differences between some robot types on their acceptances and relationships with experiences and knowledge in the current stage, it cannot be clarified whether and how these acceptances and experiences changed since the introduction of the Japanese government's "New Robot Strategy".

Caroline L. van Straten et al., in 'The Wizard and I' (this volume), undertook an experimental study on the effects of transparent teleoperation (using a Wizard of Oz exemplar) and self-description on children's perception, and relationship formation with a robot. The finding of the study indicates that children may consider robots

as potential friends regardless of their knowledge of the robot's teleoperated working and its engagement in self-description. The authors conclude that a societally important implication of this finding is that it may be possible to reach potential benefits of child-robot relationship formation (e.g., in education and healthcare applications) without 'deceiving' children into thinking robots are more capable and social than they currently are. They suggest that future research should investigate whether their findings with respect to the emergence of children's initial sense of relationship with robots extend to situations in which children interact with robots on a long-term basis. It is proposed that further elucidation of the boundary conditions of child-robot relationship formation would advance our understanding of the characteristics of robots that are necessary or sufficient to support children—whether as a complement or a temporary replacement of interpersonal interaction.

Nidal Al Said and Khaleel M. Al-Said, in 'The effect of visual and informational complexity of news website designs on comprehension and memorisation amongst undergraduate students' (this volume), study how the basic web designs aesthetically affect users. The study engaged students from Arab universities to determine their levels of perception and recall. The results revealed that interactive sites with multiple aesthetic elements representing the message of the news item enable users to perceive and remember the information better. The findings of this study can be useful in creating news website templates that improve user comprehension and recall.

Jon Dron in 'Educational technology: what it is and how it works' (this volume), elucidates the nature of educational technology and, in the process, sheds light on a number of phenomena in educational systems, from the no-significant-difference phenomenon to the singular lack of replication in studies of educational technologies. The author concludes that education is, primarily, not a process of instilling skills and facts, but of preparing human beings to live, work, and play with other humans in society. It is as fundamentally human as art and, just as it would make little sense to build a machine to make art, it makes little sense to build a machine to educate. Just as machines can extend and enable what an artist can create, so machines can support the educational process, but it is not the machine itself that achieves this. It is the ways the case that the machine is orchestrated by humans, with humans, and for humans that makes it educational. We are all co-participants in this deeply human, highly distributed educational machine, not just users but—necessarily—both creators and parts of its ever unfolding form. Being parts of machines is part of what it means to be human, and being part-human is part of what it means to be a machine. If we can better understand how the machines work then, as co-participants in them, we can make each

one a thing of beauty and value rather than a vehicle of oppression.

Ahmed Izzidien, in ‘Word Vector Embeddings Hold Social Ontological Relations Capable of Reflecting Meaningful Fairness Assessments’ (this volume), explores the challenge of defining cross-cultural fairness assessments of texts through top–down rules, bottom–up training, and hybrid approaches. The paper proposes the plausibility of using Word Embedding vectors to make fairness assessments. This approach is based on the premise that an inherent aversion to harmful gainless activity introduces a pro-social bias into Word Embeddings, whereby acts that meet this propensity are qualified as being closer in the vector space to the latent concept of fairness. Here, the recognition of loss, pain, and punishment are seen as blameworthy, whereas gain, joy, and liberty as praiseworthy, but only when filtered according to their associated score of being responsible, or irresponsible, respectively. In this approach, the ethical assessment relies on the efficacy of Word Embeddings, rather than deontic rules or training data.

Marion Maisonobe, in ‘The future of urban models in the Big Data and AI era’ (this volume), examines the effects on urban research dynamics of the Big Data and AI in urban management, in relation to two urban systems: transportation and water. The author argues that although transportation studies are increasingly focussing on AI and Big Data and that traffic flow studies are arousing a growing interest amongst computer scientists, this interest is less pronounced in the water research area, and more especially regarding water quality. The differences observed between research on transportation and that on water confirm the multifaceted nature of the developments at work and encourage us to reject overly hasty and simplistic generalisations about the transformations underway.

Nader Ghotbi et al., in ‘Attitude of college students towards ethical issues of artificial intelligence in an international university in Japan’ (this volume), undertook a survey of college students in a multicultural university and suggest that their worry over unemployment is the most serious concern about AI technologies. However, a sentiment analysis of the texts written by students demonstrated a generally more positive attitude towards AI. Trust was the most common emotion, which may sound naïve considering the concern of AI experts over the use of AI for surveillance and a decline of privacy for citizens of the future. The second most common emotion was fear, which reflects the concern shared by AI experts. The authors note that an interesting finding of this study was that many students are concerned over emotional AI issues. These include the impact of AI on human behaviour as well as the emotions and rights of future robots. The authors suggest that AI engineers may want to consider more seriously the emotional aspects of AI in research and development.

Rashid Minhas et al., in ‘Protecting Victim and Witness Statement’ (this volume), examine the effectiveness of a chatbot (the AICI), that uses artificial intelligence (AI) and a cognitive interview (CI) to help record statements following an incident. The authors point out that the AICI provides means of effectively and efficiently recording high-quality evidential statements from victims and witnesses as an alternative to the traditional information collection procedures in forensic investigations.

Antonio Jimenez-Mavillard and Juan Luis Suarez, in ‘A computational approach for creativity assessment of culinary products’ (this volume), propose an AI-based method (Random Forest Classifier) for assessing culinary product creativity using the renowned high cuisine restaurant elBulli as a case study to understand the proliferation and scale of an entity's creativity and innovation. To limit the scope of the analysis and make answering the research questions feasible, the approach included specific propositions that resulted in the assumption that elBulli's recipes can be binded by their list of ingredients and techniques; and their level of creativity can be assessed by estimating their creation year. Seeing creativity assessment of culinary products as art, the authors posit how in the areas of art and fashion, new trends in fashion can influence the genre of department stores and, therefore, our closets; or how specific avant-garde artistic manifestations are present in exhibitions and art galleries in the world.

Lea Köstler and Ringo Ossewaarde, in ‘The Making of AI Society: AI Futures Frames in German Political and Media Discourses’ (this volume), shed light on the emergence, diffusion, and use of socio-technological future visions of AI within a German context. They note that the German Government perspective is that the German AI future is an unquestioned reality to which the entire German nation must adapt, and that German citizens are expected to adjust their roles within the framework provided by the German government. AI is framed as the cure-all for present and future problems in Germany (“*AI as panacea*”). In the German government's envisioning of the AI future, the German past is projected on the German future. However, the German media question basic assumptions concerning the balance of power in the future, and suggest that there is little political-administrative willingness to design German AI futures that significantly diverge from the past or present. The authors suggest that the greatest danger in the German AI future stems from insufficient anticipation, inaction, and the threat of a declining German economy (“*Uncertainty as main menace*”). The authors argue that further research on dominant AI future visions and the frames integral to them are needed to examine political, corporate, and societal interests in detail. The ultimate aim of this research should be to finally start an open debate on the use of technological

innovation in the creation of possible futures that are, first and foremost, filled by public interests.

Michael R. Scheesele, in ‘The hard limit on human non-anthropocentrism’ (this volume), argues that our existential concerns about super AIs, achieving human-level machine intelligence, may make us reflect on the limit of human anthropocentric tendencies of control, and this may prove to be for our own good.

Štěpán Cvik, in ‘Categorization and Challenges of Utilitarianisms in the Context of Artificial Intelligence’ (this volume), suggests that although, in machine ethics, there is a tendency to divide utilitarianisms into two categories of static utilitarianisms and dynamic utilitarianisms, there is a need to explore the possibility of using a combination of the two categories of utilitarianisms to resolve most of the challenges without the need to abandon the concept of utilitarianisms.

Malesela John Lamola in ‘The Future of Artificial Intelligence, Posthumanism and the Inflection of Pixley Isaka Seme’s African Humanism’ (this volume), explores the Eurocentric genitive basis of the philosophical anthropology that underpins the technological post-humanism. As an alternative to the dystopian conception of post-humanism, the author proposes an Africanist civilizational humanism proclaimed by Pixley ka Isaka, Seme, as a plausible alternative

paradigm for humanity’s technological advancement. Seme’s is being provided as a contribution to the archive on cross-cultural ethics of artificial intelligence. Concerned with the de-centring of the human in both her sociality and biological being as a feature of a technological age, the author argues that Seme’s humanistic ethical-spiritualist conception of technological progress constitutes a credible contribution to the debate and search for the ideal state of being human in the rapidly evolving technological age.

Karamjit S Gill

Editor, AI&Society

January, 2022

## References

- Cooley M (2020) *The Search for Alternatives: Liberating Human Imagination*. Spokesman Books, Nottingham
- Nowotny H (2021) *IN AI WE TRUST: power, illusion and control of predictive algorithms*. Polity Press, Cambridge

**Publisher’s Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.