



Artificial Intelligence for Social Good: Avoiding the Solutionist Trap

EU INSTITUTIONS & POLICIES

Andreas TSAMADOS

PhD Researcher, Oxford Internet Institute, University of Oxford



July 2021

Policy Paper #76/2021

Artificial Intelligence for Social Good: Avoiding the Solutionist Trap

Copyright © 2021 | All Rights Reserved
HELLENIC FOUNDATION FOR EUROPEAN & FOREIGN POLICY (ELIAMEP)
49, Vasilissis Sofias Ave., 10676, Athens, Greece
Tel.: +30 210 7257 110 | Fax: +30 210 7257 114 | www.eliamep.gr | eliamep@eliamep.gr

ELIAMEP offers a forum for debate on international and European issues. Its non-partisan character supports the right to free and well-documented discourse. ELIAMEP publications aim to contribute to scholarly knowledge and to provide policy relevant analyses. As such, they solely represent the views of the author(s) and not necessarily those of the Foundation.

Andreas TSAMADOS

PhD Researcher, Oxford Internet Institute, University of Oxford

Summary

- AI can play a meaningful role in solving global problems that are increasingly interconnected and complex.
- The distinct advantages of predictions, classifications and decisions made by AI systems can introduce unique challenges just as it can exacerbate existing problems in society.
- AI for social good still remains poorly understood as a global phenomenon. It is also missing a robust framework for assessing the value and the success of such efforts.
- To seek to exploit the potential of AI to help the international community meet some of its most urgent issues without a sound governance framework can distort one's assessment of the realistic and worthwhile opportunities at hand.
- Because of the reach and quasi universal status of the UN Sustainable Development Goals, projects seeking to leverage AI for socially beneficial outcomes –as well as researchers monitoring them– can use the SDGs as a starting point and assessment benchmark.
- The existing body of research on SDGs already includes studies and metrics on how to measure progress in attaining each of the 17 SDGs, and the 169 associated targets defined in the 2030 Agenda for Sustainable Development. These metrics can be applied to evaluate the impact of AI×SDGs projects.
- The analysis indicates that the use of AI×SDGs is an increasingly global phenomenon—with projects operating from five continents—but also that the phenomenon may not be equally distributed across the SDGs. SDG 3 ('Good Health and Well-Being') leads the way, while SDGs 5 ('Gender Equality'), 16 ('Peace, Justice and Strong Institutions'), and 17 ('Partnerships for the Goals') appear to be addressed by fewer than five projects.

Introduction

Artificial Intelligence (AI) is slowly moving past the public bewilderment phase that is common to most revolutionary technologies. In the past decade, AI has often been portrayed as either a sentient technology with an apocalyptic destiny, or as a silver bullet for most if not all conceivable problems ranging from the climate emergency to global crime prevention. This article contributes to a more recent and realistic assessment that is beginning to see the light: AI is a powerful technology that can help us address a wide range of well-defined problems given proper design, development, integration and monitoring.

“AI can play a meaningful role in solving global problems that are increasingly interconnected and complex.”

Building on previous research presented in (Cowls *et al.*, 2021a) and (Cowls *et al.*, 2021b), this article analyses the complex issue of using AI to deliver socially beneficial outcomes. In the first two sections I offer an overview of the various limitations of AI systems and of existing projects seeking to use AI for a loosely defined “social good”. In turn, I argue that despite these socio-technical challenges, AI can play a meaningful role in solving global problems that are increasingly interconnected and complex. I dedicate the last section of this article to presenting the results of a two year-long research (Cowls *et al.*, 2021a) on AI that supports the United Nations’ Sustainable Development Goals (UN SDGs), and advocating for the use of the UN SDGs as a benchmark against which to assess AI for social good.

Neither sentient nor perfect

To start thinking about the potential of this powerful new technology it is important to first agree on what we mean when we use the term “AI”. Here, I chose to align the nomenclature and definition with that of the European Commission, which has since 2016 proven to be leading the international charge with regards to AI strategies and regulations (Roberts *et al.*, 2021). The European Unions’ (EU) High-Level Expert Group on AI (AI HLEG), provides (AI HLEG, 2019) a well-rounded definition of AI, which can be broken down as follows:

“AI is not some sentient entity capable of taking over the planet if it “escaped human control” as some headlines would have you think.”

- i) software (and possibly also hardware) systems
- ii) designed by humans that,
- iii) given a complex goal,
- iv) act in the physical or digital dimension by perceiving their environment through data acquisition,
- v) interpreting the collected structured or unstructured data,
- vi) reasoning on the knowledge, or processing the information, derived from this data
- vii) and deciding the best action(s) to take to achieve the given goal.

AI is not some sentient entity capable of taking over the planet if it “escaped human control” as some headlines would have you think (Friend, 2018). One way to understand the general limitations of AI is by thinking of the technology as capacity without comprehension. An AI system may give you the impression that it can paint by using its own imaginative powers, or come up with its own ideas and arguments, but these feats are not understood by the system itself. Instead, the system simply produces outputs that are expressed in probabilistic terms (Tsamados *et al.*, 2021). For example, during my research on mass produced political propaganda, I trained a transformer language model (a form of AI) to write poems in the style of Konstantinos Kavafis. The system was able to create poems quite similar to those of Kavafis, not because it understood how to evoke vivid

“The distinct advantages of predictions, classifications and decisions made by AI systems can introduce unique challenges just as it can exacerbate existing problems in society. Acknowledging and addressing some of these (depending on the context) is a precondition to any projects or effort seeking to leverage AI for socially beneficial outcomes.”

images and emotions like the literary genius did, but because every time I gave the system a word (eg “Church”), it would simply look at everything Kavafis ever wrote and calculate what word had the highest probability of being used after “Church” in Kavafis’s corpus.

As more AI systems are integrated into critical infrastructure and virtually all areas of human activity, from electrical grids to hospitals to choosing friends and music on the Internet, the record of those systems failing and/or harming people keeps growing at a worrying rate, (Tsamados *et al.*, 2021). The distinct advantages of predictions, classifications and decisions made by AI systems can introduce unique challenges just as it can exacerbate existing problems in society. Acknowledging and addressing some of these (depending on the context) is a precondition to any projects or effort seeking to leverage AI for socially beneficial outcomes. For the purpose of this article, I will not list all possible ethical risks involved in the development and integration of AI systems (see (Tsamados *et al.*, 2021) for more), but I will present what I consider to be most relevant to the ensuing discussion on AI for social good.

First, what we describe above in (v) as interpreting data can create issues of unwanted bias and discrimination. This can be due to the data not being accurate, timely or adequate for a specific context. It can also be due to the flawed assumptions that guided the collection and curation of that data, or of data being representative of an already unfair and discriminatory system (Tsamados *et al.*, 2021). No amount of data –if they suffer from the above flaws– fed to a system will be able to somehow fix the outputs of that system. People have been treated differently in hospitals or given unjustly longer prison sentences because of the color of their skin and due to AI systems’ perpetuation of long-standing structural inequalities present in their training data (Benjamin, 2019; Obermeyer *et al.*, 2019).

A second and related issue concerns the risk of AI infringing on human autonomy (Cows *et al.*, 2021b). Our increasing reliance on recommendations provided by AI systems, such as in the intake of political information or even the allocation of government welfare programs, may have severe implications for human agency. The illusion of mechanistic objectivity we have wrapped around “smart machines” has only served to perpetuate this loss of autonomy.

Third, literature on AI has long grappled with the issue of transparency (Ananny and Crawford, 2018). The opacity of certain AI systems can be due to various reasons, including the obfuscation of training data and the lack of tools to visualise and track the dataset or evolving parameters of the system. In turn, this can lead to a lack of accountability and the erosion of trust in those systems or their developers (Fink, 2018).

Fourth, the collection of personal data to train AI systems poses serious challenges to the privacy of individuals and groups (Taylor, Floridi and van der Sloot, Forthcoming). Even for unequivocally good causes such as tracking the spread of the COVID-19 virus, the collection and use of personal data is a delicate issue with fundamental rights implications that the EU’s General Data Protection Regulation covers extensively.

Fifth, the current development of AI requires a significant amount of energy (Taddeo *et al.*, 2021). This is especially the case for the “training” process of AI systems as most breakthroughs in AI since 2012 have required an ever growing amount of computing power. From 2012 to 2018, the computing power used to train large AI systems has doubled every 3.4 months (Amodei and Hernandez, 2018). This means that if the energy-efficiency of these systems and their supporting infrastructure (eg. data centers and electrical grids) do not follow at the same pace as the growth of computing power

“From 2012 to 2018, the computing power used to train large AI systems has doubled every 3.4 months. [...] This means that if the energy-efficiency of these systems and their supporting infrastructure do not follow at the same pace as the growth of computing power demand, we could see the carbon footprint of AI research and development grow exponentially.”

demand, we could see the carbon footprint of AI research and development grow exponentially. An associated issue that is worth mentioning is that most of AI research and development happening today is situated in Western industrial nations, thus perpetuating the environmental injustice already inflicted by these nations’ through colonial exploitation and mass pollution (Cho, 2020; Taddeo *et al.*, 2021).

The sixth and last risk we are concerned with in this article is that of good governance, or lack thereof, on AI –specifically AI for social good. This can be considered a meta-risk with the potential to magnify other risks and challenges. The large number of “AI for social good” projects being deployed across the world and experimenting at a rapid pace with new applications of AI creates a risk that “harmful practices “fall through the cracks” and have adverse effects on populations and/or the environment (Cowls *et al.*, 2021b). The following section will explore aspects of the multi-faceted issue of good AI governance. This will then allow the article to propose a coherent framework to improve efforts seeking to leverage AI for socially beneficial outcomes.

Governance gaps and the solutionist trap

Since the advent of deep learning –a form of AI– in 2012, the research on and development of AI have been experiencing a new ‘summer’, both in terms of the quantity of applications enabled by the technology and the interest it has received from governments, academics and civil society alike (Perrault *et al.*, 2019). Commercial and non-commercial projects purporting to leverage AI to deliver socially beneficial outcomes have also grown rapidly (McKinsey, 2018; Vinuesa *et al.*, 2020). These are projects that use the underlying technology, in some form or another, to deliver socially good outcomes, which were previously either “unfeasible, unaffordable or simply less achievable in terms of efficiency and effectiveness” (Cowls *et al.*, 2021a).

Indeed, AI has been used to reduce the energy consumption of data centers; to produce reliable forecasts on the spread of forest fires or the melting of Arctic ice using satellite imagery; or even giving people in certain regions of the world access to health services and basic diagnoses via their mobile phones (Cowls *et al.*, 2021a, 2021b). However, AI for social good still remains poorly understood as a global phenomenon. It is also missing a robust framework for assessing the value and the success of such efforts. This is a governance issue that requires a clear understanding of the aforementioned limitations and risks of AI and of the type of thinking (mis)guiding many projects leveraging AI. It also requires international alignment on a set of well-defined and realistic objectives.

AI for social good has so far been developed on an ad hoc basis whereby only the specific areas of application are analysed (Vinuesa *et al.*, 2020). As we explain in (Cowls *et al.*, 2021a):

“this approach can indicate the presence of a phenomenon, but it cannot explain what exactly makes AI socially good nor can it indicate how [AI for social good] solutions could and should be designed and deployed to harness the full potential of the technology”.

In turn, these issues raise risks of projects failing in an unexpected manner, of opportunities being neglected, of unwarranted interventions, and of “ethics dumping”. Ethics dumping refers to the exploitative tendency to perform unethical research or use unethical research practices in countries with either less rules designed to protect research subjects or less capacity to monitor infringement and enforce the rules, compared to richer

countries (Floridi, 2019).

“To seek to exploit the potential of AI to help the international community meet some of its most urgent issues without a sound governance framework can distort one’s assessment of the realistic and worthwhile opportunities at hand.”

To seek to exploit the potential of AI to help the international community meet some of its most urgent issues without a sound governance framework can distort one’s assessment of the realistic and worthwhile opportunities at hand. The aforementioned risks often go hand-in-hand with the so-called “tech solutionism” or “AI solutionism” mentality. To paraphrase Evgeny Morozov (Morozov, 2013) who popularised the concept of “tech solutionism”, solutionism can be understood as “*an intellectual pathology*” that recasts

“all complex social situations either as neatly defined problems with definite, computable solutions or as transparent and self-evident processes that can be easily optimized — if only the right algorithms are in place!”.

This leads people to reach for the known solution (in this case AI) to address anything resembling a problem, regardless of the context. Another tendency of people or groups espousing this mentality is to “avoid engaging with unsolvable problems” or to mis-categorise them as solvable and thus “producing cascades of unintended consequences” (Morozov, 2013).

These challenges are not fatalities and will not necessarily result in harm. However, it is crucial to put them front and center in discussions on AI for social good and good governance of AI, if only to inform as many people as possible (policy-makers especially) about the difficulty of getting this right. It is worth mentioning that studies have shown that in a large majority of top AI research papers, “considerations of negative consequences are extremely rare” (Birhane *et al.*, 2021) .

“The global and existential threats, such as the climate emergency and pandemics, we are facing will require ambitious, even radical responses from the international community. However, neither the scale of these threats nor our hopeful -if not naive- view of the potential of AI should blind people (especially policymakers) to the importance of existing and evolving rules and social expectations over the use of potentially high-risk technology.”

The global and existential threats, such as the climate emergency and pandemics, we are facing will require ambitious, even radical responses from the international community. However, neither the scale of these threats nor our hopeful –if not naive– view of the potential of AI should blind people (especially policymakers) to the importance of existing and evolving rules and social expectations over the use of potentially high-risk technology (Cowls *et al.*, 2021b).

Common and fertile ground: the UN SDGs and AI

Despite all the aforementioned limitations and risks, discarding the potential of AI for socially beneficial outcomes and as a technology that can help align international efforts to solve global problems would be equivalent to shooting ourselves in the foot. As such, it presents the findings of our (Cowls *et al.*, 2021a) two year long research on AI for social good, beginning by proposing a sound definition of what can be considered AI for social good. The section also argues that the 17 UN SDGs establish a robust benchmark against which the international community can assess and monitor such AI-focused projects that seek to produce socially beneficial outcomes.

As my colleagues and I have argued (Cowls *et al.*, 2021a), AI for social good projects can be analysed on the basis of their outcomes. The project can be considered successful insofar as it manages to help mitigate or even solve a given issue while not creating new ones or exacerbating existing ones. Hence, AI for social good is formally defined (Cowls *et al.*, 2021a) as “the design, development and deployment of AI systems in ways that help to:

“The project can be considered successful insofar as it manages to help mitigate or even solve a given issue while not creating new ones or exacerbating existing ones.”

- i) prevent, mitigate and/or resolve problems adversely affecting human life and/or the wellbeing of the natural world, and/or
- ii) enable socially preferable or environmentally sustainable developments, while
- iii) not introducing new forms of harm and/or amplifying existing disparities and inequities.

The next step in establishing a robust framework and assessment benchmark on AI for social good becomes the “effective identification of problems that are deemed to affect human life or the wellbeing of the environment negatively” (Cowls *et al.*, 2021a). For a number of reasons, the 17 UN SDGs provide fertile ground to do so. Perhaps the most important factor is that these goals are shared by 193 nations and enjoy a quasi universal status.



Fig 1. The 17 UN Sustainable Development Goals.

The SDGs were signed by the international community to provide the world with one common set of goals that cut across economic, social, and environmental dimensions of sustainable development. Because of the reach and popularity of these goals, projects seeking to leverage AI for socially beneficial outcomes –as well as researchers monitoring them– can use the SDGs as a starting point and assessment benchmark.

“...the SDGs offer clear, well defined and shareable boundaries to identify positively what is socially good AI (what should be done, as opposed to what should be avoided), although they should not be understood as indicating what is not socially good AI.”

Analysis

The UN SDGs constitute a sufficiently empirical and reasonably uncontroversial benchmark to evaluate the positive social impact of AI for social good globally. This move, to set “AI for social good” as AI that supports Sustainable Development Goals (AIXSDGs), may seem restrictive because there are undoubtedly a multitude of examples of socially good uses of AI outside the scope of the SDGs. Nonetheless, the approach carries clear advantages, of which five are paramount.

First, the SDGs offer clear, well defined and shareable boundaries to identify positively what is socially good AI (what should be done, as opposed to what should be avoided), although they should not be understood as indicating what is not socially good AI.

Second, the SDGs are internationally agreed goals for development, and have begun informing relevant policies worldwide, so they raise fewer questions about relativity and cultural dependency on values. Although they are of course improvable, they are nonetheless the closest thing we have to a humanity-wide consensus on what ought to be done to promote positive social change and the conservation of our natural environment.

Third, the existing body of research on SDGs already includes studies and metrics on how to measure progress in attaining each of the 17 SDGs, and the 169 associated targets defined in the 2030 Agenda for Sustainable Development. These metrics can be applied to evaluate the impact of AI×SDGs projects (Vinuesa *et al.*, 2020).

Fourth, focusing on the impact of AI-based projects across different SDGs can improve existing, and lead to new, synergies between projects addressing different SDGs, further leveraging AI to gain insights from large and diverse datasets, and can pave the way to more ambitious collaborative projects.

Fifth and finally, understanding AI4SG in terms of AI×SDGs enables better planning and resource allocation, once it becomes clear which SDGs are under-addressed and why.

In view of the advantages of using the UN SDGs as a benchmark for assessing AI4SG, we conducted an international survey of AI×SDG projects. The survey ran between July 2018 and November 2020, and it involved collecting data on AI×SDG projects that met the following five criteria:

“...focusing on the impact of AI-based projects across different SDGs can improve existing, and lead to new, synergies between projects addressing different SDGs, further leveraging AI to gain insights from large and diverse datasets, and can pave the way to more ambitious collaborative projects.”

- i) Only projects that actually addressed (even if not explicitly) at least one of the 17 SDGs;
- ii) Only real-life, concrete projects relying on some actual form of AI (symbolic AI, neural networks, machine learning, smart robots, natural language processing, and so on), not merely referring to AI (an observed problem among AI start-ups more generally);
- iii) Only projects built and used in the field for at least six months, rather than theoretical projects or research projects yet to be developed (for example, patents, or grant programmes);
- iv) Only projects with documented positive impact, for example through a web site, a newspaper article, a scientific article, a non-governmental organization (NGO) report, and so on;
- v) Only projects with no or minimal evidence of counter-indications or negative side-effects.

From a larger pool, the survey identified 108 projects in English, Spanish and French matching these criteria. The data about the AI×SDG projects collected in this study are publicly available in the aforementioned database (<https://www.aiforsdgs.org/all-projects>), which is part of the Oxford Research Initiative on Sustainable Development Goals and Artificial Intelligence. We presented a preliminary version of our analysis in September 2019 at a side-event during the annual UN General Assembly.

Our analysis (see below) shows that every SDG is already being addressed by at least one AI-based project. The analysis indicates that the use of AI×SDGs is an increasingly global phenomenon— with projects operating from five continents—but also that the phenomenon may not be equally distributed across the SDGs (Fig. 2). SDG 3 (‘Good Health and Well-Being’) leads the way, while SDGs 5 (‘Gender Equality’), 16 (‘Peace, Justice and Strong Institutions’), and 17 (‘Partnerships for the Goals’) appear to be addressed by fewer than five projects (see Fig. 2). It is important to note that the use of AI to tackle at least one of the SDGs does not necessarily result in success. We also note that a project could address multiple SDGs simultaneously or at different timescales and in different ways. Moreover, even complete success for a given project would be exceedingly unlikely to result in the eradication of all of the challenges associated with an SDG. This is chiefly because each SDG concerns entrenched challenges that are widespread and structural in nature. This is well reflected by the way SDGs are organized: all 17 SDGs have several

targets, and some targets in turn have more than one metric of success. The survey shows which SDGs are being addressed by AI at a high level, but more finely grained analysis is required to assess the extent of the positive impact of AI-based interventions with respect to specific SDG indicators, as well as possible cascade effects and unintended consequences.

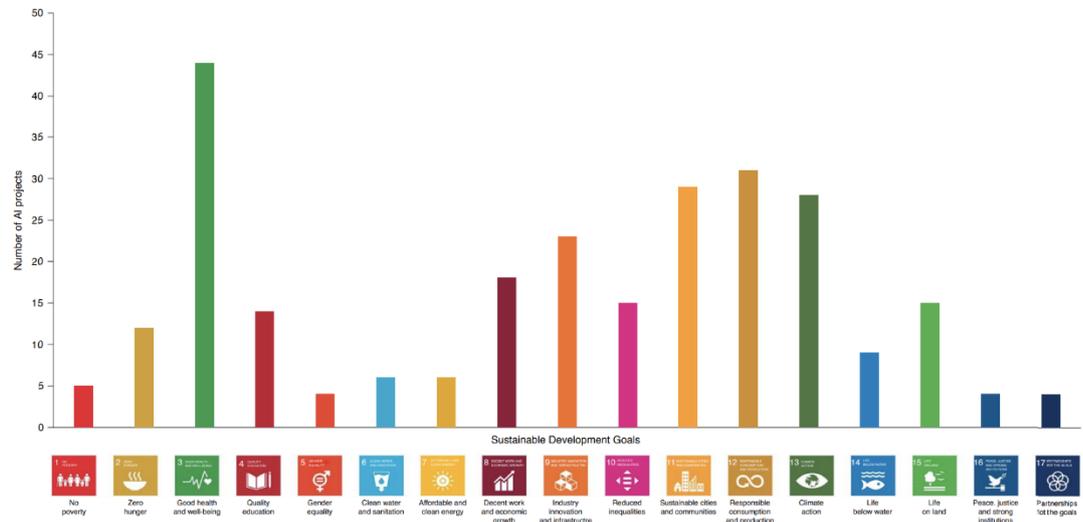


Fig. 2 Projects addressing the SDGs. A survey sample of 108 AI projects found to be addressing the SDGs globally.

The analysis suggests that the 108 projects meeting the criteria correspond with seven essential factors for socially good AI, identified in (Floridi *et al.*, 2020): falsifiability and incremental deployment; safeguards against the manipulation of predictors; receiver-contextualized intervention; receiver-contextualized explanation and transparent purposes; privacy protection and data subject consent; situational fairness; and human-friendly semanticization. Each factor relates to at least one of the five ethical principles of AI—beneficence, nonmaleficence, justice, autonomy and explicability—identified in the comparative analysis of (Floridi and Cows, 2019). This coherence is crucial: AI that supports Sustainable Development Goals cannot be inconsistent with ethical frameworks guiding the design and evaluation of any kind of AI.

Conclusion

This article was written to put to rest the apocalyptic narrative around the potential development of AI while avoiding the other extreme of presenting AI as the solution to all the world's problems. Both of these narratives have been recycled in news headlines or championed by profit-seeking or ethics dumping organisations and do more harm than good.

To summarise this article as succinctly as possible: No, AI is not a solution in and of itself – beware of tech solutionism that serves only those that develop these systems. Yes, AI systems can be used to significantly improve some of our efforts to solve global issues that are increasingly interconnected and complex. Yes, AI systems can and do introduce a number of ethical issues that vary in severity depending on the context of their use. No, this does not mean we should freeze all types of AI research and development. The UN SDGs provide a robust framework to all projects seeking to advance social good by leveraging AI. The framework enjoys a quasi universal status due to the support it has received from 193 countries, and serves as a rallying force for AI-focused projects to collaborate and maximise their positive impact.

Bibliography

AI HLEG (2019) *A Definition of AI: Main Capabilities and Disciplines*. Available at:

https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=56341.

Amodei, D. and Hernandez, D. (2018) *AI and Compute, OpenAI*. Available at: <https://openai.com/blog/ai-and-compute/> (Accessed: 31 August 2020).

Ananny, M. and Crawford, K. (2018) 'Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability', *New Media & Society*, 20(3), pp. 973–989. doi: 10.1177/1461444816676645.

Benjamin, R. (2019) *Race after technology: abolitionist tools for the new Jim code*. Medford, MA: Polity.

Birhane, A. et al. (2021) 'The Values Encoded in Machine Learning Research', *arXiv:2106.15590 [cs]*. Available at: <http://arxiv.org/abs/2106.15590> (Accessed: 8 July 2021).

Cho, R. (2020) 'Why Climate Change is an Environmental Justice Issue', *State of the Planet*, 22 September. Available at: <https://blogs.ei.columbia.edu/2020/09/22/climate-change-environmental-justice/> (Accessed: 14 April 2021).

Cowls, J. et al. (2021a) 'A definition, benchmark and database of AI for social good initiatives', *Nature Machine Intelligence*, 3(2), pp. 111–115. doi: 10.1038/s42256-021-00296-0.

Cowls, J. et al. (2021b) 'The AI Gambit — Leveraging Artificial Intelligence to Combat Climate Change: Opportunities, Challenges, and Recommendations', *SSRN Electronic Journal*. doi: 10.2139/ssrn.3804983.

Eubanks, V. (2017) *Automating inequality: how high-tech tools profile, police, and punish the poor*. First Edition. New York, NY: St. Martin's Press.

Fink, K. (2018) 'Opening the government's black boxes: freedom of information and algorithmic accountability', *Information, Communication & Society*, 21(10), pp. 1453–1471. doi: 10.1080/1369118X.2017.1330418.

Floridi, L. (2019) 'Translating Principles into Practices of Digital Ethics: Five Risks of Being Unethical', *Philosophy & Technology*, 32(2), pp. 185–193. doi: 10.1007/s13347-019-00354-x.

Floridi, L. et al. (2020) 'How to Design AI for Social Good: Seven Essential Factors', *Science and Engineering Ethics*, 26(3), pp. 1771–1796. doi: 10.1007/s11948-020-00213-5.

Floridi, L. and Cowls, J. (2019) 'A Unified Framework of Five Principles for AI in Society', *Harvard Data Science Review*. doi: 10.1162/99608f92.8cd550d1.

Friend, T. (2018) *How Frightened Should We Be of A.I.? | The New Yorker*. Available at: <https://www.newyorker.com/magazine/2018/05/14/how-frightened-should-we-be-of-ai> (Accessed: 8 July 2021).

McKinsey (2018) *Marrying artificial intelligence and the sustainable development goals: The global economic impact of AI | McKinsey*. Available at: <https://www.mckinsey.com/mgi/overview/in-the-news/marrying-artificial-intelligence-and-the-sustainable> (Accessed: 8 July 2021).

Morozov, E. (2013) *To Save Everything, Click Here: The Folly of Technological Solutionism*.

Obermeyer, Z. et al. (2019) 'Dissecting racial bias in an algorithm used to manage the health of populations', *Science*, 366(6464), pp. 447–453. doi: 10.1126/science.aax2342.

Perrault, R. et al. (2019) *Artificial Intelligence Index Report 2019*.

Ram, A. (2019) 'Europe's AI start-ups often do not use AI, study finds', 5 March. Available at: <https://www.ft.com/content/21b19010-3e9f-11e9-b896-fe36ec32aece>.

Roberts, H. *et al.* (2021) *Achieving a 'Good AI Society': comparing the aims and progress of the EU and the US*. SSRN Scholarly Paper ID 3851523. Rochester, NY: Social Science Research Network. doi: 10.2139/ssrn.3851523.

Taddeo, M. *et al.* (2021) 'Artificial intelligence and the climate emergency: Opportunities, challenges, and recommendations', *One Earth*, 4(6), pp. 776–779. doi: 10.1016/j.oneear.2021.05.018.

Taylor, L., Floridi, L. and van der Sloot, B. (eds) (Forthcoming) *Group Privacy: New Challenges of Data Technologies*. Hildenberg: Philosophical Studies, Book Series, Springer.

Tsamados, A. *et al.* (2021) 'The ethics of algorithms: key problems and solutions', *AI & SOCIETY*. doi: 10.1007/s00146-021-01154-8.

Vinuesa, R. *et al.* (2020) 'The role of artificial intelligence in achieving the Sustainable Development Goals', *Nature Communications*, 11(1), p. 233. doi: 10.1038/s41467-019-14108-y.