AI for Whom?
Shedding Critical Light on AI for Social Good

Nyalleng Moorosi*
DAIR Institute
Lesotho

Raesetje Sefala
DAIR Institute
South Africa

Alexandra Sasha Luccioni
Hugging Face
Canada

Abstract

In recent years, AI for Social Good (AI4SG) projects have grown in scope and popularity, covering a variety of topics from climate change to education and being the subject of numerous workshops and conferences at a global scale. In the current article, we reflect upon AI4SG, its definition and its current limitations. We propose ways to address these limitations, from connecting with relevant disciplines to a better consideration of the constraints and context of project deployment. We conclude with a proposal to refocus the field of AI4SG around the concept of sustainability from a variety of angles, arguing that this will help the field evolve while taking its own impacts into account.

1 Introduction

In recent years, a large number of Artificial Intelligence (AI) projects have been deployed. In these, the aim is to utilize cutting-edge technologies to alleviate critical issues such as poverty, hunger, crime, and climate change under the broader 'AI for Social Good' (AI4SG) umbrella [Chui et al., 2018, Taddeo and Floridi, 2018, Floridi et al., 2021, Aula and Bowles, 2023]. For AI researchers, these projects not only promise considerable impact in tackling difficult societal issues, but they also provide exciting fundamental research opportunities. Thus it is not surprising that entire organizations have been created to scope out and deploy AI4SG projects, both independently and as part of non-profit and large for-profit organizations such as Microsoft and Google, while workshops and tracks dedicated to AI4SG have appeared in conferences such as AAAI, ICLR and NeurIPS.

In this paper, we take a step back to take a critical look at AI4SG projects, from the way in which they are defined, including their roots in Information and Communication Technology for Development (ICT4D) programs, to the inherent issues that they have when it comes to their scope, deployment and maintenance. We propose ways to address these shortcomings, honing in on specific case studies in AI4SG that have overcome them. Finally, we conclude with a discussion around how AI4SG can be conceptualized to bring sustainability to its heart, and how that will help bridge the gap between well-intentioned AI4SG technologies and lasting positive impacts.

1.1 Defining AI for Social Good

AI4SG is generally referred to as an umbrella term for a sub-discipline of computing whose aim is to utilise AI to alleviate some of societies greatest challenges such as health [Bullock et al., 2020],

* All authors contributed equally to this research.

Computational Sustainability Workshop at NeurIPS 2023.

humanitarian crises [Beduschi, 2022], food security [Eli-Chukwu, 2019] and climate change [Rolnick et al., 2022]. Several formal definitions have been proposed over the years, many anchoring upon the Sustainable Development Goals (SDGs), which were defined by the United Nations in 2015 as a "blueprint" to be achieved by 2030 by all UN member states in order to improve global peace and prosperity [2]. Despite these unifying proposals, AI4SG initiatives remain pursued in a distributed manner under a variety of banners ranging from the generic (e.g. 'AI for Humanity') to the specific (e.g. 'AI for Development', 'AI for Global Health'), with a lack of an agreed upon definition of what exactly constitutes AI4SG work and how to measure its success.

While AI4SG research is fairly recent (mostly from 2010 onward), the ideas of tech for good are decades old [Heeks, 2009, Rashid and Rahman, 2010]. In the mid 1980s, many of the research that now falls under AI4SG were part of ICT4D [Aula and Bowles, 2023, Tomašev et al., 2020]. Like AI4SG, the boundaries of ICT4D were open and some of the issues we discuss in the following sections were already being identified. A set of four principles, published in 1990, were drafted to guide ICT4D research [Bhatnagar and Bjørn-Andersen, 1990, Walsham, 2017, Heeks, 2008]; these were: context awareness, participatory and cooperative design, indigenous development, and recognizing that information technology is only one element of the necessary efforts. We refer back to this literature because of the similarity between the current problems that plague AI4SG and these mentioned in ICT4D literature.

2 What's Wrong with AI for Social Good?

While we do not doubt the good intentions of AI4SG projects, we find that many of the ways in which they are implemented and deployed suffer from several limitations, which we dive into in this section. Some of these points are similar to those we mention in Section 1.1, and for those we update and expand them for our current computing environment. Additionally, we look at some new issues which are endemic to AI technologies.

Epistemological issues: Given the inherently subjective definition of '*good*' (which can benefit one community while harming another) [Schroeder, 2021], what exactly is included in AI4SG projects is also unclear. For instance, whether topics such as language preservation and predictive policing fall under the AI4SG umbrella, since they do not pursue any of the SDGs specifically, but can help empower communities retain autonomy [Hao et al., 2023, Whaanga et al., 2023] as well as counter mechanisms of oppression [West et al., 2019, Crawford and Calo, 2016]. Even when anchored in the seemingly universal UN SDGs, which provide specific goals, AI4SG projects espouse a specific set of values and priorities that have already been critiqued from a Global South perspective due to their homogeneous representation and understanding of multiple issues ranging from wealth to sustainability [Waldmueller, 2015, Durokifa and Ijeoma, 2018, Ziai, 2016].

Ignoring research from other disciplines: While AI4SG projects propose AI as a *solution* for major societal problems like global health or climate change, they often ignore the fact that social problems are inherently economical, historical, political, cultural and therefore require many different types of interventions. While it may seem that many of society's problems are ripe to be addressed with a seamless computational solution (i.e. a techno-solutionist approach), reviewing related work in other fields such as linguistics, sociology, development studies, geographic information systems or political science (e.g. [Sartori and Theodorou, 2022, Kusters et al., 2020]) is important to both have a better understanding of the issue at hand and what factors to consider when proposing an AI-based "solution" to a societal problem. It is undoubtedly more challenging to both find and analyze research done in other disciplines, both because of the unfamiliar publishing venues and the divergent vocabulary Baum [2021]. However, it is paramount to examine these works to ensure that any technical solution proposed to a societal problem fits with the way in which the problem has been studied and framed by domain experts.

Emphasizing deployment over maintenance: AI research in general tends to place values such as novelty and performance over ones like efficiency, interpretability and representativity [Birhane et al., 2022] and AI4SG is no exception to this. While the former may be more logical in contexts such as scientific publications – which are evaluated based on criteria such as novelty – the latter are much more important for AI systems deployed in the real world. Given that AI4SG publication venues are often co-located with AI conferences such as AAAI, ICML and NeurIPS, the values of these venues

[2] `https://sdgs.un.org/goals`

influence both the scope and formulation of AI4SG (to maximize their chances of acceptance). This inadvertently penalizes work that uses existing approaches in new contexts, or that which employs more efficient methods over those that are 'state-of-the-art' which are often untested on the ground.

Imperfect technologies result in imperfect solutions: AI for good solutions often get built with the intention to get deployed into the real world. However, the question of whether the AI technology behind these solutions actually works as intended is an important one that is often not rigorously explored. As any technology, AI exists within a broader societal and environmental context. Taking this context into account is crucial for ensuring that AI4SG projects have a chance of succeeding in reaching their intended goals. Recent work by Raji et al. [2022] extensively covers real world examples of AI systems being deployed into the world when the technology itself does not function as intended. This results in harm towards the communities it was intended to help (also see Obermeyer et al. [2019], Szalavitz [2021], Strickland [2019], Wong et al. [2021], Richardson et al. [2019], Freeman et al. [2021]).

3 Making AI for Social Good Better

The critiques that we present to the framing and execution of AI4SG project can be addressed with a series of complementary actions, as illustrated by the case studies that we mention alongside each of the proposals below:

Mindful problem scoping and data collection: At the beginning of any AI4SG project, it is primordial to reflect upon the way in which it is defined – while the overarching goal of a given endeavor may be to reduce poverty or solve climate change, the scope of any individual project will necessarily be much narrower. Addressing specific societal problems such as detecting illegal logging or identifying students at risk of attrition, instead of general concepts such as "social good", can help make project outcomes more tangible [Baum, 2021]. Once a specific problem has been identified, it is then possible to establish whether relevant data is available. For instance, in cases where satellite imagery is used in humanitarian contexts, instead of using GDP, which overlooks aspects such as sustainability or well-being, [Ivković, 2016, Osberg and Sharpe, 2005], it is possible to use a more diverse set of indicators such as the presence of bathroom facilities and cooking stoves in residences, as indicators of wealth (see Bansal et al. [2020], Salas et al. [2021]).

Giving power to the people: The way in which AI4SG is currently practiced is more akin to consulting than co-creation, which may address a specific issue at a given point in time but is unsustainable in the long term because it fails to build local capacity [Nightingale, 2012]. However, the inclusion of communities ranging from non-governmental organizations (NGOs) to charities is important for the success of AI4SG projects because for AI4SG solutions to truly be effective, they should include and empower local stakeholders who are the people for whom the solution should work and are also the people who will maintain the system when the project ends and who have real-world expertise regarding the problem at hand and the solution it requires (see Tomašev et al. [2020], Vinuesa et al. [2020], Mohamed et al. [2020]). A prime example of this is the Masakhane initiative [Orife et al., 2020] – which provides compute, datasets and ongoing knowledge support for communities who want to build language preservation tools for African languages. The workshops and projects they organize empower local experts such as linguists, dataset curators and NLP researchers to explore using AI as a tool for language preservation without losing control or governance over the results of their research both in terms of models and datasets.

Additionally, new research in machine learning indicates there are many levers that are available to AI researchers which result in vastly different products. By 2018, AI researchers noted the reproducibility crisis in ML Hutson [2018]. This means that several published ML models can not be independently verified by other researchers Belz et al. [2021]. More importantly, there are results that show that the same models can output vastly different outputs for the same data point, in deployment than in training [Chen et al., 2020]. This phenomena is even more visible in large models [D'Amour et al., 2020]. For sensitive applications, this is not sustainable.

Connecting to relevant disciplines: As stated in Section 2, a major shortcoming of AI4SG projects is that they predominantly focus on relevant AI-specific literature, which results in the perpetuation of simplifications and outright inaccuracies. For example, while in theory reducing the risk of violent crime is a worthwhile endeavor, in practice using a technology such as AI to predict the location of the next carjacking in a particular region is an intractable problem [Sankin and Mattu, 2023]. Doing effective predictive work like this often requires usable, complete, accurate, unbiased and clean historical data which is almost impossible to get [Ferguson, 2016]. However, these requirements have been long studied in other disciplines like law and criminology and there is evidence to show how a lot of the crimes go unreported and how police-reported data in some places has been shown to be inaccurate, misleading and sometimes manipulated [Morganteen, 2013, Rudovsky, 2001, Rashbaum, 2010]. Working in silos also means not paying attention to the story the data comes with and that contributes to the kind of impact being delivered to communities.

Considering deployment constraints: More often than not, state of the art, deep learning-based solutions are too complex, too expensive and too technologically demanding to be deployed *in situ* for various reasons from lack of access to sufficient computational resources [Kshirsagar et al., 2021] to the absence of reliable electricity grid [Beduschi, 2022]. It is therefore paramount to take into account the context and constraints of deployment when proposing AI-based solutions. This can include first trying other, more simpler, approaches such as random forests or Bayesian networks, as well as approaches that can run offline or on-edge on devices with limited memory. For example, CNN-based bio-acoustic monitoring solutions are used to detect illegal logging deforestation, which run on disused mobile phones powered by solar panels, enabling them to cover wide regions of the Amazon rain-forest [Liu et al., 2019].

In many of the top AI conferences, the AI4SG tracks and workshops often only have projects that are at the level of a minimum-viable product. We wish to make a call to the community to also explicitly seek projects that highlight how to run sustainable AI4SG projects post-deployment. These conferences may very well exist in other disciplines but due to how siloed academic research can be, it would make a meaningful contribution to include these discussions in our own proceedings given that we add a unique AI aspect to software engineering deployments. Lastly, we propose that in addition to authors submitting works on the technical details of a project, they should also include a statement on the sustainability of their project - stating measures they have put or will put in place to make sure that their project has impact and longevity.

4 Discussion and Conclusion

As discussed in previous sections, AI4SG suffers from several limitations, from an unclear definition of what constitutes 'good' to a siloed practice of AI4SG projects that is often isolated from both relevant literature from other domains as well as from the contexts and communities who are meant to use this technology. To overcome these limitations, we propose putting *sustainability* at the core of AI4SG work instead. As many definitions of sustainability often rely on several pre-defined 'pillars' (e.g. human, social, economic and environmental) (see Sverdrup and Svensson [2002], Vos [2007], Basiago [1995], *inter alia*), similarly AI4SG projects should consider the impacts of different aspects of their projects on different parts of society and the environment, depending on the project at hand and its ramifications. These can range from:

Environmental impacts: AI technologies require energy and engender greenhouse gas emissions, which should be quantified and communicated [Schmidt et al., 2021], and efforts should be made to reduce them, for instance by choosing low-carbon compute [Luccioni and Hernandez-Garcia, 2023] or strategies that allow them to run jobs at times when carbon intensity is low [Dodge et al., 2022].

Human impacts: By this, we refer to the importance of carrying out capacity building and ensuring the sustainability of AI4SG projects in the long run, since maintenance needs to be carried out over time after the initial research is finished. The importance of capacity-building has historically been under-explored in AI in general and AI4SG in particular, but recent work has highlighted its importance to ensure the perennity of AI projects (e.g. [Sey and Mudongo, 2021]).

Economic impacts: In cases where proposed projects intend to partially or fully automate labor currently carried out by human workers, the economic impacts of this automation should be considered Klinova and Korinek [2021], Solaiman et al. [2023]. When projects rely upon grants or funding, the perennity of external sources of funding also becomes an important factor for consideration. While there are different structures to ensure the funding of AI4SG projects, ranging from for-profit elements of an otherwise not-for-profit enterprise to adoption of projects by NGOs or corporations, considering these possibilities from the start can help maximize economic sustainability.

Social impacts: While 'social' is a very broad term, in the context of AI4SG projects, we may choose to focus on aspects relating to the power dynamics between the different actors involved. With AI technologies increasingly being influenced by industry players (see Abdalla and Abdalla [2021], Abdalla et al. [2023] for an overview), it is important to consider these when making partnerships across for-profit and non-profit institutions. Other aspects such as licensing and governance over the code and data, rights to publish and communicate results, and accountability are all important to consider, and there is no one-size-fits-all solution. However, being aware of the power differentials at play and the different options that are available is important to ensure sustainability.

Finally, none of the factors described above exists in isolation, and trade-offs must be made between several of them. For instance, if we take the specific example of creating an AI chatbot for farmers that will allow them to ask questions about weather and yield, this can present a choice between investing money into creating custom solution in-house versus paying for a generic solution such as a commercial API, which can be cheaper. This also requires the consideration of the social impacts of building capacity locally versus outsourcing the solution to a large corporation, further aggravating the existing power dynamics between multi-national technology companies versus NGOs and local enterprises. There is also a multitude of more context-specific questions, such as how to evaluate whether a given technology can take into account the specific details of its deployment in practice (such as regional weather patterns), as well as issues of accountability and responsibility if any of the proposed technologies provide erroneous advice and cost a farmer their crop.

With AI maturing as both a field of academic study as well as for-profit deployment, this produces both new opportunities for AI4SG work as well as new challenges such as the ones described above. In the present article, we have endeavored to describe the current state of AI4SG, its limitations and future directions for improvement, and we remain convinced that putting sustainability at the center of AI4SG projects will contribute to improving the sustainability of the field as a whole.

References

Mohamed Abdalla and Moustafa Abdalla. The grey hoodie project: Big tobacco, big tech, and the threat on academic integrity. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 287–297, 2021.

Mohamed Abdalla, Jan Philip Wahle, Terry Ruas, Aurélie Névéol, Fanny Ducel, Saif M Mohammad, and Karën Fort. The elephant in the room: Analyzing the presence of big tech in natural language processing research. *arXiv preprint arXiv:2305.02797*, 2023.

Ville Aula and James Bowles. Stepping back from data and ai for good–current trends and ways forward. *Big Data & Society*, 10(1):20539517231173901, 2023.

Chahat Bansal, Arpit Jain, Phaneesh Barwaria, Anuj Choudhary, Anupam Singh, Ayush Gupta, and Aaditeshwar Seth. Temporal prediction of socio-economic indicators using satellite imagery. In *Proceedings of the 7th ACM IKDD CoDS and 25th COMAD*, pages 73–81. 2020.

Andrew D Basiago. Methods of defining 'sustainability'. *Sustainable development*, 3(3):109–119, 1995.

Seth D Baum. Artificial interdisciplinarity: Artificial intelligence for research on complex societal problems. *Philosophy & Technology*, 34(Suppl 1):45–63, 2021.

Ana Beduschi. Harnessing the potential of artificial intelligence for humanitarian action: Opportunities and risks. *International Review of the Red Cross*, 104(919):1149–1169, 2022.

Anya Belz, Shubham Agarwal, Anastasia Shimorina, and Ehud Reiter. A systematic review of reproducibility research in natural language processing. *arXiv preprint arXiv:2103.07929*, 2021.

Subhash C Bhatnagar and Niels Bjørn-Andersen. Information technology in developing countries. *Proceedings of the Ifip Tc9/tc8 Working Conference on the Impact of Information Systems on Developing Countries,New Delhi, India, 24-26 November 1988*, 1990.

Abeba Birhane, Pratyusha Kalluri, Dallas Card, William Agnew, Ravit Dotan, and Michelle Bao. The values encoded in machine learning research, 2022.

Joseph Bullock, Alexandra Luccioni, Katherine Hoffman Pham, Cynthia Sin Nga Lam, and Miguel Luengo-Oroz. Mapping the landscape of artificial intelligence applications against covid-19. *Journal of Artificial Intelligence Research*, 69:807–845, 2020.

Zhe Chen, Yuyan Wang, Dong Lin, Derek Zhiyuan Cheng, Lichan Hong, Ed H. Chi, and Claire Cui. Beyond point estimate: Inferring ensemble prediction variation from neuron activation strength in recommender systems, 2020.

Michael Chui, Martin Harryson, Silicon Valley, James Manyika, and Roger Roberts. Notes from the AI frontier applying ai for social good. 2018.

Kate Crawford and Ryan Calo. There is a blind spot in AI research. *Nature*, 538(7625):311–313, 2016.

Alexander D'Amour, Katherine Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel, Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D. Hoffman, Farhad Hormozdiari, Neil Houlsby, Shaobo Hou, Ghassen Jerfel, Alan Karthikesalingam, Mario Lucic, Yian Ma, Cory McLean, Diana Mincu, Akinori Mitani, Andrea Montanari, Zachary Nado, Vivek Natarajan, Christopher Nielson, Thomas F. Osborne, Rajiv Raman, Kim Ramasamy, Rory Sayres, Jessica Schrouff, Martin Seneviratne, Shannon Sequeira, Harini Suresh, Victor Veitch, Max Vladymyrov, Xuezhi Wang, Kellie Webster, Steve Yadlowsky, Taedong Yun, Xiaohua Zhai, and D. Sculley. Underspecification presents challenges for credibility in modern machine learning, 2020.

Jesse Dodge, Taylor Prewitt, Remi Tachet des Combes, Erika Odmark, Roy Schwartz, Emma Strubell, Alexandra Sasha Luccioni, Noah A Smith, Nicole DeCario, and Will Buchanan. Measuring the carbon intensity of ai in cloud instances. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 1877–1894, 2022.

Anuoluwapo Abosede Durokifa and Edwin Chikata Ijeoma. Neo-colonialism and millennium development goals (mdgs) in africa: A blend of an old wine in a new bottle. *African Journal of Science, Technology, Innovation and Development*, 10(3):355–366, 2018.

Ngozi Clara Eli-Chukwu. Applications of artificial intelligence in agriculture: A review. *Engineering, Technology & Applied Science Research*, 9(4), 2019.

Andrew Guthrie Ferguson. Policing predictive policing. *Wash. UL Rev.*, 94:1109, 2016.

Luciano Floridi, Josh Cowls, Monica Beltrametti, Raja Chatila, Patrice Chazerand, Virginia Dignum, Christoph Luetge, Robert Madelin, Ugo Pagallo, Francesca Rossi, et al. An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Ethics, governance, and policies in artificial intelligence*, pages 19–39, 2021.

Karoline Freeman, Julia Geppert, Chris Stinton, Daniel Todkill, Samantha Johnson, Aileen Clarke, and Sian Taylor-Phillips. Use of artificial intelligence for image analysis in breast cancer screening programmes: systematic review of test accuracy. *bmj*, 374, 2021.

Karen Hao, Heidi Swart, Andrea Paola Hernández, and Nadine Freischlad. AI Colonialism. 2 2023. URL `https://www.technologyreview.com/supertopic/ai-colonialism-supertopic`.

Richard Heeks. Ict4d 2.0: The next phase of applying ict for international development. *Computer*, 41(6):26–33, 2008.

Richard Heeks. The ict4d 2.0 manifesto: Where next for icts and international development? *Development Informatics Working Paper*, (42), 2009.

Matthew Hutson. Artificial intelligence faces reproducibility crisis, 2018.

Anita Frajman Ivković. Limitations of the gdp as a measure of progress and well-being. *Ekonomski vjesnik/Econviews-Review of Contemporary Business, Entrepreneurship and Economic Issues*, 29(1):257–272, 2016.

Katya Klinova and Anton Korinek. AI and shared prosperity. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 645–651, 2021.

Meghana Kshirsagar, Caleb Robinson, Siyu Yang, Shahrzad Gholami, Ivan Klyuzhin, Sumit Mukherjee, Md Nasir, Anthony Ortiz, Felipe Oviedo, Darren Tanner, et al. Becoming good at AI for good. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 664–673, 2021.

Remy Kusters, Dusan Misevic, Hugues Berry, Antoine Cully, Yann Le Cunff, Loic Dandoy, Natalia Díaz-Rodríguez, Marion Ficher, Jonathan Grizou, Alice Othmani, et al. Interdisciplinary research in artificial intelligence: challenges and opportunities. *Frontiers in big data*, 3:577974, 2020.

Yuan Liu, Zhongwei Cheng, Jie Liu, Bourhan Yassin, Zhe Nan, and Jiebo Luo. Ai for earth: rainforest conservation by acoustic surveillance. *arXiv preprint arXiv:1908.07517*, 2019.

Alexandra Sasha Luccioni and Alex Hernandez-Garcia. Counting carbon: A survey of factors influencing the emissions of machine learning. *arXiv preprint arXiv:2302.08476*, 2023.

Shakir Mohamed, Marie-Therese Png, and William Isaac. Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33:659–684, 2020.

Jeff Morganteen. What the compstat audit reveals about the nypd. *The New York World*, 3, 2013.

Katherine Nightingale. Building the future of humanitarian aid: Local capacity and partnerships in emergency assistance. Technical report, Christian Aid, 2012.

Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464):447–453, 2019.

Iroro Orife, Julia Kreutzer, Blessing Sibanda, Daniel Whitenack, Kathleen Siminyu, Laura Martinus, Jamiil Toure Ali, Jade Abbott, Vukosi Marivate, Salomon Kabongo, et al. Masakhane–machine translation for africa. *arXiv preprint arXiv:2003.11529*, 2020.

Lars Osberg and Andrew Sharpe. How should we measure the "economic" aspects of well-being? *Review of Income and Wealth*, 51(2):311–336, 2005.

Inioluwa Deborah Raji, I Elizabeth Kumar, Aaron Horowitz, and Andrew Selbst. The fallacy of ai functionality. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 959–972, 2022.

William K Rashbaum. Retired officers raise questions on crime data. *New York Times*, 6, 2010.

Nahid Rashid and Shams Rahman. An investigation into critical determinants of e-government implementation in the context of a developing nation. In *Proceedings of the First International Conference on Electronic Government and the Information Systems Perspective*, EGOVIS'10, page 9–21, Berlin, Heidelberg, 2010. Springer-Verlag. ISBN 364215171X.

Rashida Richardson, Jason M Schultz, and Vincent M Southerland. Litigating algorithms: 2019 us report. *AI Now Institute*, 2019.

David Rolnick, Priya L Donti, Lynn H Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, et al. Tackling climate change with machine learning. *ACM Computing Surveys (CSUR)*, 55(2):1–96, 2022.

David Rudovsky. Law enforcement by stereotypes and serendipity: Racial profiling and stops and searches without cause. *U. Pa. J. Const. L.*, 3:296, 2001.

Joaquín Salas, Pablo Vera, Marivel Zea-Ortiz, Elio-Atenogenes Villaseñor, Dagoberto Pulido, and Alejandra Figueroa. Fine-grained large-scale vulnerable communities mapping via satellite imagery and population census using deep learning. *Remote Sensing*, 13(18):3603, 2021.

Aaron Sankin and Surya Mattu. Predictive policing software terrible at predicting crimes, 2023. URL `https://www.wired.com/story/plainfield-geolitica-crime-predictions/`.

Laura Sartori and Andreas Theodorou. A sociotechnical perspective for the future of ai: narratives, inequalities, and human control. *Ethics and Information Technology*, 24(1):4, 2022.

Victor Schmidt, Kamal Goyal, Aditya Joshi, Boris Feld, Liam Conell, Nikolas Laskaris, Doug Blank, Jonathan Wilson, Sorelle Friedler, and Sasha Luccioni. Codecarbon: estimate and track carbon emissions from machine learning computing. *Cited on*, page 20, 2021.

Mark Schroeder. Value Theory. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2021 edition, 2021.

Araba Sey and Oarabile Mudongo. Case studies on ai skills capacity building and ai in workforce development in africa. 2021.

Irene Solaiman, Zeerak Talat, William Agnew, Lama Ahmad, Dylan Baker, Su Lin Blodgett, Hal Daumé III, Jesse Dodge, Ellie Evans, Sara Hooker, et al. Evaluating the social impact of generative ai systems in systems and society. *arXiv preprint arXiv:2306.05949*, 2023.

Eliza Strickland. Ibm watson, heal thyself: How ibm overpromised and underdelivered on ai health care. *IEEE Spectrum*, 56(4):24–31, 2019.

Harald Sverdrup and Mats GE Svensson. Defining sustainability. In *Developing principles and models for sustainable forestry in Sweden*, pages 21–32. Springer, 2002.

Maia Szalavitz. The pain was unbearable. so why did doctors turn her away. *Wired. August*, 11, 2021.

Mariarosaria Taddeo and Luciano Floridi. How AI can be a force for good. *Science*, 361(6404):751–752, 2018.

Nenad Tomašev, Julien Cornebise, Frank Hutter, Shakir Mohamed, Angela Picciariello, Bec Connelly, Danielle CM Belgrave, Daphne Ezer, Fanny Cachat van der Haert, Frank Mugisha, et al. Ai for social good: unlocking the opportunity for positive impact. *Nature Communications*, 11(1):2468, 2020.

Ricardo Vinuesa, Hossein Azizpour, Iolanda Leite, Madeline Balaam, Virginia Dignum, Sami Domisch, Anna Felländer, Simone Daniela Langhans, Max Tegmark, and Francesco Fuso Nerini. The role of artificial intelligence in achieving the sustainable development goals. *Nature communications*, 11(1):1–10, 2020.

Robert O Vos. Defining sustainability: a conceptual orientation. *Journal of Chemical Technology & Biotechnology: International Research in Process, Environmental & Clean Technology*, 82(4):334–339, 2007.

Johannes M Waldmueller. Agriculture, knowledge and the 'colonial matrix of power': approaching sustainabilities from the global south. *Journal of Global Ethics*, 11(3):294–302, 2015.

Geoff Walsham. ICT4D research: reflections on history and future agenda. *Information Technology for Development*, 23(1):18–41, 2017. doi: 10.1080/02681102.2016.1246406. URL `https://doi.org/10.1080/02681102.2016.1246406`.

Sarah Myers West, Meredith Whittaker, and Kate Crawford. Discriminating systems. *AI Now*, pages 1–33, 2019.

Hēmi Whaanga, Te Taka Keegan, and Petera Hudson. Computing Technologies for Resilience, Sustainability and Resistance. 2 2023. doi: 10.36227/techrxiv.22146500.v1. URL `https://www.techrxiv.org/articles/preprint/Computing_Technologies_for_Resilience_Sustainability_and_Resistance/22146500`.

Andrew Wong, Erkin Otles, John P Donnelly, Andrew Krumm, Jeffrey McCullough, Olivia DeTroyer-Cooley, Justin Pestrue, Marie Phillips, Judy Konye, Carleen Penoza, et al. External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. *JAMA Internal Medicine*, 181(8):1065–1070, 2021.

Aram Ziai. *Development discourse and global history: From colonialism to the sustainable development goals*. Taylor & Francis, 2016.