FOUNDATION MODEL PLATFORMS AND BOTTOM-OF-THE-PYRAMID INNOVATION

Kush R. Varshney
IBM Research – Thomas J. Watson Research Center
Yorktown Heights, NY 10598, USA
krvarshn@us.ibm.com

ABSTRACT

We analyze whether artificial intelligence (AI) foundation model platforms meet the twelve principles of bottom-of-the-pyramid innovation. Perhaps surprisingly, we find that they do satisfy several of the principles. On many of the ones they do not, active AI research is underway. Some platform functionality beyond the actual foundation model and education in the language of low-resourced users is still missing. This analysis is supported by a case study of ChatGPT to perform semantic mapping of sentences in a social good context. A call to action for grantmaking institutions is provided.

1 INTRODUCTION

The two components of the title of this paper may seem incongruous or even oxymoronic. Foundation models for artificial intelligence (AI), also known as large language models when applied to natural language data modalities, are behemoths that require herculean efforts to train and run (Bender et al., 2021). For example, it required an army of at least 31 people (Brown et al., 2020) and cost several millions of dollars to train the GPT-3 model (Li, 2020). It is estimated that training the BLOOM model emitted 50.5 tonnes of carbon dioxide equivalent and keeping its API deployed and responsive is very costly as well (Luccioni et al., 2022).

In contrast, bottom-of-the-pyramid innovation is the idea of creating products of world-class quality for the world's socioeconomically poorest 2.7 billion people who live on less than 2.50 USD per day at one-fiftieth the cost (or less) of those for developed markets (Prahalad, 2012). Extant examples include biomass stoves, shampoo, solar lanterns, electrocardiogram machines, cars, and cataract eye surgery. The key to such innovation is not forcing products or processes from developed contexts onto developing ones, but to start a fresh design that begins with immersion into the lives of the consumers to obtain insights and appropriate constraints. The argument is that constraints are a necessary ingredient for breakthrough innovation and that the solutions obtained may be even better than top-of-the-pyramid solutions.

In the last decade, there has been a movement of using AI for social impact — to address poverty, hunger, ill health, poor education, and inequities of various kinds. The primary paradigm has been volunteer data scientists partnering with social change organizations (non-profits, social enterprises, government agencies, etc.) on short-term, one-off projects that may or may not be truly impactful. Varshney & Mojsilović (2019) argue that to make true impact, what is needed is a reusable platform approach rather than bespoke projects. A platform is a group of technologies used as a base upon which applications and processes are developed. Platforms, instead of custom-tailored solutions, reduce the time and effort in creating solutions, avoid the limited asset reuse that pervades AI for social good, and reduce the work to integrate, deploy, monitor, and maintain solutions.

Since that paper predates the explosion of foundation models (and the coining of the term by Bommasani et al. (2021)), the authors point to single AI algorithmic formulations that can be repurposed for multiple tasks, cf. Shi et al. (2020), and to AI toolkits and libraries that contain several closely related algorithms applicable to different tasks, cf. Bellamy et al. (2019), as the starting point for platforms. Today, however, we would argue that foundation models are the ultimate platform.

Non-profit organizations and social enterprises (in contrast to large social media, financial services, information technology, and retail companies) are low-resourced just like the individuals they serve. They are bottom-of-the-pyramid *organizations*. In this paper, we will investigate whether the way foundation models are being developed and deployed now leave gaps in being appropriate for low-resourced organizations, especially those working to make positive social impact in the developing world. In doing so, we will point out some potential calls to action. The discussion, although grounded at the bottom of the pyramid, will be instructive for organizations, industries, and sectors that span the entire gamut except for perhaps the three or four companies currently synonymous with foundation models.

2 Twelve Principles of Bottom-of-the-Pyramid Innovation

The so-called bottom of the pyramid is highly heterogeneous, ranging from people dwelling in urban slums to those spread out in rural villages on different continents having different cultures and ways of life. Nevertheless, creating products and ecosystems to serve these disparate populations commonly requires a focus on awareness, accessibility, affordability, and availability (Prahalad, 2012) by adhering to the following twelve principles (Prahalad, 2006):

1. focus on (quantum jumps in) price performance;
2. hybrid solutions, blending old and new technology;
3. scalable and transportable operations across countries, cultures and languages;
4. reduced resource intensity: eco-friendly products;
5. identify appropriate functionality;
6. build logistical and manufacturing infrastructure;
7. deskill (services) work;
8. educate (semi-literate) customers in product usage;
9. products must work in hostile environments;
10. adaptable user interface to heterogeneous consumer bases;
11. distribution methods designed to reach both highly dispersed rural markets and highly dense urban markets; and
12. focus on broad architecture, enabling quick and easy incorporation of new features.

3 Principles in the Context of Foundation Models

Some of these principles are more and less relevant to AI and foundation models. Let us go through them one by one and discuss.

First, price. Although some foundation models may presently be called for free or at a nominal cost for users, the truth of the matter is that they are still extremely extremely expensive, and no organization in the world outside of a select handful have the resources to run them. If the end-user maintains the ability to run inferences with the model, and the cost is amortized across many users, then perhaps the price performance is actually reasonable. Nevertheless, reducing the costs of foundation models (e.g. by somehow intelligently reducing the computation of training or finding smaller, pruned or compressed models with equivalent performance) are called upon from the perspective of bottom of the pyramid innovation, and are being pursued by AI researchers.

Second, foundation models are already beginning to combine the new technology of the transformer architecture with older technologies like symbolic AI, information retrieval and human feedback. These hybrid approaches are achieving improved performance and also reduced costs.

Third, large language models currently struggle with obtaining equivalent performance on low-resourced languages and dialects as on high-resourced languages such as general American English. Much of this discrepancy is because of much smaller available dataset size, which initiatives such as Masakhane are addressing (Orife et al., 2020). Additionally, there may be inherent differences between languages, such as in their morphology, that make some easier to model and others harder to model (Wan, 2022). Similar phenomena exist in foundation models for other modalities. Research is underway to ameliorate this lack of transportability across cultures and languages.

Fourth, as mentioned in the introduction, training and running foundation model platforms incurs tremendously high computational and thus energy and environmental costs. Moreover, one should ask where the data centers are located at which the computations are performed. They are often in rural towns that bear the brunt of environmental and social costs. Reducing environmental costs and spreading them equally will require new innovation.

The fifth principle is appropriate functionality. For social change organizations to be able to use AI and foundation models to satisfy their needs, simplicity is imperative. Some of the appropriate functionality includes: model and data catalog and management, ease of applying models to slightly different settings or desired predictions, batch processing, outputs that can be ingested by common dashboard tools, and ease of integrating with other apps. The ease of applying models to different settings or predictions is a key inherent property of foundation models, but the other kinds of appropriate functionality require additional components in a foundation model platform than the foundation model itself.

With AI being an information technology, the sixth principle, manufacturing infrastructure, is not so relevant.

Seventh, since foundation models may be prompted in natural language, foundation model platforms can deskill sophisticated data and other analysis that would normally require social change organizations to have advanced skills. Furthermore, the software-as-a-service paradigm that is becoming the prevailing form factor for foundation models pushes tasks such as deployment, monitoring, and maintenance away from end-users within social change organizations.

The eighth principle, education, is critical. Although some foundation models such as Stable Diffusion and ChatGPT have caught the public imagination, there is still a gap in the imagination of how to apply such powerful models in useful ways to solve operational problems that non-profits and social enterprises face. Teaching material, tutorials, examples, etc. in the language of the social sector are needed.

Related to the ninth principle, hostile environments require a focus on safety and mitigations of hazards and harms (Varshney & Alemzadeh, 2017; Varshney, 2022). Foundation models, including large language models, have the ability to induce large socio-technical harms (Shelby et al., 2022), including unwanted discrimination, toxicity, and more. Algorithms for fairness and other pillars of trustworthy AI, specifically tailored for foundation models are being developed (Sattigeri et al., 2022; Basu et al., 2023). Moreover, like most organizations, social change organizations do not have the ability to verify the training of a foundation model because it is computationally prohibitive (Brundage et al., 2020), unless clever methods are included from the outset (Bore et al., 2019).

The tenth principle calls for an easy user experience, likely with a no-code natural language interface that is resilient to heterogeneous registers of language. It should not require computer programming to run useful jobs. The ChatGPT interface with all interaction in natural language is a good example of such an interface.

Eleventh, on distribution. As mobile phones have spread to the furthest reaches of humanity, including among social change workers serving clients in rural villages and urban slums, they are the vehicle for distributing the interfaces to foundation models for both fine-tuning (a form of extra training on a specialized dataset) and inferencing. Uploading of new data and fine-tuning should be performed when network connectivity is available. However, waiting for connectivity at inference time may not be possible. Therefore, pushing foundation models to edge devices remains an appropriate technological need (that some researchers are pursuing).

Twelfth, the software-as-a-service or platform-as-a-service approach for interacting with foundation models already allows updated re-trained models and new innovations to quickly disseminate.

To summarize, many of the principles of bottom-of-the-pyramid innovation are inherently met by foundation models delivered as platforms-as-a-service, but others require additional work. Some of that work is underway in research laboratories (although not necessarily in the same labs developing and releasing the foundation models attracting the greatest attention), such as cost reduction, hybrid approaches, performance improvement on low-resource languages, harm mitigation, and placement

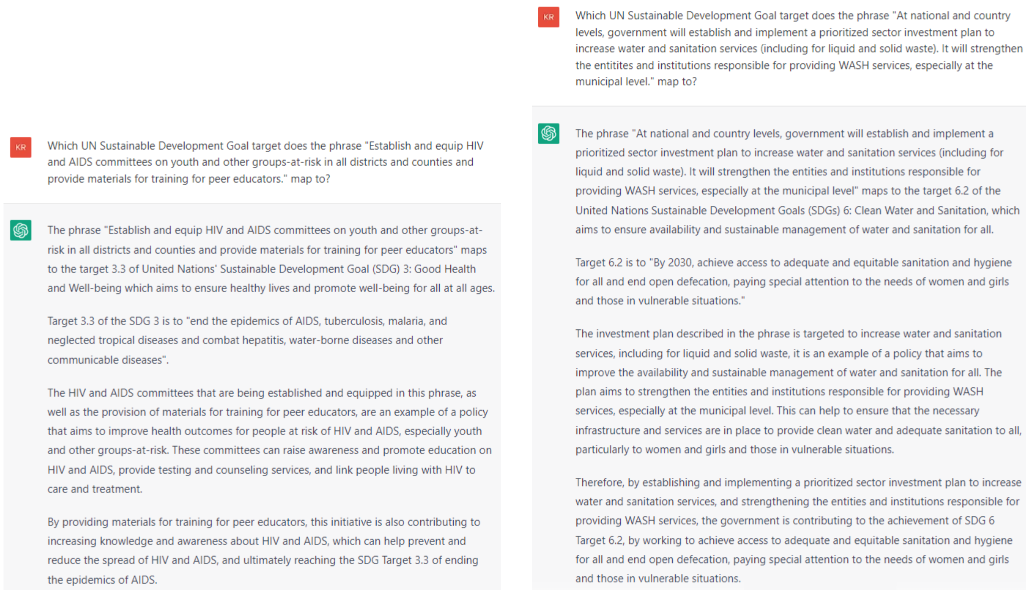

Figure 1: ChatGPT solution to the semantic mapping of sentences from the national development plan of Liberia to United Nations Sustainable Development Goal targets.

of foundation models on edge devices. Other work is not underway, such as quality model management, ingestible outputs, integrability, and education in the language of social change organizations.

4 Case Study

As a case study of how a current foundation model platform can be used in a social good application, we considered using ChatGPT to semantically map sentences in the national development plans of countries to one of 169 possible United Nations Sustainable Development Goal targets they refer to. This is a key step in a process known as the Rapid Integrated Assessment by the United Nations Development Programme (UNDP). In 2017, this precise problem was tackled by a team of six AI specialists in a bespoke 3-month-long research and development engagement with the UNDP using modern natural language processing techniques at the time (Galsurkar et al., 2018). In 2023, a low-skilled user obtained results of the same quality in seconds using the ChatGPT foundation model platform with no fine-tuning and nothing but out-of-the-box use required. Results for two example sentences from the national development plan of Liberia with ChatGPT are shown in Figure 1. They exactly match the results of Galsurkar et al. (2018).

To ensure that there was no data leakage (ChatGPT containing the Galsurkar et al. paper in its training set), we further tested the ChatGPT platform on sentences from the United Nations Human Rights Council's Universal Periodic Review in collaboration with the non-profit International Center for Advocates Against Discrimination (ICAAD). In this case, the semantic mapping task was to a plurality of the 17 higher-level United Nations Sustainable Development Goals. The results, completed by a non-AI specialist in seconds with ChatGPT out-of-the-box without requiring any computer programming, are shown in Figure 2. The mapping of these sentences to Sustainable Development Goals has never appeared before in print or online; the results exactly match human-validated results of ICAAD completed before the release of ChatGPT.

This case study is impressive, but also brought forward some shortcomings of the ChatGPT foundation model platform, which are congruent with those identified in Section 3. In terms of education, the non-AI specialist from ICAAD had heard about ChatGPT, but had not imagined that it could be used to solve the semantic mapping problem that ICAAD needs to carry out its mission. This awareness only arose in a meeting with an AI expert. Also, the interaction pattern which does not require programming, was appreciated by the ICAAD employee, but was very slow to use since every sentence along with its prompt had to be typed in. The output, although correct, was not amenable

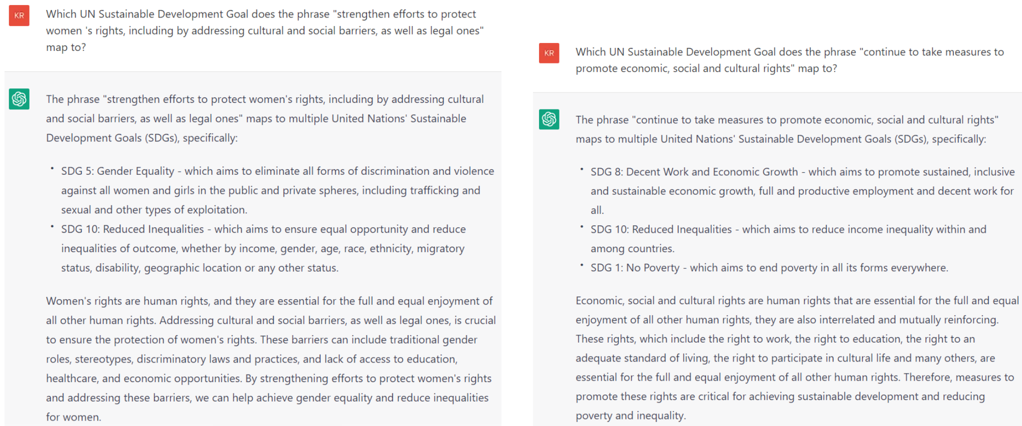

Figure 2: ChatGPT solution to semantic mapping of Sustainable Development Goals problem.

to further analysis in the tools used by ICAAD: Excel and Tableau. The results are not necessarily reproducible, which is a consideration in ICAAD's work.

5 Conclusion

It is perhaps surprising that a costly technology, AI foundation models, already abides by a large number of tenets of bottom-of-the-pyramid innovation. In most areas that it does not, research is already underway to close the gap. The biggest gaps that remain are in educating the social sector and in the ancillary platform components for model management and integration with other tools. Given this state of affairs, grantmaking organizations (the funders of social change organizations) should invest in the development of open-access foundation models in modalities that would be used by their grantees and educate the sector on the possibilities that this technology holds. Furthermore, it is likely that foundation model software companies will soon enough start developing integration kits and connectors to other applications. If grantmakers do not exert their (purchasing) power and also support technology developers to immerse themselves with non-profits to obtain insights on which applications those are, the social sector and the developing world may be left behind.

Some modalities of foundation models that would be particularly useful for social good include the following examples. A foundation model for satellite imagery (Martineau, 2023) would support analysis and decision making for organizations working on problems such as poverty estimation and intervention targeting, planting advice for shareholder farmers, cleaning ocean plastic pollution, climate action, and many more. A foundation model for the United States Census and American Community Survey would support several use cases encountered by social service providers, whether it is in planning emergency food distribution or working toward housing equity. A foundation model of tabular data from the operations management domain could support the logistics decision making of various non-profits, social enterprises, and government agencies.

References

Sourya Basu, Prasanna Sattigeri, Karthikeyan Natesan Ramamurthy, Vijil Chenthamarakshan, Kush R. Varshney, Lav R. Varshney, and Payel Das. Equi-tuning: Group equivariant fine-tuning of pretrained models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, February 2023.

Rachel K. E. Bellamy, Kuntal Dey, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilović, Seema Nagar, Karthikeyan Natesan Ramamurthy, John Richards, Diptikalyan Saha, Prasanna Sattigeri, Moninder Singh, Kush R. Varshney, and Yunfeng Zhang. AI fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias. *IBM Journal of Research and Development*, 63(4/5):4, July–September 2019.

Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency*, pp. 610–623, March 2021.

Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avanika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, Julian Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Rob Reich, Hongyu Ren, Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishnan Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. On the opportunities and risks of foundation models. arXiv:2108.07258, 2021.

Nelson Kibichii Bore, Ravi Kiran Raman, Isaac M. Markus, Sekou L. Remy, Oliver Bent, Michael Hind, Eleftheria K. Pissadaki, Biplav Srivastava, Roman Vaculin, Kush R. Varshney, and Komminist Weldemariam. Promoting distributed trust in machine learning and computational simulation. In *Proceedings of the IEEE International Conference on Blockchain and Cryptocurrency*, pp. 311–319, Seoul, Korea, May 2019.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, pp. 1877–1901, 2020.

Miles Brundage, Shahar Avin, Jasmine Wang, Haydn Belfield, Gretchen Krueger, Gillian Hadfield, Heidy Khlaaf, Jingying Yang, Helen Toner, Ruth Fong, Tegan Maharaj, Pang Wei Koh, Sara Hooker, Jade Leung, Andrew Trask, Emma Bluemke, Jonathan Lebensold, Cullen O'Keefe, Mark Koren, Théo Ryffel, JB Rubinovitz, Tamay Besiroglu, Federica Carugati, Jack Clark, Peter Eckersley, Sarah de Haas, Maritza Johnson, Ben Laurie, Alex Ingerman, Igor Krawczuk, Amanda Askell, Rosario Cammarota, Andrew Lohn, David Krueger, Charlotte Stix, Peter Henderson, Logan Graham, Carina Prunkl, Bianca Martin, Elizabeth Seger, Noa Zilberman, Seán Ó hÉigeartaigh, Frens Kroeger, Girish Sastry, Rebecca Kagan, Adrian Weller, Brian Tse, Elizabeth Barnes, Allan Dafoe, Paul Scharre, Ariel Herbert-Voss, Martijn Rasser, Shagun Sodhani, Carrick Flynn, Thomas Krendl Gilbert, Lisa Dyer, Saif Khan, Yoshua Bengio, and Markus Anderljung. Toward trustworthy AI development: Mechanisms for supporting verifiable claims. arXiv:2004.07213, 2020.

Jonathan Galsurkar, Moninder Singh, Lingfei Wu, Aditya Vempaty, Mikhail Sushkov, Devika Iyer, Serge Kapto, and Kush R. Varshney. Assessing national development plans for alignment with sustainable development goals via semantic search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 7753–7758, New Orleans, LA, USA, February 2018.

Chuan Li. OpenAI's GPT-3 language model: A technical overview. https://lambdalabs.com/blog/demystifying-gpt-3, June 2020.

Alexandra Sasha Luccioni, Sylvain Viguier, and Anne-Laure Ligozat. Estimating the carbon footprint of BLOOM, a 176B parameter language model. arXiv:2211.02001, 2022.

Kim Martineau. IBM and NASA team up to spur new discoveries about our planet. https://research.ibm.com/blog/ibm-nasa-foundation-models, February 2023.

Iroro Orife, Julia Kreutzer, Blessing Sibanda, Daniel Whitenack, Kathleen Siminyu, Laura Martinus, Jamiil Toure Ali, Jade Abbott, Vukosi Marivate, Salomon Kabongo, Musie Meressa, Espoir Murhabazi, Orevaoghene Ahia, Elan van Biljon, Arshath Ramkilowan, Adewale Akinfaderin, Alp Öktem, Wole Akin, Ghollah Kioko, Kevin Degila, Herman Kamper, Bonaventure Dossou, Chris Emezue, Kelechi Ogueji, and Abdallah Bashir. Masakhane – machine translation for Africa. arXiv:2003.11529, 2020.

C. K. Prahalad. *The Fortune at the Bottom of the Pyramid: Eradicating Poverty Through Profits*. Prentice Hall, Upper Saddle River, NJ, USA, 2006.

C. K. Prahalad. Bottom of the pyramid as a source of breakthrough innovations. *Journal of Product Innovation Management*, 29(1):6–12, January 2012.

Prasanna Sattigeri, Soumya Ghosh, Inkit Padhi, Pierre Dognin, and Kush R. Varshney. Fair infinitesimal jackknife: Mitigating the influence of biased training data points without refitting. In *Advances in Neural Information Processing Systems*, December 2022.

Renee Shelby, Shalaleh Rismani, Kathryn Henne, AJung Moon, Negar Rostamzadeh, Paul Nicholas, N'Mah Yilla, Jess Gallegos, Andrew Smart, Emilio Garcia, and Gurleen Virk. Sociotechnical harms: Scoping a taxonomy for harm reduction. arXiv:2210.05791, 2022.

Zheyuan Ryan Shi, Zhiwei Steven Wu, Rayid Ghani, and Fei Fang. Bandit data-driven optimization: AI for social good and beyond. arXiv:2008.11707, 2020.

Kush R. Varshney. *Trustworthy Machine Learning*. Independently Published, Chappaqua, NY, USA, 2022.

Kush R. Varshney and Homa Alemzadeh. On the safety of machine learning: Cyber-physical systems, decision sciences, and data products. *Big Data*, 5(3):246–255, 2017.

Kush R. Varshney and Aleksandra Mojsilović. Open platforms for artificial intelligence for social good: Common patterns as a pathway to true impact. In *Proceedings of the ICML AI for Social Good Workshop*, Long Beach, CA, USA, June 2019.

Ada Wan. Fairness in representation for multilingual NLP: Insights from controlled experiments on conditional language modeling. In *Proceedings of the International Conference on Learning Representations*, April 2022.