

From Artificial Intelligence Bias to Inequality in the Time of COVID-19

Miguel Luengo-Oroz

United Nations Global Pulse, New York,
NY 10017 USA

Joseph Bullock

United Nations Global Pulse
Durham University, Durham DH1 3LE, U.K.

Katherine Hoffmann Pham

United Nations Global Pulse
NYU Stern School of Business, New York,
NY 10012 USA

Cynthia Sin Nga Lam

United Nations Global Pulse

Alexandra Luccioni

Université de Montréal, Montréal, QC H3T 1J4,
Canada

■ **AS SECRETARY GENERAL** of the United Nations, Antonio Guterres said during the 2020 Nelson Mandela Annual Lecture, “COVID-19 has been likened to an X-ray, revealing fractures in the fragile skeleton of the societies we have built.” Without a doubt, the COVID-19 pandemic has exposed and exacerbated existing global inequalities. Whether at the local, national, or international scale, the gap between the privileged and the vulnerable is growing wider, resulting in a broad increase in inequality across all dimensions of society. The disease has strained health systems, social support programs, and the economy as a whole, drawing an ever-widening distinction between those with access to treatment, services, and job opportunities and those without. Global lockdown restrictions have led to increases in childcare and housework responsibilities, and most of the burden has fallen on women, further increasing existing gender inequality [1], [2]. Indigenous populations worldwide find themselves more vulnerable to infection, many times with less access

to health services or hygiene measures and limited updated scientific information about the virus and measures that can be taken to mitigate it [3]. Inequality has also pervaded the education sector, with only a subset of students able to attend safe in-person schooling or access online education when needed.

The pandemic has also increased the differences between countries, distinguishing between those which are able to access tests and diagnostic tools, personal protective equipment (PPE), medical equipment such as ventilators, and (eventually) vaccines, and those which cannot. International cooperation is critical, and as new diagnostics, medicines, and vaccines come through the development pipeline, there must be collaboration and mechanisms for joint procurement and pooling of risk. Support for initiatives such as the COVID-19 Vaccine Global Access Facility (COVAX) will help avoid myopic nationalist vaccine strategies which will result in increasing inequality and worsen outcomes for humanity as a whole. When planning public health and economic interventions, governments need to take existing inequalities into account, lest

Digital Object Identifier 10.1109/MTS.2021.3056282
Date of current version: 16 March 2021.

they end up further worsening the situation [4]. To help address these issues, the United Nations Development Programme (UNDP) has proposed a comprehensive framework for providing socioeconomic support to countries and societies in the face of COVID-19 which covers five complementary dimensions, including: 1) essential health services; 2) social protection; 3) microeconomic response and recovery programs; 4) macroeconomic fiscal and financial stimuli and policies; and 5) community-led resilience and response systems, while ensuring social equality and inclusion [5]. They stress that all of these dimensions must be taken into account when designing and deploying interventions to ensure that all necessary aspects of human and societal well-being are addressed.

AI, COVID-19, and inequality

In a recent article reviewing AI solutions against COVID-19 [6], we identified three main scales of AI applications: the molecular scale, the clinical scale, and the societal scale. We also highlighted the importance of assessing the maturity and feasibility of proposed interventions at each of these scales before deploying AI solutions. At the molecular level, much of the current research involves assisting drug discovery and development, and improving molecular diagnosis. Many of these applications remain at the research stage, given the challenges in synthesizing and testing compounds, as well as limitations in data and model sharing. Nonetheless, at least four candidate vaccines that reported the use of machine learning (ML) in their development have advanced to the clinical evaluation stage according to the World Health Organization (WHO) [7]. At a clinical level, AI has mostly been used to assist and improve patient-level assessment of COVID-19 via analysis of medical imagery and patient records. However, the extent to which the analysis from these techniques alone can be used for the diagnosis of COVID-19 is still debated by the medical community [8], and transparency and explainability of the proposed diagnoses remain overlooked by most AI approaches. Finally, on a societal scale, two crucial and complementary lines of research focus on modeling the spread of the pandemic across territories and regions, as well as its accompanying “infodemic” of misinformation. Despite reservations regarding whether models trained on data from one context are applicable in another one, data-driven models

remain paramount in analyzing the spread of the virus. Overall, while AI can be an important tool in fighting the pandemic, context-sensitive deployment is key. Working with stakeholders who have necessary domain knowledge, in addition to developing appropriate model and data sharing solutions, is critical [9], [10].

Will the AI systems applied in the fight against COVID-19 increase or decrease inequality? While the AI community is still in the early stages of application development, we have identified some persistent sources of bias which run the risk of exacerbating inequality [9]. We believe, however, that some of these biases can be mitigated (and potentially overcome) if proper assessment is included during the application development process. Generally, AI applications against COVID-19 have been developed in and for countries in the Global North, resulting in a lack of information about how the resulting tools will impact the evolution of the disease and related policies in the Global South [11]. Even within highly developed countries, many modeling and data collection efforts overlook or neglect underrepresented minorities [12]–[14].

AI systems can be inherently prone to bias which can be introduced at several different points in the application development pipeline. First, bias can be present during problem scoping, that is, by influencing the way in which the problem to be addressed by AI is framed, as well as the extent of the work itself. Second, bias can be present in the data used to train an AI system, or the way in which the data was labeled. Especially in healthcare situations, using AI approaches can reproduce and amplify existing biases in medical data sets, oftentimes not accounting for data from minority groups [15], [16]. Third, bias can be present in the choice of the algorithms themselves and how its intricate configuration parameters are tuned. Finally, bias may influence how results are evaluated and interpreted, which in turn affects how AI outputs are used. A further challenge arises when models trained and tested in one setting are then applied to another, and results are interpreted without acknowledging any necessary corrections specific to the new context. While this is not an exhaustive list of all potential sources of bias, these steps of an AI pipeline are all subject to both conscious and unconscious bias, which can reinforce existing inequalities as well as promote new ones if not properly addressed.

Examples

In the sections below, we illustrate examples of applications in which biases result (or might result) in increased inequality.

Diagnosis

There are many algorithmic approaches for supporting COVID-19 diagnosis from computer tomography (CT) and X-ray scans which frame the diagnostic problem as a classification task (i.e., identifying healthy vs. COVID-19-positive individuals), to training neural networks to detect masses and patterns in lung scans [17]–[19]. This can be particularly problematic when attempting to make meaningful conclusions based solely on medical imagery and, in particular, in geographies with a high prevalence of other diseases not included in the training data sets which might also affect the lungs (such as tuberculosis or HIV/AIDS [20]–[22]), which can be confused with COVID-19 and lead to its misdiagnosis. In addition, much of the existing medical imaging research relies on small and poorly balanced data sets that mix data from several populations without proper traceability. Prior research has shown that when training data sets are imbalanced on gender, the performance of deep learning models in radiology decreases, in particular, in the case of X-ray image data sets used to diagnose thoracic diseases [23]. Such biases might also reduce the performance of AI applications on CT and X-ray imaging of COVID-19 if not properly taken into account.

Another structural source of bias is that not all regions can afford scanner equipment. As a result, the data used to train the related AI models will not be representative and the approaches proposed may not be systematically deployed in disadvantaged regions. On the other end of the technology requirement spectrum, there are some proposals involving mobile-based diagnosis approaches [24] that are potentially promising, but warrant further exploration and validation.

Much of the existing work that utilizes AI to analyze medical imagery does not provide transparency or interpretability, delivering a categorical verdict based solely on an incoming image. This “black-box” approach may be acceptable for human-in-the-loop deployment where potential COVID-19-positive images are flagged for expert radiologists, who will then carry out further analysis manually. However, such approaches can be problematic in contexts

where medical experts are lacking or do not have enough time. Fully automated AI pipelines must be assessed in the context of their clinical impact before being deployed, which includes carefully considering and mitigating the risks in addition to assessing biases which might result into incorrect and unfair functioning.

Utilizing complementary medical data such as information regarding a patient’s gender, age, and comorbidities, as well as clinical indicators, does not only improve the accuracy of image-only approaches, but also produces results that are more interpretable for clinicians. For instance, a hybrid approach that merges algorithmic analysis of both CT scans as well as clinical features to predict the severity of COVID-19 [25] reported high accuracy rates, and the set of clinical features identified as relevant by the algorithm were coherent with those identified by previous studies. This overlap is promising for eventual clinical monitoring of COVID-19 severity, both manually and using AI-infused approaches. However, these kinds of hybrid studies should also be replicated in other regions and complemented with clinical data from incoming cases around the world, achieving better global coverage and reproducibility.

Treatment

Research has consistently shown higher rates of infection, hospitalization, and death in ethnic minorities during the COVID-19 pandemic [12], [13], [26], [27]. Nonetheless, our understanding of these inequalities remains poor, which hinders the development of solutions to combat COVID-19 and the disparities it magnifies. Gathering sufficient and accurate data on all of the social determinants of health, including race and ethnicity, is critical for effective research and development of both medical and public health interventions. However, inadequate and biased data collection is prevalent in practice. For example, data collection mechanisms are often poorly designed, with inconsistent ethnicity and race labeling [28], [29]. In fact, a systematic review has found that of 1518 COVID-19-related clinical trials registered on ClinicalTrials.gov, only one randomized controlled trial and five observational studies collected data on ethnicity [30]. Overall, fragmented and incomplete data makes it challenging for AI to succeed in furthering our understanding of COVID-19 and devising appropriate interventions.

The inequitable selection of participants as well as inconsistent presentation of demographic data has been a prevalent issue in pharmaceutical trials long before COVID-19 [31]. Prior research has shown that demographic differences often translate into differences in physical outcomes for individuals and communities. For example, immune profiles differ from person to person as a result of genetic, evolutionary, and environmental factors such as age, ethnicity, comorbidities, geographic location, nutrition, and so on [32], [33]. Hence, these dimensions should be properly accounted for when designing ML-based methodologies. We must be particularly cautious to ensure that data limitations do not lead us to develop biased algorithms that use readily available and seemingly effective but in fact, problematic proxies, such as using healthcare costs as a proxy to assess health needs, which was found to be racially biased [34]. These kinds of AI applications risk not only making inaccurate predictions, but also reinforcing systemic injustices.

Although the private sector is playing a critical role in devising interventions against COVID-19, its lack of transparency may reinforce biases and inequalities. Most of the vaccine candidates that utilized ML in their development came from corporations that made little information on their ML approach available, making it difficult for independent researchers to inspect the biases that might come from the collection and measurement of data as well as the evaluation, aggregation, and deployment processes [35]. Research conducted in private sector largely relies on data such as knowledge graphs which are not publicly accessible, and proprietary algorithms hinder the discovery, evaluation, and rectification of algorithmic biases [34].

All in all, given that AI is playing an increasingly important role in expediting the discovery and development of medical solutions, it is particularly important to ensure that the data gathered is representative of global populations and that appropriate mechanisms to audit algorithms and dissect any possible biases are in place. Otherwise, solutions developed and tested in ways that are subject to systemic, clinical, and algorithmic biases may have unintended consequences in vulnerable populations.

Epidemics

The modeling of epidemics is relevant for understanding potential infection trajectories and to

inform operational planning. Epidemiological models need careful fine-tuning to regional variables to capture age, socioeconomic status, and cultural norms. One of the most frequently used approaches is susceptible-infected-recovered (SIR) modeling, in which models are usually designed based on regional statistics capturing age, sex, and potentially other variables if data are available. However, these approaches often average over demographic and societal structures, meaning that differences in behavior between different groups are missed. Furthermore, developing epidemiological models in certain settings, and applying them to others with different cultural norms, can result in incorrect predictions. In fact, we have already seen examples of models failing to capture and explain differences in COVID-19 transmission trends between European and African contexts [36], [37]. If models are not developed in collaboration with regional experts, local populations are at a severe risk of being left behind.

Disparities in modeling efforts have also grown out of differences in data availability. For example, agent-based approaches to modeling disease spread, while able to better account for heterogeneous transmission dynamics, often require detailed data inputs. These can be challenging to acquire in under-served and ill-documented settings such as informal settlements. Similarly, many modeling approaches are designed to assess the potential for certain policy interventions, yet many of these interventions, such as social distancing or household quarantine, are impossible in highly crowded spaces where individuals need to leave their homes for water or aid collection.

In addition, many data sets which have served as important indicators for possible disease spread and recovery, as well as tools to aid in SARS-CoV-2 transmission mitigation efforts, are reliant on the access to certain technology. For example, Google mobility data [38] has provided valuable insights into the movement patterns of individuals and has been used in several modeling efforts [39], [40]. However, this fails to capture portions of society which do not have access to the mobile technology used in creating such data sets. Newly developed tools to fight the pandemic can also be reliant on technology which is unavailable to certain members of society such as contact tracing applications which

may have certain smartphone and operating system requirements.

We have recently seen a rise in the use of ML for modeling the epidemiological trends of COVID-19 [41], [42], which might be heavily influenced by data limitations and biases [6], [43]. Indeed, while many classical modeling approaches can be interrogated mathematically and computationally, we are still much less able to interpret the reasoning behind the outputs of an AI system. Therefore, while AI models may be powerful as prediction tools, we need the ability to better explore models which have the potential to significantly impact people's lives.

Infodemics

The propagation of mis- and disinformation around COVID-19 may act as yet another driver of inequality, and online social media platforms represent a particularly rich environment for the spread of infodemics. Research has shown that information is disseminated on these platforms in a viral manner, reproducing at a rate similar to that of a pandemic (i.e., exponentially) [44]. Although the rate at which true versus false claims are amplified appears to vary across platforms [44], research on Twitter has found that social media users may be more likely to share false information because it is novel [45]. The reach of misinformation may be further extended by the automated recommendation algorithms underlying many online social media platforms, which seek to identify and promote popular or viral posts to maximize engagement and garner attention [46]. Despite the widespread proliferation of misinformation, only a fraction of this content is flagged and ultimately removed or corrected [47].

Some of the most vulnerable populations in our society—such as the elderly, minorities, and other populations with low health literacy—may also be the most susceptible to the infodemic [48]. On the one hand, since much of the false information about COVID-19 involves fake cures, rumors about invincibility, and/or ineffective preventive measures, misinformation could lead these already-vulnerable populations to engage in unnecessarily risky behavior. On the other hand, these populations may be more likely to believe misinformation due to a tendency toward “inequality-driven mistrust” in which a historical legacy of discrimination and mistreatment at the hands of society and the medical

community makes them more likely to question official information and believe in conspiracy theories [49]. In sub-Saharan Africa, researchers have argued that this skepticism extends to “distrust of philanthropic institutions, distrust of developed nations, and even distrust of leaders in their own respective countries” [50].

Systematic fact-checking efforts to proactively identify and correct false information may fail to reach these vulnerable populations because these populations may engage with misinformation through private or interpersonal channels [51], and in local languages or dialects [52]. This makes it difficult for fact checkers to identify the relevant pieces of false information and to reach these populations with corrected information. Many efforts to combat the infodemic so far have focused on data sources that are easy to mine with AI (e.g., Facebook, Twitter, and online news media), rather than on offline channels such as radio or TV which may act as primary sources of information in the developing world [53]. These efforts also tend to focus on widely spoken languages such as English, Spanish, Arabic, French, and Portuguese, rather than on under-represented regional languages for which natural language processing tools may not yet be available.

Another repercussion of the infodemic has been the targeting and stigmatization of vulnerable groups. For example, there has been a reported rise in hate crimes [54] and online hate speech [55] against Chinese and other Asian minorities. There has also been a rise in xenophobic and antimigrant sentiment more generally [56], both because there is a fear that arriving migrants will transmit the virus, and because of the perceived strain that these migrants will place on already-overburdened.

Discussion

AI applications have the potential to positively contribute to the fight against COVID-19. However, AI can also amplify the biases and inequalities that have emerged or become more extreme during the pandemic. As in any other application of AI, bias can be introduced at multiple points throughout the modeling pipeline. As COVID-19 furthers inequality along every dimension of society, it is important to identify and address AI applications which might contribute to this trend.

AI researchers and practitioners should pay special attention to bias introduced in the problem definition and data collection stages. Public health algorithms play important roles in decision-making, from shaping critical care protocols to deciding how to distribute vaccines. Women and minorities are also often not properly represented in data sets whose use may result in medical treatments and services. In addition, nonpharmaceutical interventions are increasingly based on measures constructed from digital traces (e.g., mobility patterns recorded from cell phone usage), although these do not account for those individuals without access to mobile networks or the internet. Finally, the spread of mis- and disinformation can be amplified by AI, which impacts more those with less health literacy.

In all of these examples, AI-driven biases and their implications for inequality should be assessed from the initial project development stages through the presentation and use of outputs. Furthermore, AI systems should be adapted to local contexts, following and accounting for cultural and social norms. It is also important to recognize that in some cases, the proposed solutions might be fair and efficient with respect to a particular problem statement, but may increase inequality along other dimensions. Contrary to other AI applications where the goal is to identify a single model to beat a single metric and benchmark, in situations where vulnerable populations are at risk, a more sensitive approach should be taken with a focus on outliers as metrics to assess failure as well as success. This analysis may require the problem to be approached from different angles. For example, outliers in a patient group might represent diversity and, rather than focusing on average performance, the more relevant task for modelers might be to understand for whom the model is performing worst.

SOME KEY CONCEPTS to help reduce bias in AI applications include: the transferability of models and their adaptation to local contexts; federated learning strategies which may make it possible to include patient data from multiple cohorts in a privacy-preserving way; and the deployment of interpretable models which allow users to interrogate the patterns and reasons for model decisions, which is particularly relevant in high-risk settings. We hope this article will stimulate some

reflections in the AI community aiming to support the COVID-19 response and will encourage researchers to assess whether biases in their AI applications will amplify inequality from a health, economic, and social perspective. ■

Acknowledgments

The United Nations Global Pulse is supported in part by the Governments of Sweden and Germany and in part by the William and Flora Hewlett Foundation. The work of Joseph Bullock was supported by the United Kingdom Research and Innovation - Science and Technology Facilities Council (UKRI-STFC) under Grant ST/P006744/1. The work of Alexandra Luccioni was supported by funding from IVADO and Mila.

References

- [1] L. Farré et al., "How the COVID-19 lockdown affected gender inequality in paid and unpaid work in Spain," IZA Inst. Labor Econ., Bonn, Germany, Tech. Rep. IZA DP no. 13434, 2020.
- [2] C. Rivera et al., *Gender Inequality and the COVID-19 Crisis: A Human Development Perspective*. New York, NY, USA: United Nations Development Programme, 2020.
- [3] T. Power et al., "COVID-19 and indigenous peoples: An imperative for action," *J. Clin. Nursing*, vol. 29, nos. 15–16, pp. 2737–2741, Aug. 2020, doi: 10.1111/jocn.15320.
- [4] N. A. Ashford et al., "Addressing inequality: The first step beyond COVID-19 and towards sustainability," *Sustainability*, vol. 12, no. 13, p. 5404, Jul. 2020.
- [5] UN. (2020). *A UN Framework for the Immediate Socio-Economic Response to COVID-19*. [Online]. Available: <https://unsdg.un.org/sites/default/files/2020-04/UN-framework-for-the-immediate-socio-economic-response-to-COVID-19.pdf>
- [6] J. Bullock et al., "Mapping the landscape of artificial intelligence applications against COVID-19," *J. Artif. Intell. Res.*, vol. 69, pp. 807–845, 2020.
- [7] WHO. (2020). *Draft Landscape of COVID-19 Candidate Vaccines*. [Online]. Available: <https://www.who.int/who-documents-detail/draft-landscape-of-covid-19-candidate-vaccines>
- [8] M. Weinstock, A. Echenique, and J. E. A. Russell, "Chest X-ray findings in 636 ambulatory patients with COVID-19 presenting to an urgent care center: A normal chest X-ray is no guarantee," *J. Urgent Care Med.*, vol. 14, no. 7, pp. 13–18, 2020.

- [9] A. Luccioni et al., "Considerations, good practices, risks and pitfalls in developing AI solutions against COVID-19," in *Proc. Harvard CRCS Workshop AI Social Good*, 2020.
- [10] M. Luengo-Oroz et al., "Artificial intelligence cooperation to support the global response to COVID-19," *Nature Mach. Intell.*, vol. 2, pp. pages295–297, May 2020.
- [11] J. A. Oldekop et al., "COVID-19 and the case for global development," *World Develop.*, vol. 134, Oct. 2020, Art. no. 105044.
- [12] M. K. Evans, "Covid's color line—Infectious disease, inequity, and racial justice," *New England J. Med.*, vol. 383, no. 5, pp. 408–410, Jul. 2020.
- [13] E. J. Williamson et al., "Factors associated with COVID-19-related death using OpenSAFELY," *Nature*, vol. 584, no. 7821, pp. 430–436, 2020.
- [14] R. Mathur et al., "Ethnic differences in COVID-19 infection, hospitalisation, and mortality: An OpenSAFELY analysis of 17 million adults in England," *medRxiv*, 2020. [Online]. Available: <https://www.medrxiv.org/content/early/2020/09/23/2020.09.22.20198754>
- [15] I. Y. Chen, S. Joshi, and M. Ghassemi, "Treating health disparities with artificial intelligence," *Nature Med.*, vol. 26, no. 1, pp. 16–17, Jan. 2020.
- [16] N. Turner Lee, "Detecting racial bias in algorithms and machine learning," *J. Inf. Commun. Ethics Soc.*, vol. 16, no. 3, pp. 252–260, Aug. 2018.
- [17] S. Wang et al., "A deep learning algorithm using CT images to screen for corona virus disease (COVID-19)," *medRxiv*, 2020.
- [18] J. Chen et al., "Deep learning-based model for detecting 2019 novel coronavirus pneumonia on high-resolution computed tomography: A prospective study," *medRxiv*, 2020.
- [19] O. Gozes et al., "Coronavirus detection and analysis on chest CT with deep learning," 2020, *arXiv:2004.02640*. [Online]. Available: <http://arxiv.org/abs/2004.02640>
- [20] A. Khan et al., "Imaging lung manifestations of HIV/AIDS," *Ann. Thoracic Med.*, vol. 5, no. 4, p. 201, 2010.
- [21] T. Franquet, "Imaging of pulmonary viral pneumonia," *Radiology*, vol. 260, no. 1, pp. 18–39, Jul. 2011.
- [22] E. Skoura, A. Zumla, and J. Bomanji, "Imaging in tuberculosis," *Int. J. Infectious Diseases*, vol. 32, pp. 87–93, Mar. 2015.
- [23] A. J. Larrazabal et al., "Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis," *Proc. Nat. Acad. Sci. USA*, vol. 117, no. 23, pp. 12592–12594, 2020.
- [24] X. Li, C. Li, and D. Zhu, "COVID-MobileXpert: On-device COVID-19 patient triage and follow-up using chest X-rays," 2020, *arXiv:2004.03042*. [Online]. Available: <http://arxiv.org/abs/2004.03042>
- [25] W. Shi et al., "Deep learning-based quantitative computed tomography model in predicting the severity of COVID-19: A retrospective study in 196 patients," *Ann. Transl. Med.*, vol. 9, no. 3, p. 216, 2021, doi: 10.21037/atm-20-2464.
- [26] CDC. (2020). *CDC COVID Data Tracker Maps, Charts, and Data Provided by the CDC*. [Online]. Available: <https://covid.cdc.gov/covid-data-tracker/#demographics>
- [27] K. Khunti et al., "Is ethnicity linked to incidence or outcomes of COVID-19?" *BMJ*, vol. 369, p. m1548, Apr. 2020. [Online]. Available: <https://www.bmj.com/content/369/bmj.m1548>
- [28] S. M. Meraji. (2020). *The News Beyond the COVID Numbers*. [Online]. Available: <https://www.npr.org/2020/04/21/840609912/the-news-beyond-the-covid-numbers>
- [29] Mississippi State Department of Health. (2020). *Mississippi Coronavirus Disease 2019 (COVID-19) Data*. [Online]. Available: https://msdh.ms.gov/msdhsite/_static/14,0,420.html#Mississippi
- [30] D. Pan et al., "The impact of ethnicity on clinical outcomes in COVID-19: A systematic review," *EClin. Med.*, vol. 23, Jun. 2020, Art. no. 100404.
- [31] F. Zhang and J. Finkelstein, "Inconsistency in race and ethnic classification in pharmacogenetics studies and its potential clinical implications," *Pharmacogenomics Personalized Med.*, vol. 12, pp. 107–123, Jul. 2019.
- [32] J. Sanz, H. E. Randolph, and L. B. Barreiro, "Genetic and evolutionary determinants of human population variation in immune responses," *Current Opinion Genet. Develop.*, vol. 53, pp. 28–35, Dec. 2018.
- [33] P. Zimmermann and N. Curtis, "Factors that influence the immune response to vaccination," *Clin. Microbiol. Rev.*, vol. 32, no. 2, Mar. 2019.
- [34] Z. Obermeyer et al., "Dissecting racial bias in an algorithm used to manage the health of populations," *Science*, vol. 366, no. 6464, pp. 447–453, Oct. 2019. [Online]. Available: <https://science.sciencemag.org/content/366/6464/447>
- [35] H. Suresh and J. V. Gutttag, "A framework for understanding unintended consequences of machine learning," 2020, *arXiv:1901.10002*. [Online]. Available: <https://arxiv.org/abs/1901.10002>

- [36] L. Nordling, "The pandemic appears to have spared Africa so far. Scientists are struggling to explain why," *Science*, vol. 26, 2020.
- [37] M. Mbow et al., "COVID-19 in Africa: Dampening the storm?" *Science*, vol. 369, no. 6504, pp. 624–626, Aug. 2020.
- [38] Google. (2020). *See How Your Community Is Moving Around Differently Due to COVID-19*. [Online]. Available: <https://www.google.com/covid19/mobility/>
- [39] B. Jeffrey et al., "Report 24: Anonymised and aggregated crowd level mobility data from mobile phones suggests that initial compliance with COVID-19 social distancing interventions was high and geographically consistent across the UK," MCR Centre Global Infectious Disease Anal., London, U.K., COVID-19 Rep. 24, 2020.
- [40] P. Nouvellet et al., "Report 26 reduction in mobility and COVID-19 transmission," MCR Centre Global Infectious Disease Anal., London, U.K., COVID-19 Rep. 26, 2020.
- [41] V. Lampos et al., "Tracking COVID-19 using online search," 2020, *arXiv:2003.08086*. [Online]. Available: <http://arxiv.org/abs/2003.08086>
- [42] R. Dandekar and G. Barbastathis, "Neural network aided quarantine control model estimation of COVID spread in Wuhan, China," 2020, *arXiv:2003.09403*. [Online]. Available: <http://arxiv.org/abs/2003.09403>
- [43] B. McCall, "COVID-19 and artificial intelligence: Protecting health-care workers and curbing the spread," *Lancet Digital Health*, vol. 2, no. 4, pp. 166–167, 2020.
- [44] M. Cinelli et al., "The COVID-19 social media infodemic," 2020, *arXiv:2003.05004*. [Online]. Available: <http://arxiv.org/abs/2003.05004>
- [45] S. Vosoughi, D. Roy, and S. Aral, "The spread of true and false news online," *Science*, vol. 359, no. 6380, pp. 1146–1151, Mar. 2018.
- [46] C. Meserole. (2018). *How Misinformation Spreads on Social Media—And What to Do About It*. [Online]. Available: <https://www.brookings.edu/blog/orderfromchaos/2018/05/09/how-misinformation-spreads-on-social-media-and-what-to-do-about-it/>
- [47] (2020). *Types, Sources, and Claims of COVID-19 Misinformation*. [Online]. Available: <https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2020-04/Brennen%20-%20COVID%2019%20Misinformation%20FINAL%20%283%29.pdf>
- [48] L. Paakkari and O. Okan, "COVID-19: Health literacy is an underestimated problem," *Lancet Public Health*, vol. 5, no. 5, pp. e249–e250, May 2020.
- [49] J. Jaiswal, C. LoSchiavo, and D. C. Perlman, "Disinformation, misinformation and inequality-driven mistrust in the time of COVID-19: Lessons unlearned from AIDS denialism," *AIDS Behav.*, vol. 24, pp. 2776–2780, Oct. 2020.
- [50] E. Nwankwo, C. Okolo, and C. Habonimana, "Topic modeling approaches for understanding COVID-19 misinformation spread in sub-Saharan Africa," in *Proc. AI Social Good Workshop*. Cambridge, MA, USA: Harvard Center for Research on Computation and Society, 2020. [Online]. Available: <https://crs.seas.harvard.edu/publications/topic-modeling-approaches-understanding-covid-19-misinformationspread-sub-saharan>
- [51] I. V. Pasquetto et al. (2020). *Understanding Misinformation on Mobile Instant Messengers (MIMs) in Developing Countries*. [Online]. Available: <https://shorensteincenter.org/wp-content/uploads/2020/06/Misinfo-on-MIMs-Shorenstein-Center-May-2020.pdf>
- [52] K. Traoré, "Who fact checks online disinformation in West Africa's Bambara language?," *Rising Voices*, Jun. 2020. [Online]. Available: <https://rising.globalvoices.org/blog/2020/06/15/who-fact-checks-online-disinformation-in-west-africas-bambara-language/>
- [53] L. Newman, P. Hutchinson, and D. Meekers, "Key findings of the 3-2-1 service COVID-19 surveys: Information (Wave 1)," Viamo, Saskatchewan, Canada, Tech. Rep., 2020. [Online]. Available: https://viamo.io/wp-content/uploads/2020/09/Updated_-3-2-1-Service-COVID-19-Survey_-Information-1.pdf
- [54] A. R. Gover, S. B. Harper, and L. Langton, "Anti-asian hate crime during the COVID-19 pandemic: Exploring the reproduction of inequality," *Amer. J. Criminal Justice*, vol. 45, no. 4, pp. 647–667, Aug. 2020.
- [55] L. Schild et al., "'Go eat a bat, Chang!': An early look on the emergence of sinophobic behavior on Web communities in the face of COVID-19," 2020, *arXiv:2004.04046*. [Online]. Available: <http://arxiv.org/abs/2004.04046>
- [56] UN News. (2020). *COVID-19 Stoking Xenophobia, Hate and Exclusion, Minority Rights Expert Warns*. [Online]. Available: <https://news.un.org/en/story/2020/03/1060602>

Miguel Luengo-Oroz is currently the Chief Data Scientist with the United Nations Global Pulse, New York, NY, USA, and the Artificial Intelligence (AI) and Big Data Innovation Initiative of the United Nations (UN) Secretary-General. He is also with the Universidad Politécnica de Madrid, Madrid.

Over the last decade, as the first data scientist at the UN, his research has focused on bringing AI to operations and policy in domains, including poverty, food security, epidemics, refugees, conflict, human rights, gender, infodemics, and climate. He has participated in multiple advisory bodies as the expert group on AI strategy for the government of Spain. He is the Founder of the global health social enterprise SpotLab, Madrid, Spain, and the Inventor of the games-for-health platform MalariaSpot.org.

Dr. Luengo-Oroz is an Ashoka Fellow. He received the MIT TR35 Award and the European Responsible Research and Innovation Award.

Joseph Bullock is currently pursuing the Ph.D. degree with the Department of Physics, Institute for Data Science, Durham University, Durham, U.K.

He is also a Researcher with United Nations Global Pulse, New York, NY, USA. He also holds a position of Research Associate at the RiskEcon Lab, Courant Institute for Mathematical Sciences, New York University, New York. His research focuses on developing artificial intelligence and numerical modeling techniques in the context of humanitarian and physics-based scenarios, with specific applications in particle physics, automated satellite image analysis, and disease modeling. In addition, he works on designing theoretical frameworks for quantifying and communicating uncertainties in machine learning systems and understanding why certain models fail.

Katherine Hoffmann Pham received the B.A. degree in international relations and economics and the M.A. degree in international policy studies from Stanford University, Stanford, CA, USA, both in 2011. She is currently pursuing the Ph.D. degree in information systems with the NYU Stern School of Business, New York, NY, USA.

She spent four years working at the Innovations for Poverty Action, Washington, DC, USA. She is

currently a Researcher with the United Nations Global Pulse, New York. Her research examines how big data, machine learning, and technology can be used to address policy problems, with a focus on the digital future of urban mobility and international migrant and refugee movements.

Cynthia Sin Nga Lam is currently an Infodemic Analyst at the United Nations Global Pulse, New York, NY, USA, supporting the COVID-19 response, a Contract Consultant at World Health Organization, Geneva, Switzerland, advising on digital health communications, and an agile-certified (PMI-ACP) practitioner managing various innovation projects. She is passionate about her variety of work connecting the dots and strives to innovate through making connections across different disciplines. Her research interests include the application of artificial intelligence in various areas of health (public health, health education, and consumer health) and health literacy.

Alexandra Luccioni received the Ph.D. degree in cognitive computing from the Université du Québec à Montréal, Montréal, QC, Canada, in 2018.

She is currently a Researcher working on Artificial Intelligence for Humanity Initiatives at the Mila Institute, Montréal, QC, Canada, where she leads projects at the nexus of artificial intelligence (AI) and social issues, such as climate change, education, and health. She spent two years working in applied research, applying deep learning in industries such as customer service and banking. Since joining Mila in 2019, she has organized and led many AI for social good initiatives, conferences, and workshops. She is also highly involved in her community, volunteering for initiatives such as Women in Machine Learning, Climate Change AI, and Kids Code Jeunesse.

■ Direct questions and comments about this article to Miguel Luengo-Oroz, United Nations Global Pulse, New York, NY 10017, USA; miguel@unglobalpulse.org.