

---

## 33. Addressing global inequity in AI development

*Chinasa T. Okolo*

---

### A GLOBAL OUTLOOK ON ARTIFICIAL INTELLIGENCE

Global inequality renders itself visible, especially in the development of artificial intelligence (AI) and machine learning (ML). In an analysis of publications at two major ML conference venues, NeurIPS 2020 and ICML 2020, Chuvpilo (2020) found that of the top ten countries in terms of publication index (calculated by treating a publication as a unit of one and splitting up the unit equally by authorship), none were in Latin America, Africa, or Southeast Asia. Vietnam, the highest-placed country of these groups, comes in 27th place. Of the top ten institutions by publication index, eight out of ten were based in the United States, including American tech giants like Google, Microsoft, and Facebook. Indeed, the full lists of the top 100 universities and top 100 companies by publication index include no companies or universities based in Africa or Latin America. Although conference publications are just one metric, they remain the predominant medium in which progress in AI is disseminated, and as such serve to be a signal of who is generating research.

Other work such as the Global AI Index from Tortoise Media (2020), which claims to be the “first index to benchmark nations on their level of investment, innovation and implementation of artificial intelligence”, ranks the United States, China, and the United Kingdom in the first, second, and third spots, respectively. Within the top 50, other countries in the Global South include India (#17), Brazil (#39), Malaysia (#43), Mexico (#44), Chile (#45), Argentina (#46), Colombia (#49), and Uruguay (#50). The Global AI Vibrancy Ranking produced by researchers in the Stanford Institute for Human-Centered Artificial Intelligence (2021) uses 23 economic and research and development indicators to rank 29 countries to highlight global progress made towards AI. Their rankings only include four countries in the Global South with China, India, Brazil, and Malaysia being ranked in the 2nd, 3rd, 10th, and 17th places, respectively. The predominance of the United States in these rankings is consistent with its economic and cultural dominance, just as the appearance of China with the second-highest index is a marker of its growing might. Also comprehensible is the relative absence of countries in the Global South, given the exploitation and underdevelopment of these regions by European colonial powers (Frank, 1967; Rodney, 1972; Jarosz, 2003; Bruhn & Gallego, 2012). While India is highlighted as a standout in AI research in Southeast Asia, the appearance of Malaysia on both rankings indicates a possible expansion point for AI development. Additionally, the inclusion of a significant number of Latin American countries on the Global AI Index suggests that this region could potentially be a significant hub for AI research and development within the Global South.

Current global inequality in AI development involves both a concentration of profits and a danger of ignoring the contexts in which AI is applied. As AI systems become increasingly integrated into society, those responsible for developing and implementing such systems stand to profit to a large extent. If these players are predominantly located outside of the

Global South, a disproportionate share of economic benefit will fall also outside of this region, exacerbating extant inequality. Furthermore, the ethical application of AI systems requires knowledge of the contexts in which they are to be applied. As recent research (Grush, 2015; De La Garza, 2020; Coalition for Critical Technology, 2020; Beede et al., 2020; Sambasivan et al., 2021) has highlighted, work that lacks such contextual knowledge can fail to help the targeted individuals, and can even harm them (e.g., misdiagnoses in medical applications, denied loans, incorrect crop yields).

Whether explicitly in response to these problems or not, calls have been made for broader inclusion in the development of AI (Asemota, 2018; Lee et al., 2019). At the same time, some have acknowledged the limitations of inclusion. Sloane et al. (2020) describe and argue against participation-washing, whereby the mere fact that somebody has participated in a project lends it moral legitimacy. In this work, we employ a post-colonial critical development studies approach to focus on the implications of participation for global inequality, concentrating particularly on the limitations in which inclusion in AI development is practiced in the Global South. We look specifically at how this plays out in the construction of datasets and establishment of research labs and conclude with a discussion of opportunities for ameliorating the power imbalance in AI development.

## DATASETS

Given the centrality of large amounts of data in today's ML systems, there would appear to be substantial opportunities for inclusion in data collection and labeling processes. While there are benefits to more diverse participation in data-gathering pipelines, this approach does not go far enough in addressing global inequality in AI development. As we consider ways in which to improve the inclusion of stakeholders from the Global South, this section discusses what kinds of problems can be alleviated, what forms of data labeling currently look like, barriers to participation, and the deep problems with this form of inclusion.

Data collection itself is a practice fraught with problems of inclusion and representation. Two large, publicly available image datasets, ImageNet (Deng et al., 2009; Russakovsky et al., 2015) and OpenImages (Krasin et al., 2017), are US- and Euro-centric (Shankar et al., 2017). Shankar et al. (2017) further argue that models trained on these datasets perform worse on images from the Global South. For example, images of grooms are classified with lower accuracy when they come from Ethiopia and Pakistan, compared with images of grooms from the United States. Along this vein, DeVries et al. (2019) show that images of the same word, like "wedding" or "spices", look very different when queried in different languages, as they are presented distinctly in different cultures. Thus, publicly available object recognition systems fail to correctly classify many of these objects when they come from the Global South. Representative training datasets are crucial to allowing models to learn how certain objects and concepts are represented in different cultures.

The importance of data labeling in machine-learning research and development has led to crowdsourcing, whereby anonymous individuals are remunerated for completing this work. Large tech companies such as Uber and Alphabet rely heavily on these services, with some paying millions of dollars monthly to third-party firms (Synced, 2019). A major venue for crowdsourcing work is Amazon Mechanical Turk; according to Difallah et al. (2018), less than 2% of Mechanical Turk workers come from the Global South (a vast majority come from

the United States and India). Other notable companies in this domain, Samasource, Scale AI, and Mighty AI also operate in the United States, but crowdsource workers from around the world, primarily relying on low-wage workers from sub-Saharan Africa and Southeast Asia (Murgia, 2019). In the Global South, local companies have begun to proliferate, like Fastagger in Kenya, Sebenz.ai in South Africa, and Supahands in Malaysia. As AI development continues to scale, the expansion of these companies opens the door for low-skilled laborers to enter the workforce but also presents a chance for exploitation to continue to occur.

### **Barriers to Data Labeling**

There exist many barriers to equitable participation in data labeling. First, a computing device (laptop, desktop, or tablet) and stable internet access are required for access to most data labeling platforms. These goods are highly correlated with socioeconomic status and geographic location, thus serving as a barrier to participation for many people situated in low-resource settings (Harris et al., 2017). A reliable internet connection is necessary for finding tasks to complete, completing those tasks, and accessing the remuneration for those tasks. Further, those in the Global South pay higher prices for Internet access compared with their counterparts in the Global North (Nzekwe, 2019). Another barrier lies in the method of payment for data labeling services on some platforms. For example, Amazon Mechanical Turk, a widely used platform for finding data labelers, only allows payment to a US bank account or in the form of an Amazon.com gift card (Amazon, 2020). Such methods of payment may not be desired by a potential worker and can serve as a deterrent to work for platforms that employ similar restrictive payment methods.

### **Problems with Data Collection and Labeling**

After having discussed the benefits of incorporating data labeling as one part of inclusion, as well as some of the barriers to participation it has, we finish this section by discussing issues associated with data labeling. At a cursory glance, having labelers who represent a diversity of backgrounds might appear largely beneficial, as it would allow for objects that might not be recognized and labeled appropriately (e.g., “wedding”) by one group of people to be done by another. Additionally, data labelers are prone to bringing their own stereotypes and biases to the task at hand. Diversifying the labeler population could help dilute the pool of shared biases that may propagate into a dataset. For example, it has been shown that MSCOCO (Lin et al., 2014), a commonly used object detection and image captioning dataset, contains strong gender biases in the image captions (Hendricks et al., 2018; Bhargava & Forsyth, 2019). If the population of dataset labelers for MSCOCO consisted of people more aware of the problems with gender stereotypes, or even people with very different gender identities, perhaps the biases in the captions might not manifest with that level of prevalence.

With respect to data collection, current practices often neglect consent and poorly represented areas of the Global South. Image datasets are often collected without consent from the people involved, even in pornographic contexts (Birhane & Prabhu, 2021; Paullada et al., 2021), while companies, academic institutions, and other entities benefit from their use. Work from Jo and Gebru (2020) suggests drawing from the long tradition or archives when collecting data because this is a discipline that has already been thinking about challenges like consent and privacy. Indeed, beyond a possible honorarium for participation in the data collection

process, no large-scale, successful schema currently exists for compensating users for the initial and continued use of their data in machine-learning systems, although some efforts are currently underway (Kelly, 2020). However, the issue of compensation eludes the question of whether such large-scale data collection should occur in the first place. Indeed, the process of data collection can contribute to an “othering” of the subject and cement inaccurate or harmful beliefs. Even if data comes from somewhere in the Global South, it is often from the perspective of an outsider (Wang et al., 2020) who may not understand the respective context or may have an agenda counter to the interest of the subject. Such values can be reflected in the data captured, as has been extensively studied in the case of photography (Ranger, 2001; Batziou, 2011; Thompson, 2016). Ignorance of context can cause harm, as Sambasivan et al. (2021) discuss in the case of fair ML in India, where distortions in the data (e.g., a given sample corresponds to multiple individuals because of shared device usage) distort the meaning of fairness definitions that were formulated in Western contexts. Furthermore, the history of phrenology reveals the role that the measurement and classification of colonial subjects had in justifying domination (Bank, 1996; Poskett, 2013). Denton et al. (2020) argue the need to interrogate more deeply the norms and values behind the creation of datasets, as they are often extractive processes that benefit only the dataset collector and users.

As another significant part of the data collection pipeline, data labeling is a low-paying job involving rote, repetitive tasks that offer no room for upward mobility. Individuals may not require many technical skills to label data, but they do not develop any meaningful technical skills either. The anonymity of platforms like Amazon Mechanical Turk inhibits the formation of social relationships between the labeler and the client that could otherwise have led to further educational opportunities or better remuneration. Although data is central to the AI systems of today, data labelers receive only a disproportionately tiny portion of the profits of building these systems. In parallel with colonial projects of resource extraction, data labeling as the extraction of meaning from data is no way out of a cycle of colonial dependence. In reference to this parallel, Couldry and Mejias (2019) characterize the term “data colonialism” as a system that exploits human capital through the production of data for technological processes. Exporting these kinds of jobs follows in the long history of colonialism (Mohamed et al., 2020), with the groups on the receiving end of the labor showing great gains in the form of strong AI models while the groups on the giving end of the labor receive few benefits from their work.

The people doing the work of data labeling have been termed “ghost-workers” (Gray & Suri, 2019). The labor of these unseen workers generates massive earnings that others capture, leading to a significant disparity between the millions in profits earned by data labeling companies and worker income; for example, data workers at Samasource earn around US\$8 a day (Lee, 2018) while the company reportedly made \$25 million in revenue during 2019 (Craft .co, 2019). While Lee (2018) notes that US\$8 may well be a living wage in certain areas, the massive disparity is poignant given the importance of these workers to the core businesses of these companies. While data labeling is not as physically intensive as traditional factory labor, workers report the pace and volume of their tasks as “mentally exhausting” and “monotonous” due to the strict requirements needed for labeling images, videos, and audio to client specifications (Gent, 2019; Croce & Musa, 2019). The lack of protections seen for data labelers in the Global South emphasizes the need for labor protection laws to address these power imbalances and improve working conditions (Kaye, 2019). As large tech companies continue to establish AI research labs within the Global South, such protections will be essential in enforcing safeguards for all workers across the ML development lifecycle.

## RESEARCH LABS

Establishing research labs has been essential for major tech companies to advance the development of their respective technologies while providing valuable contributions to the field of computer science (Nature, 1915). In the United States, General Electric (GE) Research Laboratory is widely accepted as the first industrial research lab, providing early technological achievements to GE and establishing them as a leader in industrial innovation (Center, 2011). As AI becomes more important to the bottom lines of many large tech companies, industrial research labs have spun out that solely focus on AI and its respective applications. Companies from Google to Amazon to Snapchat have doubled down in this field and opened labs leveraging AI for web search, language processing, video recognition, voice applications, and much more. As AI becomes increasingly integrated into the livelihoods of consumers around the world, tech companies have recognized the importance of democratizing AI development and moving it outside the bounds of the Global North. Of five notable tech companies developing AI solutions (Google, Microsoft, IBM, Facebook, and Amazon), Google, Microsoft, and IBM have research labs in the Global South, and all have either development centers, customer support centers, or data centers within these regions. Despite their presence throughout the Global South, AI research centers tend to be concentrated in certain countries. Within Southeast Asia, the representation of lab locations is limited to India; in South America, representation is limited to Brazil. Sub-Saharan Africa has a wider spread in locations, with AI labs established in Accra, Ghana; Nairobi, Kenya; and Johannesburg, South Africa.

### **Barriers to Establishing Research Labs**

For a company to choose to establish an AI research center, the company must believe this initiative to be in its financial interest. Unfortunately, several barriers exist. The necessity of generating reliable returns for shareholders precludes ventures that appear too risky, especially for smaller companies. The perception of risk can take a variety of forms and is at risk of being influenced by stereotypes to differing extents. Two such factors are political and economic instability or a relatively lower proportion of tertiary formal education in the local population, which can be traced to the history of colonial exploitation and underdevelopment (Rodney, 1972; Jarosz, 2003; Bruhn & Gallego, 2012), whereby European colonial powers extracted labor, natural resources, and economic surplus from colonies, while at the same time subordinating their economic development to that of the metropolises. Given this history, the establishment of top-tier research institutions to advance technical training and AI development in the Global South will be a significant challenge without sufficient investment from local governments and private entities.

Although several tech companies have established research facilities across the world and in the Global South, these efforts remain insufficient at addressing long-term problems in the AI ecosystem. A recent report from Georgetown University's Center for Security and Emerging Technologies (CSET) details the establishment of AI labs by US companies, namely Facebook, Google, IBM, and Microsoft, abroad (Heston & Zwetsloot, 2020). The report notes that while 68% of the 62 AI labs are located outside of the United States, 68% of the staff are located within the United States. Therefore, the international offices remain half as staffed on average relative to the domestic locations. Additionally, none of these offices are located in South America and only four are in Africa. To advance equity within AI and

improve inclusion efforts, it is imperative that companies not only establish locations in under-represented regions but also hire employees and include stakeholders from those regions in a proportionate manner.

While the opening of data centers and AI research labs in the Global South initially appears beneficial for the local workforce, these positions may require technical expertise which members of the local population might not have. This would instead introduce opportunities for displacement by those from the Global North who have had more access to specialized training needed to develop, maintain, and deploy AI systems. Given the unequal distribution of AI development globally, it is common for AI researchers and practitioners to work and study in places outside of their home countries (i.e., outside of the Global South). For example, the current director of Google AI Accra, originally from Senegal and a doctoral graduate from Pierre and Marie Curie University in Paris, was recruited to Google from Facebook (now Meta) AI Research in Menlo Park, CA. However, Cisse's return to the African continent is particularly notable given the significant "brain drain" in fields such as medicine and engineering. While the directors of many research labs established in the Global South have experience working in related contexts, we find that local representation is sorely lacking at both the leadership and general workforce levels. Grassroots AI education and training initiatives by communities such as Deep Learning Indaba, Data Science Africa, and Khipu in Latin America aim to increase local AI talent, but since these initiatives are less than five years old, it is hard to measure their current impact on improving the pipeline of AI researchers and ML engineers. However, with the progress made by these organizations publishing novel research at premier AI conferences, hosting conferences of their own, and much more, the path to inclusive representation in the global AI workforce is strengthening.

### **Formation of International AI Labs**

The CSET report also notes that AI labs form abroad generally in one of three ways: through the acquisition of startups; by establishing partnerships with local universities or institutions; and by relocating internal staff or hiring new staff in these locations (Heston & Zwetsloot, 2020). The first two of these methods may favor locations with an already-established technological or AI presence, as many AI startups are founded in locations where a financial and technological support system exists for them. Similarly, the universities with whom tech companies choose to partner are often already leaders in the space, as evidenced by Facebook's partnership with Carnegie Mellon professors and MIT's partnerships with both IBM and Microsoft. The general strategy of partnering with existing institutions and of acquiring startups has the potential to reinforce existing inequities by investing in locations with already thriving tech ecosystems. An exception to this is Google's establishment of an AI research office along with its investment in infrastructure, skills training, and startups in Ghana (Asemota, 2018). Long-term investment and planning in the Global South can form the stepping stones for broadening AI to include underrepresented and marginalized communities.

Even with long-term investment into regions in the Global South, the question remains whether local residents are provided opportunities to join management and contribute to important strategic decisions. Several organizations have emphasized the need for AI development within a country to happen at the grassroots level so that those implementing AI as a solution understand the context of the problem being solved (Mbayo, 2020; Gul, 2019). The necessity of Indigenous decision-making is just as important in negotiating the values that

AI technologies are to instill, such as through AI ethics declarations that are at the moment heavily Western-based (Jobin et al., 2019). Although this is critical not only to the success of individual AI solutions but also to equitable participation within the field at large, more can and should be done. True inclusion necessitates that underrepresented voices be found in all ranks of a company's hierarchy, including in positions of upper management and senior leadership. Tech companies establishing a footprint in regions within the Global South are uniquely positioned to offer such opportunities to natives of these respective regions. Taking advantage of this ability will be critical to ensuring that the benefits of AI apply not only to technical problems that arise in the Global South but to socioeconomic inequalities that persist around the world.

## **Opportunities**

In the face of global inequality in AI development, there are a few promising opportunities to engage diverse stakeholders in more inclusive, fulfilling ways. After examining existing barriers and challenges to equitable participation in AI development, this section discusses the range of grassroots AI initiatives within the Global South, what large tech companies can learn from such approaches to AI research and development, and suggestions on providing work opportunities beyond data labeling as means of representation in the development of ML models and AI systems.

## **Grassroots AI Initiatives**

While AI and technology in general have long excluded marginalized populations, there has been a strong emergence of grassroots efforts by organizations to ensure that Indigenous communities are actively involved as stakeholders of AI. Black in AI, a nonprofit organization with worldwide membership, was founded to increase the global representation of Black-identifying students, researchers, and practitioners in the field of AI, and has made significant improvements in increasing the number of Black scholars attending and publishing in NeurIPS and other premier AI conferences (Earl, 2020; Silva, 2021). Inclusion in AI is extremely sparse in higher education and recent efforts by Black in AI have focused on instituting programming to support members in applying to graduate programs and in pursuing postdoctoral and faculty positions. Other efforts such as Khipu, based in Latin America, have been established to provide a venue to train aspiring AI researchers in advanced ML topics, foster collaborations, and actively participate in how AI is being used to benefit Latin America. Other communities based on the African continent such as Data Science Africa and Deep Learning Indaba have expanded their efforts, establishing conferences, workshops, and dissertation awards, and developing curricula for the broader African AI community. These communities are clear about their respective missions and the focus of collaboration. Notably, Masakhane, a grassroots organization focused on improving the representation of African languages in the field of natural language processing, shares a powerful sentiment on how AI research should be approached (Masakhane, 2021):

Masakhane are not just annotators or translators. We are researchers. We can likely connect you with annotators or translators but we do not support shallow engagement of Africans as only data generators or consumers.

As such initiatives grow across the Global South, large institutions and technology companies should seek mutually beneficial partnerships with and adopt the values of these respective organizations to ensure AI developments are truly representative of the global populace.

### **Research Participation**

One key component of AI inclusion efforts should be to elevate the involvement and participation of those historically excluded from technological development. Many startups and several governments across the Global South are creating opportunities for local communities to participate in the development and implementation of AI programs (Mbayo, 2020; Gul, 2019; Galperin & Alarcon, 2017). Currently, data labelers are often wholly detached from the rest of the ML pipeline, with workers oftentimes not knowing how their labor will be used or for what purpose (Graham et al., 2017; Graham & Anwar, 2019). Little sense of fulfillment comes from menial tasks, and by exploiting these workers solely for their produced knowledge without bringing them into the fold of the product that they are helping to create, a deep chasm exists between workers and the downstream product (Rogstadius et al., 2011). Thus, in addition to policy that improves work conditions and wages for data labelers, workers should be provided with educational opportunities that allow them to contribute to the models they are building in ways beyond labeling (Gray & Suri, 2019). Instead of roles that have traditionally focused on data labeling and collection, strides should be taken to add model development, research, and design roles to the catalog of opportunities. Valuing data labeling work in this economic form, of increasing wages and allowing professional mobility, backs up otherwise empty statements of valuation. Similarly, where participation in the form of model development is the norm, employers should seek to involve local residents in the ranks of management and in the process of strategic decision-making. The advancement of an equitable AI workforce and ecosystem requires that those in positions of data collection and training be afforded opportunities to lead their organizations. Including these voices in positions of power has the added benefit of ensuring the future hiring and promotion of local community members, while increasing the possibility that AI is developed in alignment with local values, traditions, and needs.

### **CONCLUSION**

As the development of AI continues to progress, the exclusion of those from communities most likely to bear the brunt of algorithmic inequity only stands to worsen. This chapter addresses this concern by exploring the challenges and benefits of increasing broader inclusion in the field of AI. It also examines the limits of current AI inclusion methods, problems of participation regarding AI research labs situated in the Global South, and discusses opportunities for AI to accelerate development within disadvantaged regions. It is essential that communities in the Global South move from being beneficiaries and subjects of AI systems to active, engaged stakeholders. Having true agency over the AI systems integrated into the livelihoods of marginalized communities will maximize the impact of these systems and lead the way for global inclusion of AI.

As a limitation of this work, it is important to acknowledge the author is currently located at and has been educated at North American institutions. However, the author respects the perspectives missed in the course of writing this work.



## REFERENCES

- Aryeetey, E., & Moyo, N. (2012). Industrialisation for structural transformation in Africa: Appropriate roles for the state. *Journal of African Economies*, 21(suppl\_2), ii55–ii85. <https://doi.org/10.1093/jae/lej043>.
- Asemota, V. (2018). ‘Ghana is the future of Africa’: Why Google built an AI lab in Accra. *CNN*. <https://edition.cnn.com/2018/07/14/africa/google-ghana-ai/>.
- Bank, A. (1996). Of ‘native skulls’ and ‘noble caucasians’: Phrenology in colonial South Africa. *Journal of Southern African Studies*, 22(3), 387–403. <https://www.jstor.org/stable/2637310>.
- Batzou, A. (2011). Framing ‘otherness’ in press photographs: The case of immigrants in Greece and Spain. *Journal of Media Practice*, 12(1), 41–60. [https://doi.org/10.1386/jmpr.12.1.41\\_1](https://doi.org/10.1386/jmpr.12.1.41_1).
- Beede, E., Baylor, E., Hersch, F., Iurchenko, A., Wilcox, L., Ruamviboonsuk, P., & Vardoulakis, L. M. (2020, April). A human-centered evaluation of a deep learning system deployed in clinics for the detection of diabetic retinopathy. In *Proceedings of the 2020 CHI conference on human factors in computing systems* (pp. 1–12). <http://doi.org/10.1145/3313831.3376718>.
- Bhargava, S., & Forsyth, D. (2019). Exposing and correcting the gender bias in image captioning datasets and models. University of Illinois at Urbana-Champaign, Masters Thesis. arXiv preprint arXiv:1912.00578.
- Birhane, A., & Prabhu, V. U. (2021, January). Large image datasets: A pyrrhic win for computer vision? In *2021 IEEE winter conference on applications of computer vision (WACV)* (pp. 1536–1546). IEEE.
- Bruhn, M., & Gallego, F. A. (2012). Good, bad, and ugly colonial activities: Do they matter for economic development? *Review of Economics and Statistics*, 94(2), 433–461. [https://doi.org/10.1162/REST\\_a\\_00218](https://doi.org/10.1162/REST_a_00218).
- Center, E. T. (2011). General electric research lab. <https://edisontechcenter.org/GEResearchLab.html>.
- Chang, H. J. (2010). *Bad Samaritans: The myth of free trade and the secret history of capitalism*. Bloomsbury Publishing USA. ISBN 978-1-59691-598-5.
- Chuvpilo, G. (2020). AI research rankings 2020: Can the United States stay ahead of China? <https://chuvpilo.medium.com/ai-research-rankings-2020-can-the-united-states-stay-ahead-of-china-61cf14b1216>.
- Coalition for Critical Technology. (2020). Abolish the #TechToPrisonPipeline. <https://medium.com/@CoalitionForCriticalTechnology/abolish-the-techtoprisonpipeline9b5b14366b16>.
- Couldry, N., & Mejias, U. A. (2019). Data colonialism: Rethinking big data’s relation to the contemporary subject. *Television & New Media*, 20(4), 336–349.
- Craft.co. (2019). Samasource company profile - Office locations, competitors, revenue, financials, employees, key people, subsidiaries. <https://craft.co/samasource>.
- Croce, N., & Musa, M. (2019). The new assembly lines: Why AI needs low-skilled workers too. *World Economic Forum*. <https://www.weforum.org/agenda/2019/08/ai-low-skilled-workers/>.
- De La Garza, A. (2020). States’ automated systems are trapping citizens in bureaucratic nightmares with their lives on the line. *TIME*. <https://time.com/5840609/algorithm-unemployment/>.
- Deng, J., Dong, W., Socher, R., Li, L. J., Li, K., & Fei-Fei, L. (2009, June). Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition* (pp. 248–255). IEEE. <https://doi.org/10.1109/CVPR.2009.5206848>.
- Denton, E., Hanna, A., Amironesei, R., Smart, A., Nicole, H., & Scheuerman, M. K. (2020). Bringing the people back in: Contesting benchmark machine learning datasets. *ICML workshop on participatory approaches to machine learning*. arXiv preprint arXiv:2007.07399.
- De Vries, T., Misra, I., Wang, C., & Van der Maaten, L. (2019). Does object recognition work for everyone? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops* (pp. 52–59).
- Difallah, D., Filatova, E., & Ipeirotis, P. (2018, February). Demographics and dynamics of mechanical turk workers. In *Proceedings of the eleventh ACM international conference on web search and data mining* (pp. 135–143). <https://doi.org/10.1145/3159652.3159661>.
- Earl, C. C. (2020). Notes from the Black in AI 2019 workshop. <https://charlesearl.blog/2020/01/08/notes-from-the-black-in-ai-2019-workshop/>.

- Frank, A. G. (1967). *Capitalism and underdevelopment in Latin America* (Vol. 16). NYU Press.
- Galperin, H., & Alarcon, A. (2017). The future of work in the global south. *International Development Research Centre (IDRC), Ottawa*.
- Gent, E. (2019). The ‘ghost work’ powering tech magic. *BBC*. <https://www.bbc.com/worklife/article/20190829-the-ghost-work-powering-tech-magic>.
- Graham, M., & Anwar, M. (2019). The global gig economy: Towards a planetary labour market? *First Monday*, 24(4).
- Graham, M., Lehdonvirta, V., Wood, A., Barnard, H., Hjorth, I., & Simon, D. P. (2017). *The risks and rewards of online gig work at the global margins*. University of Oxford for the Oxford Internet Institute. <https://ora.ox.ac.uk/objects/uuid:8c791d5a-e3a5-4a59-9b93-fbabea881554>.
- Gray, M. L., & Suri, S. (2019). *Ghost work: How to stop Silicon Valley from building a new global underclass*. Eamon Dolan Books.
- Grush, L. (2015). Google engineer apologizes after Photos app tags two black people as gorillas. *The Verge*. <https://www.theverge.com/2015/7/1/8880363/google-apologizes-photosapp-tags-two-black-people-gorillas>.
- Gul, E. (2019). Is artificial intelligence the frontier solution to global south’s wicked development challenges? *Towards Data Science*. <https://towardsdatascience.com/is-artificial-intelligence-the-frontier-solution-to-global-souths-wicked-development-challenges-4206221a3c78>.
- Harris, C., Straker, L., & Pollock, C. (2017). A socioeconomic related ‘digital divide’ exists in how, not if, young people use computers. *PLoS One*, 12(3), e0175011. <https://doi.org/10.1371/journal.pone.0175011>.
- Hendricks, L. A., Burns, K., Saenko, K., Darrell, T., & Rohrbach, A. (2018). Women also snowboard: Overcoming bias in captioning models. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 771–787). [https://doi.org/10.1007/978-3-030-01219-9\\_47](https://doi.org/10.1007/978-3-030-01219-9_47).
- Heston, R., & Zwetsloot, R. (2020). Mapping U.S. Multinationals’ Global AI R&D Activity. *CSET*. <https://cset.georgetown.edu/publication/mapping-u-s-multinationals-global-ai-rd-activity/>.
- Jarosz, L. (2003). A human geographer’s response to guns, germs, and steel: The case of agrarian development and change in Madagascar. *Antipode*, 35(4), 823–828. <https://doi.org/10.1046/j.1467-8330.2003.00356.x>.
- Jo, E. S., & Gebru, T. (2020, January). Lessons from archives: Strategies for collecting sociocultural data in machine learning. In *Proceedings of the 2020 conference on fairness, accountability, and transparency* (pp. 306–316). <https://doi.org/10.1145/3351095.3372829>.
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399. <https://doi.org/10.1038/s42256-019-0088-2>.
- Kaye, K. (2019). These companies claim to provide “fairtrade” data work. Do they? <https://www.technologyreview.com/2019/08/07/133845/cloudfactory-dddsamasource-imerit-impact-sourcing-companies-for-dataannotation/>.
- Kelly, M. (2020). Andrew Yang is pushing Big Tech to pay users for data. *The Verge*. <https://www.theverge.com/2020/6/22/21298919/andrew-yang-big-tech-data-dividend-projectfacebook-google-ubi>.
- Krasin, I., Duerig, T., Alldrin, N., Ferrari, V., Abu-El-Hajja, S., Kuznetsova, A., Rom, H., Uijlings, J., Popov, S., Veit, A., Belongie, S., Gomes, V., Gupta, A., Sun, C., Chechik, G., Cai, D., Feng, Z., Narayanan, D., & Murphy, K. (2017). Openimages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset*, 2(3), 18. <https://github.com/openimages>.
- Lee, D. (2018). Why Big Tech pays poor Kenyans to teach self-driving cars. *BBC News*, 3. <https://www.bbc.com/news/technology-46055595>.
- Lee, M. K., Kusbit, D., Kahng, A., Kim, J. T., Yuan, X., Chan, A., See, D., Noothigattu, R., Lee, S., Psomas, A., & Procaccia, A. D. (2019). WeBuildAI: Participatory framework for algorithmic governance. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW), 1–35. <https://doi.org/10.1145/3359283>.
- Lin, J. Y. (2013). From flying geese to leading dragons: New opportunities and strategies for structural transformation in developing countries. In *The industrial policy revolution II* (pp. 50–70). Palgrave Macmillan. [https://doi.org/10.1057/9781137335234\\_3](https://doi.org/10.1057/9781137335234_3).

- Lin, T.-Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C. L., & Dollár, P. (2014, September). Microsoft coco: Common objects in context. In *European conference on computer vision* (pp. 740–755). Springer. [https://doi.org/10.1007/978-3-319-10602-1\\_48](https://doi.org/10.1007/978-3-319-10602-1_48).
- Masakhane. (2021). Masakhane: A grassroots NLP community for Africa, by Africans. <https://www.masakhane.io/>.
- Mbayo, H. (2020). Data and power: AI and development in the global south. *Oxford Insights*. <https://www.oxfordinsights.com/insights/2020/10/2/data-and-power-ai-and-development-in-the-global-south>.
- Mendes, A. P. F., Bertella, M. A., & Teixeira, R. F. (2014). Industrialization in Sub-Saharan Africa and import substitution policy. *Brazilian Journal of Political Economy*, 34(1), 120–138. <https://www.scielo.br/j/rep/a/BkMPVxWtx4CBxssXCWtPgy/?format=pdf&lang=en>.
- Mohamed, S., Png, M. T., & Isaac, W. (2020). Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33(4), 659–684. <https://doi.org/10.1007/s13347-020-00405-8>.
- Murgia, M. (2019). AI's new workforce: The data-labelling industry spreads globally. *Financial Times*. <https://www.ft.com/content/56dde36c-aa40-11e9-984c-fac8325aaa04>.
- Nature. (1915). *Industrial research laboratories* (pp. 419–420). *Nature* 96. <https://doi.org/10.1038/096419a0>.
- Nzekwe, H. (2019). Africans are paying more for internet than any other part of the world – Here's why. *WeeTracker*. <https://weetracker.com/2019/10/22/africans-pay-more-forinternet-than-other-regions/>.
- Our World in Data. (2018). GDP per capita, 1869 to 2016. [https://ourworldindata.org/grapher/average-real-gdp-per-capita-across-countries-and-regions?time=1869..2016&country=KOR~USA~OWID\\_WRL](https://ourworldindata.org/grapher/average-real-gdp-per-capita-across-countries-and-regions?time=1869..2016&country=KOR~USA~OWID_WRL).
- Paullada, A., Raji, I. D., Bender, E. M., Denton, E., & Hanna, A. (2021). Data and its (dis) contents: A survey of dataset development and use in machine learning research. *Patterns*, 2(11), 100336. <https://doi.org/10.1016/j.patter.2021.100336>.
- Poskett, J. (2013). Django unchained and the racist science of phrenology. Phrenology really was used to justify slavery, as portrayed in Django unchained. But it was also used to justify abolition. *The Guardian*, 5. <https://www.theguardian.com/science/blog/2013/feb/05/django-unchained-racist-science-phrenology>.
- Ranger, T. (2001). *Colonialism, consciousness and the camera* (pp. 203–215). The Past & Present Society. ISSN 0031-2746. <http://www.jstor.org/stable/3600818>.
- Rodney, W. (1972). *How Europe underdeveloped Africa*. Bogle L'Ouverture Publications. ISBN 978-0-9501546-4-0.
- Rogstadius, J., Kostakos, V., Kittur, A., Smus, B., Laredo, J., & Vukovic, M. (2011). An assessment of intrinsic and extrinsic motivation on task performance in crowdsourcing markets. In *Proceedings of the international AAAI conference on web and social media* (Vol. 5, No. 1, pp. 321–328). <https://ojs.aaai.org/index.php/ICWSM/article/view/14105>.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., & Fei-Fei, L. (2015). Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3), 211–252.
- Sambasivan, N., Arnesen, E., Hutchinson, B., Doshi, T., & Prabhakaran, V. (2021, March). Re-imagining algorithmic fairness in India and beyond. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency* (pp. 315–328). <https://dl.acm.org/doi/10.1145/3442188.3445896>.
- Shankar, S., Halpern, Y., Breck, E., Atwood, J., Wilson, J., & Sculley, D. (2017). No classification without representation: Assessing geodiversity issues in open data sets for the developing world. arXiv preprint arXiv:1711.08536.
- Silva, M. (2021). <https://blackinai.github.io/#/about>.
- Sloane, M., Moss, E., Awomolo, O., & Forlano, L. (2020). Participation is not a design fix for machine learning. arXiv preprint arXiv:2007.02423.
- Stanford Institute for Human-Centered Artificial Intelligence. (2021). Global AI vibrancy tool. *Artificial Intelligence Index*. <https://aiindex.stanford.edu/vibrancy/>.

- Synced. (2019). Data annotation: The Billion dollar business behind AI breakthroughs. <https://medium.com/syncedreview/data-annotation-the-billion-dollar-businessbehind-ai-breakthroughs-d929b0a50d23>.
- Thompson, A. (2016). Otherness and the fetishization of subject. <https://petapixel.com/2016/11/16/othernessfetishization-subject/>.
- Tortoise Media. (2020). The global AI index. <https://www.tortoisemedia.com/intelligence/global-ai/>.
- Wang, A., Narayanan, A., & Russakovsky, O. (2020, August). REVISE: A tool for measuring and mitigating bias in visual datasets. In *European conference on computer vision* (pp. 733–751). Springer. [https://doi.org/10.1007/978-3-030-58580-8\\_43](https://doi.org/10.1007/978-3-030-58580-8_43).