

Fostering a Federated AI Commons ecosystem¹

TF05 - Inclusive digital transformation. Subtopic 5.5 - Challenges, Opportunities, and Governance of Artificial Intelligence

Joana Varon, Coding Rights, Brazil

Sasha Costanza-Chock, Berkman Klein Center for Internet and Society at Harvard Law School, USA

Timnit Gebru, Distributed Artificial Intelligence Research Institute (DAIR), USA

Abstract

The development of Artificial Intelligence (AI) is currently spearheaded by a handful of Big Tech firms based in the global minority, racing to outpace each other and turbocharging surveillance capitalism, digital colonialism, and a monoculture of thought. As a result, most AI systems are being developed and deployed with a "one-model-for-everything" approach that increases inequities, automates oppression, and exacerbates the climate catastrophe. This is counter to the Agenda 2030 goals of "promoting the social, economic and political inclusion of all," and in direct opposition to the three pillars of sustainable development – social, economic and environmental - as defined in Rio-92 and reaffirmed at Rio+20. If the development of AI systems fails to acknowledge and redress structural inequities, these systems will continue to cause more harm to marginalized communities and territories.

However, the centralization of power through AI is not inevitable. For example, there are initiatives aiming to build federations of small organizations that can become part of a broader AI Commons ecosystem. This policy paper provides actionable recommendations for the G20 to foster decentralized AI development. We urge support for an alternative AI ecosystem characterized by community and public control of consensual data; decentralized, local, and federated development of small, task-specific AI models; worker cooperatives for appropriately compensated and dignified data labeling and content moderation work, and ongoing attention to minimizing the ecological footprint and the social-economic-environmental harms of AI systems. We call on the G20 to center *bienes comunes* (the commons), human rights, and the public's interest in AI development.

Keywords: AI, Big Tech monopolies, decentralization, federated AI, distributed AI, *bienes comunes*, commons, social-environmental justice

¹ Policy paper inspired by the findings of the field scan report "AI Commons: nourishing alternatives to Big Tech monoculture", available at <https://codingrights.org/docs/AICommons.pdf>

Diagnosis of the issue

Big Tech Monopolies: a monoculture of thought threatening social justice and increasing climate change

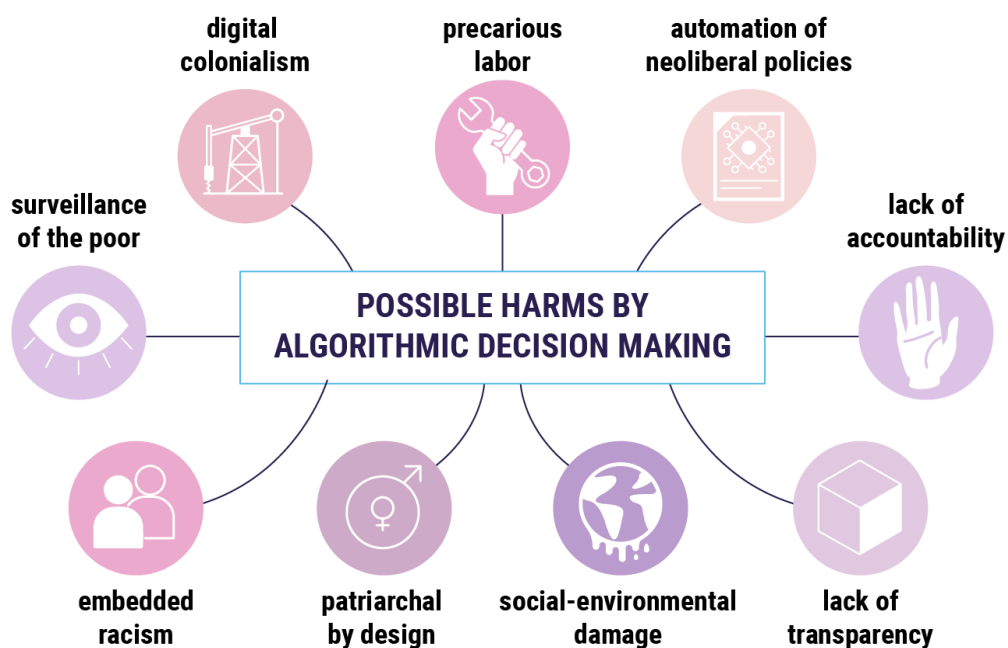
The development of Artificial Intelligence (AI) is currently spearheaded by a handful of Big Tech firms who are based in the global minority, motivated by profit, and have little regard for the socio-economic and environmental consequences of their business models. **Big Tech control of AI presents a threat to all 17 goals of the UN 2030 Agenda** for Sustainable Development that guide all working groups of the G20 Sherpa track. This policy paper provides recommendations to foster alternatives to centralized AI development under global minority Big Tech companies.

Years of collection and commodification of personal data have established the economic logic of **surveillance capitalism**.² Over the last two decades, Big Tech companies not only extracted massive amounts of information from their customers, but also built the only global infrastructure capable of collecting and processing data at their scale. Being the only organizations to have large amounts of data along with the infrastructure to “mine” it, makes **Big Tech companies the sole proprietors of key assets needed for training machine learning (ML) models** underlying current AI systems. **This reliance on large, uncured datasets has resulted in models that have racist, ableist, sexist, and otherwise biased outputs,³ and in automated inequality, poverty, xenophobia and the full spectrum of violence towards bodies and territories of historically marginalized communities.** This stands in clear contradiction to the goals of Agenda 2030: “promoting the social, economic and political inclusion of all.” In fact, in 2019, the former UN Rapporteur on Extreme Poverty, Philip Alston, delivered a report exclaiming that “the technology sector remains virtually a human rights-free zone,”⁴ noting how this is particularly alarming since the private sector is taking a leading role in designing, constructing and even operating significant parts of the Digital Welfare State. The Task Force for a Global Alliance Against Hunger and Poverty and the Sherpa Track Working Groups on Women’s Empowerment, Health, Employment, Education, Development, and Culture should take note.

Big Tech’s one-model-for-everything approach has also increased the exploitation of the data workers without whom AI systems would not exist. For instance, workers in Kenya are paid less than \$2 USD per hour to train filters for toxic text and images.⁵ These content moderators, located in Global Majority countries, are akin to first responders who clean up toxic waste so that others aren’t harmed by it. Yet, they are underpaid and left traumatized by their work, while attempting to increase the safety of products tailored to populations in global minority countries.^{6,7}

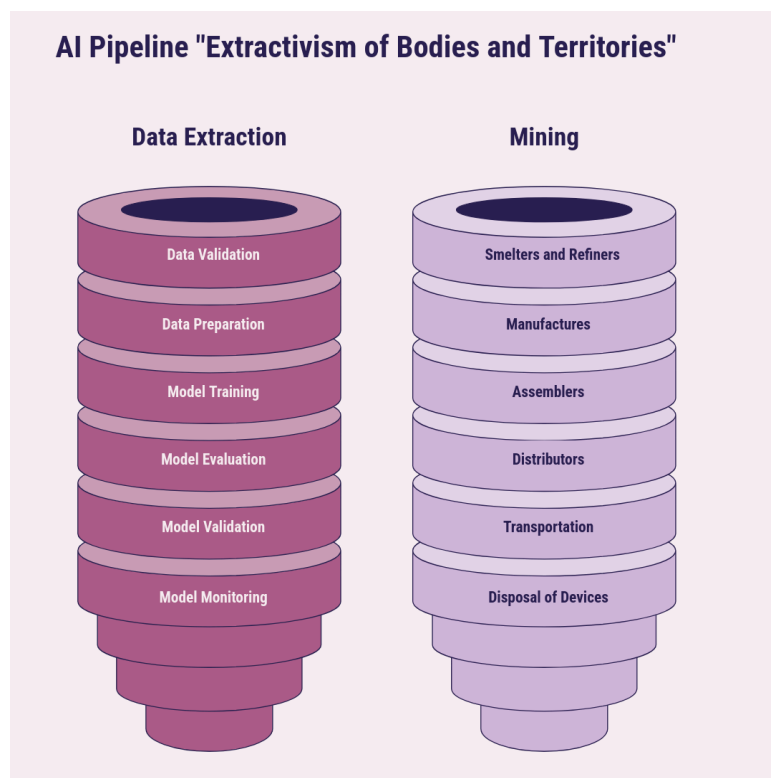
Big Tech companies evade accountability for the aforementioned harms they cause, by promoting a **narrative that AI poses an existential risk to humanity, and that the West must develop AI before China to stop this existential risk.**⁸ This discourse uses speculations of an AI apocalypse in the far future, to distract us from the real, present-day harms perpetuated by the same companies presenting themselves as saviors of humanity. Meanwhile, marginalized groups are experiencing a wide range of harms through AI systems,

as indicated by recent headlines such as “Eight Months Pregnant and Arrested After False Facial Recognition Match”,⁹ “The Case of the Creepy Algorithm That ‘Predicted’ Teen Pregnancy”¹⁰, “A deepfake nude generator reveals a chilling look at its victims”¹¹, “AI image generators often give racist and sexist results”¹², “Behind the AI boom, an army of overseas workers in digital sweatshops”¹³, “Facial recognition software regularly misgenders trans people”¹⁴, and “The Gospel’: how Israel uses AI to select bombing targets in Gaza.”¹⁵ These are not isolated harms: after an extensive literature review,¹⁷ the project notmy.ai developed a feminist framework to assess multilayered harms from algorithmic decision-making projects deployed in the public sector, represented by the infographic below.



Oppressive A.I. Framework by Joana Varon and Paz Peña. Design by Clarote for notmy.ai.

The current AI race is also accelerating the climate catastrophe, directly contradicting the goals of the UN Framework Convention on Climate Change and the Paris Agreement. The need for ever-greater computing power has increased the number of power-hungry data centers polluting the environment and consuming large amounts of water.^{18,19} But instead of having structural changes that address our **planet’s inability to survive the infinite growth of a global extractivist industry,** Big Tech companies sell the same AI technologies accelerating the climate catastrophe as solutions to climate change,²⁰ while building their products using minerals extracted through displacement of indigenous forest protectors.²¹ As recognized in the Paris Agreement and by the UN Executive Secretary of Climate Change, Patricia Espinosa, **“Indigenous people are part of the solution to climate change.”**²² Thus, **being green should include protecting the bodies and territories of those who protect the environment and practice lifestyles that keep forests alive.**



Current AI pipeline from the report, "AI Commons: nourishing alternatives to Big Tech monoculture", see the study for a detailed explanation of the pipeline.

RECOMMENDATIONS

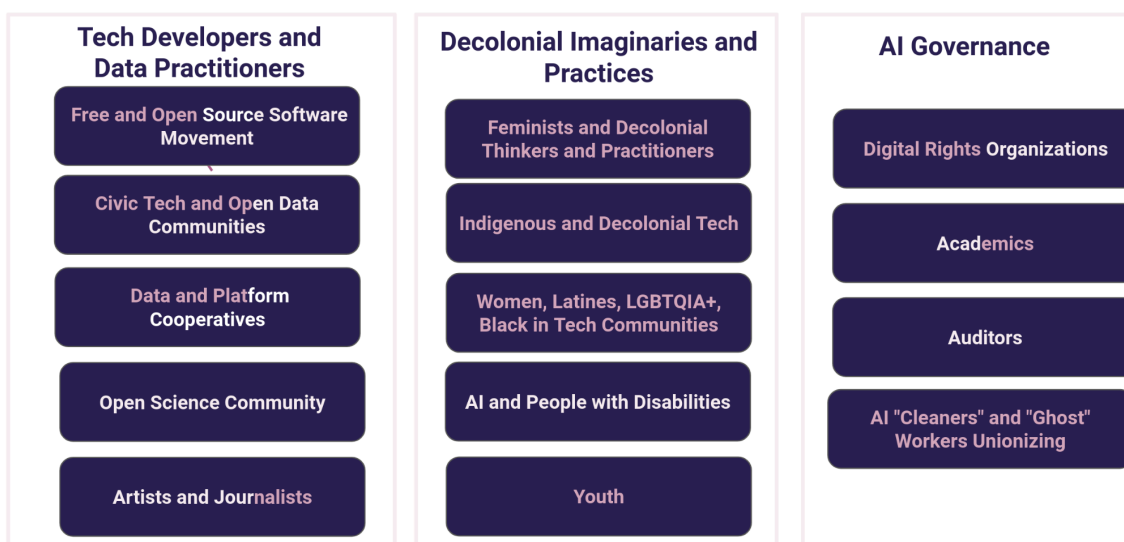
Foster a federated AI Commons ecosystem

To counteract Big Tech hegemony in AI development and the environmental and social catastrophe that it brings, **we call on the G20 to support an alternative AI ecosystem where communities work with developers to build small, task-specific models that meet their needs, rather than the "one-model-for-everything" approach taken by Big Tech companies.** As research by Hadgu et al. shows, Big Tech companies often claim to have one model that solves many problems, while in reality the quality of their products is poor for the "long tail"—the populations that aren't considered important by these companies.²⁴ For example, Meta claimed to have one model that could perform state-of-the-art machine translation across 200 languages, including 54 African languages. Yet Hadgu et al. showed that the data and model from Meta had various quality issues for Amharic, Tigrina, and two dialects of Twi, whereas data and models from local organizations like Ghana NLP and Lesan that are focused on solutions for those specific languages, outperformed Meta's products. However, hype by Meta resulted in reduced investment into these small, community-focused companies, because potential investors assumed that Meta's product had rendered these organizations obsolete. Thus, communities in the Global Majority have the double punishment of being subjected to subpar products created by the Big Tech companies that don't consider them a priority, while the marketing hype from these same Big Tech companies results in divestment from the local companies that are focused on

creating products that meet these communities’ needs. Recently, some small organizations have begun to band together to counteract Big Tech hegemony. For instance, The Distributed AI Research Institute (DAIR), Lelapa AI, Lesan, and Ghana NLP are **organizing a federation of small African companies focused on natural language processing (NLP)**.²⁵

They are not the only ones to work on alternatives to Big Tech hegemony. The recent study, “AI Commons: Is there a field to nourish?,” identified 247 entities from North America, Europe, Africa, Latin America and the Caribbean that show a **potential infrastructure for a decentralized AI Commons**. The study maps a wide range of groups that collectively envision and co-design an alternative development pathway for AI systems, using liberatory frameworks like Decolonial AI, Feminist AI, Antiracist AI, Indigenous AI, Post-Capitalist AI, and Decentralized AI among others.”²⁶ These organizations and collectives were grouped as follows:

Possible Allies for an “AI Commons” Pipeline



Source: “AI Commons: nourishing alternatives to Big Tech monoculture,” see the study for a detailed explanation of each group.

The report concluded that most of the groups that work on decolonial imaginaries and practices pertaining to AI are from the Global Majority, and stresses the need to significantly **involve these groups in all aspects of AI development, deployment and regulation. Developers should also work with groups working on decolonial imaginaries who are more likely to prioritize the socio-environmental impacts of AI systems.**²⁷

As there are hundreds of expert actors composing an emergent distributed AI ecosystem, **we call on the G20 to allocate funds and implement policies that support a decentralized AI Commons ecosystem to combat Big Tech hegemony. Specifically:**

a) **Develop public procurement policies that prioritize contracts with data cooperatives, platform cooperatives, and small and medium-sized, locally-owned AI businesses.** At a bare minimum, public procurement should require compliance with all relevant national and international laws, including data collection, privacy, copyright, anti-discrimination, and more.

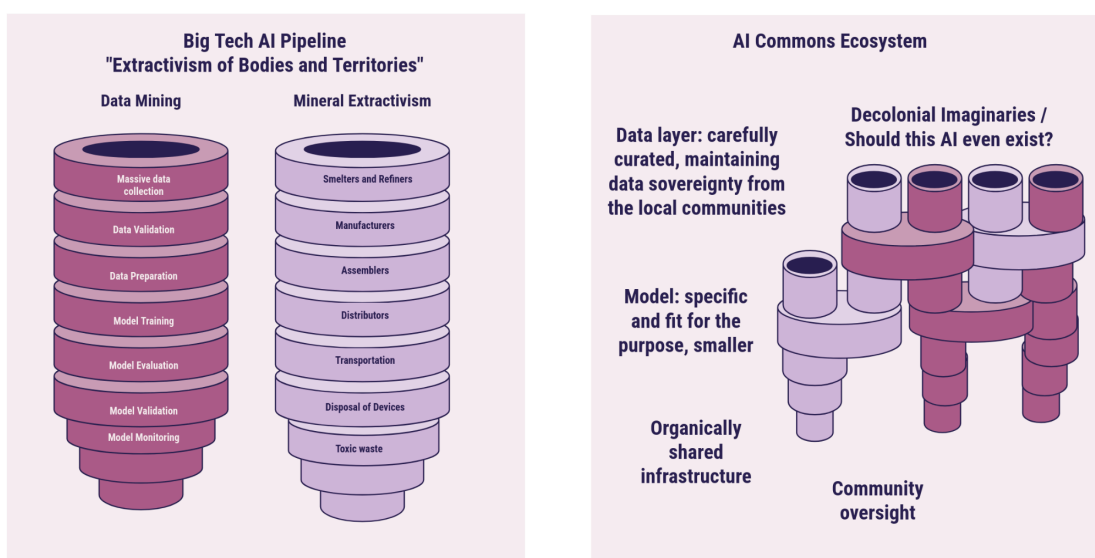
b) Create specific funding streams for education and knowledge exchange that are focused on researchers from the global majority;

c) Prioritize AI research funding initiatives for shared infrastructure, including federated AI projects and platform cooperatives of data workers;

d) Mandate that all companies built with public funding have (controlling) shares of public equity, where the public is a (majority) owner and has governance rights over the companies. At an absolute minimum, all findings from publicly-funded research must be open-licensed and available with no paywall (as recently mandated in the USA);

e) Ensure that open resources created by communities (such as language-specific datasets, image datasets focused on particular cultural styles, and health datasets held by data cooperatives) are not appropriated by Big Tech companies without consent or compensation;

f) Promote and support initiatives that recognize the data sovereignty of local communities.



Shifting from the current AI pipeline towards an AI Commons ecosystem. Graphics from the report "AI Commons: nourishing alternatives to Big Tech monoculture," see the study for a detailed explanation

Further steps to break up Big Tech Monopolies and ban harmful AI systems

Fostering a distributed AI ecosystem also means taking real action against Big Tech monopolies. Although some governments such as the EU have developed AI-specific laws, there are several loopholes arising from their focus on risk assessment. In environmental law, the risk assessment framework has failed to protect natural resources that reduce climate change. This framework is similarly failing us in the current attempts to legislate AI. For example, "high-risk" use of AI is allowed for "national security" purposes, by law enforcement, and against migrants and refugees.²⁸

T20 POLICY BRIEFING

In addition, current laws require groups with little resources to prove that they have been harmed by products from Big Tech companies, rather than the latter proving that they haven't created harmful products in the first place. So we recommend:

- **Do no harm.** Instead of expecting individuals to fight back, **we call on the G20 to put the onus on the well-resourced multinational companies to prove to us that their products are not harmful before they are allowed to deploy them.** For example, instead of artists having to prove that Big Tech companies stole their data, companies like OpenAI should have to prove, via 3rd party independent auditors,²⁹ that they haven't trained or evaluated their models on data they obtained without credit, consent, or compensation. **Harm should be assessed in terms of both human rights law (including economic and social rights) and socio-environmental impacts.**
- **Increase consequences for bad-faith AI companies.** Even if producers of AI systems are found to have violated data privacy, copyright, labor rights, environmental, anti-discrimination, or other human rights laws, their punishment is often akin to a slap on the hand, which, rather than serving as a deterrence, becomes a line item on their budget. **We call on the G20 to institute fines that are proportionate to Big Tech monopolies' revenues, to actually serve as a deterrence for the "move fast and break things" ethos of Silicon Valley.**
- **Ban military AI. A number of uses of AI should be banned.** In particular, the **G20 should ban public funding for military applications of AI.** In the context of project Nimbus³⁰ and 'The Gospel' IDF target selection system used for the genocide in Gaza,³¹ it is time to draw a bright line and pull all public funds from military AI. **The G20 should also ban the use of biometric systems by law enforcement, without having exceptions for migrants and refugees.**

Scenario of outcomes:

If our recommendations are substantially adopted by the G20, we would see a shift in the ecosystem of AI development and deployment as follows:

- **Consensual Data:** From non-consensual data extraction to fully consensual, community-controlled, data co-ops and public data trusts;
- **Fit for Purpose Models:** From monoculture models (one-for-everything) to small, specific, and fit-for-purpose models;
- **Dignified Data Work:** From exploitative ghost work to dignified, appropriately compensated data worker cooperatives;

T20 POLICY BRIEFING

- **Equitable Economic Impact:** From extreme concentration in a tiny number of global minority Big Tech firms to a flourishing ecosystem of small, local, global majority-based firms, cooperatives, public entities, and initiatives;
- **Accountable Social Impact:** From automated inequality to harm reduction and accountability;
- **Ecological Evolution:** From exploding resource consumption and ecological damage, to a smaller footprint and increased sensitivity to our relationship with the earth and its inhabitants.

References

[1] [26] [27] Varon, Joana, Costanza-Chock, Sasha, Tamari, Mari, Taye, Berhan and Koetz, Vanessa. "AI Commons: nourishing alternatives to Big Tech monoculture". Coding Rights. 2024. <https://codingrights.org/docs/AICommons.pdf>

[2] Zuboff, Shoshana, and Karin Schwandt. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. London, England: Profile Books.

[3] Timnit Gebru & Émile P. "Torres. Eugenics and the Promise of Utopia through AGI." SaTML. Feb 13, 2023. <https://www.youtube.com/watch?v=P7XT4TWLzJw>

[4] Report of the Special Rapporteur on Extreme Poverty and Human Rights submitted to the 74th session of the UN General Assembly (A/74/48037 2019)

[5] Billy Perrigo, "Exclusive: OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic." *Time*. Jan, 18, 2013.

[6] Eileen Guo Adi Renaldi. "Deception, exploited workers, and cash handouts: How Worldcoin recruited its first half a million test users." *MIT Technology Review*. April 6, 2023

[7] Shestakofsky, B. (2019). Book Review: Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass. By Mary L. Gray and Siddharth Suri. *ILR Review*, 72(5), 1283-1285.

[8] "Open letter to News Media and Policy Makers re: Tech Experts from the Global Majority," *Free Press*, accessed May 8, 2023,

[9] Kashmir Hill, "Eight Months Pregnant and Arrested After False Facial Recognition Match," *The New York Times*, Aug 06, 2024,

[10] "The Case of the Creepy Algorithm That 'Predicted' Teen Pregnancy," *Wired*, Feb 16, 2022,

[11] Caroline Haskins, "A Deep Fake Nude Generator Reveals a Chilling Look at Its Victims," *Wired*, Mar 25, 2024

[12] Ananya, "AI image generators often give racist and sexist results: can they be fixed?," *Nature* 627, (Mar 2024)

T20 POLICY BRIEFING

- [13] Rebecca Tan and Regine Cabato, "Behind the AI boom, an army of overseas workers in 'digital sweatshops,'" *The Washington Post*, Aug 28, 2023,
- [14] Matthew Gault, "Facial Recognition Software Regularly Misgenders Trans People", *Vice*, Feb 19, 2019
- [15] Harry Davies, Bethan McKernan and Dan Sabbagh, "'The Gospel': how Israel uses AI to select bombing targets in Gaza", *The Guardian*, Dec 1, 2023,
- [17] Peña, Paz and Varon, Joana. "Oppressive A.I.: Feminist Categories to Understand its Political Effects," *Coding Rights*. <https://notmy.ai>
- [18] Karen Hao. "AI Is Taking Water From the Desert". *The Atlantic*. March 1, 2024
- [19] Luccioni et al. Power Hungry Processing: Watts Driving the Cost of AI Deployment? arXiv. Nov, 20213. <https://arxiv.org/abs/2311.16863>
- [20] Nobrega, Camila and Varon, Joana. "Big Tech Goes Green(washing): Feminist Lenses to Unveil New Tools in the Master's Houses," *Global Information Society Watch*, 2020
- [21] Daniel Camargos, "Exclusivo: Apple, Google, Microsoft e Amazon usaram ouro ilegal de terras indígenas brasileiras," *Repórter Brasil*, Jul 25, 2022,
- [22] "Countries Give Voice to Indigenous Peoples through New Platform", *United Nations Climate Change*, accessed Mar 28, 2024
- [24] Hadgu et al. "Combating Harmful Hype in Natural Language Processing". ICLR 2023 https://pml4dc.github.io/iclr2023/pdf/PML4DC_ICLR2023_39.pdf
- [25] Deep Learning Indaba. Sep 4, 2023 https://www.youtube.com/watch?v=5cm_FvHmtVI
- [28] EDRI. "The EU AI Act fails migrants and people on the move". March 13, 2024
- [29] Costanza-Chock, Sasha, Emma Harvey, Inioluwa Deborah Raji, Martha Czernuszenko, and Joy Buolamwini. "Who Audits the Auditors? Recommendations from a field scan of the algorithmic auditing ecosystem." Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency. 2022. <https://arxiv.org/pdf/2310.02521.pdf>
- [30] Sam Biddle. "Documents Reveal Advanced AI Tools Google Is Selling to Israel." *The Intercept*. July 24, 2020
- [31] Yuval Abraham. "A mass assassination factory': Inside Israel's calculated bombing of Gaza". +972 Mag, Nov 30, 2023