

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/374870830>

Interim Report for Ubuntu-AI: A Bottom-up Approach to More Democratic and Equitable Training and Outcomes for Machine Learning

Conference Paper · September 2023

CITATIONS

3

READS

1,655

6 authors, including:



Micheal Nayebare

University of Michigan

7 PUBLICATIONS 14 CITATIONS

[SEE PROFILE](#)



Ron Eglash

University of Michigan

137 PUBLICATIONS 2,376 CITATIONS

[SEE PROFILE](#)



Ussen Kimanuka

Pan African University Institute for Basic Sciences, technology and Innovation (PAUSTI)

6 PUBLICATIONS 24 CITATIONS

[SEE PROFILE](#)



Joshua Mounsey

1 PUBLICATION 3 CITATIONS

[SEE PROFILE](#)

Cite as:

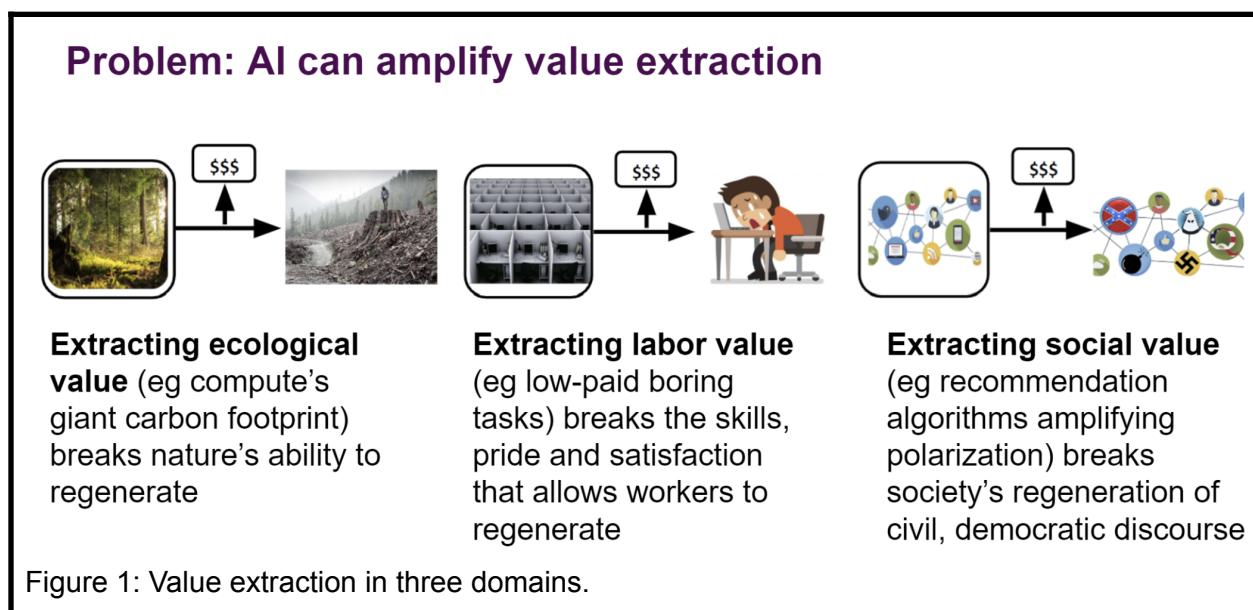
Nayebare, M., Eglash, R., Kimanuka, U., Baguma, R., Mounsey, J., Maina, C. (2023). "Interim Report for Ubuntu-AI: A Bottom-up Approach to More Democratic and Equitable Training and Outcomes for Machine Learning". Democratic Inputs for AI, OpenAI Foundation, San Francisco, Sept 29, 2023.

Correspondence: eglash@umich.edu

Interim Report for Ubuntu-AI: A Bottom-up Approach to More Democratic and Equitable Training and Outcomes for Machine Learning

Artificial Intelligence (AI) can be a threat to creative arts and design, taking data and images without permission or compensation. But with AI becoming a global portal for human knowledge access, anyone resisting inclusion in its data inputs will become invisible to its outputs. This is the *AI double bind*, in which the threat of exclusion forces us to give up any claims of ownership to our creative endeavors. To address such problems, this project develops an experimental platform designed to return value to those who create it, using a case study on African arts and design. If successful, it will allow African creatives to work with AI instead of against it, creating new opportunities for funding, gaining wider dissemination of their work, and creating a database for machine learning that results in more inclusive knowledge of African arts and design for AI outputs.

The challenges created by AI can be understood through the lens of value extraction (figure 1).



Ecosystems normally regenerate, providing sinks for carbon and sources for bioproduction. But industrial scale extraction of that value--clear cutting forests, overwhelming carbon sinks, and so on--breaks nature's ability to regenerate. Labor in its best form gives us pride in our skills and products, but extracting labor value with low-paid boring tasks exhausts the workers' ability to regenerate. And communities in their normal state regenerate themselves through civic behaviors and convivial ways of life: thus when social media is designed for extracting our attention, data and mutual respect, it too breaks the kinds of regeneration crucial to a democratic society. The same problem applies equally to both capitalism and communism: whether value is extracted for private corporations or a communist state, the outcomes are equally damaging (Eberstadt, 2017; Eglash, 2018). AI threatens to amplify value extraction in all three domains, from its giant carbon footprint, to its labor issues, to its media acquisition and impacts on society.

But there are many examples of systems that are highly productive, yet do not alienate value. Rather, they circulate value in unalienated forms. In organic agriculture, for example, the left-over plant materials and organic waste is used for composting, which enriches the soil, allowing it to regenerate (figure 2). In other words, we do not *have* to create techno-social systems for an extractive economy. It is possible to instead create designs that support a *generative* economy. Eglash (2016) defines this principle as "generative justice":

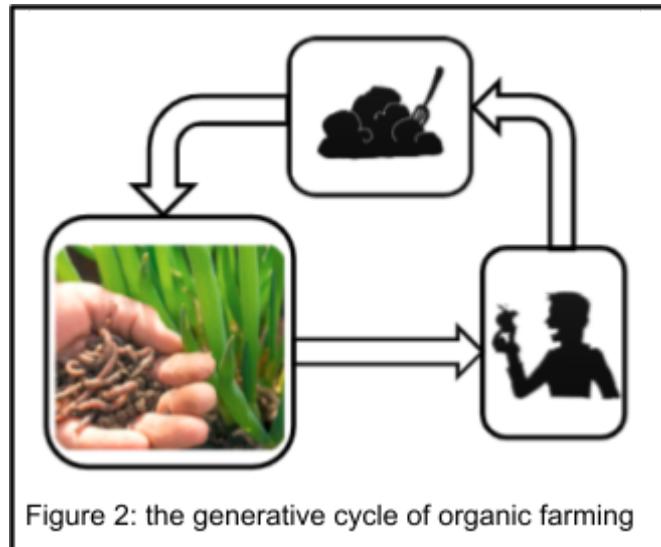


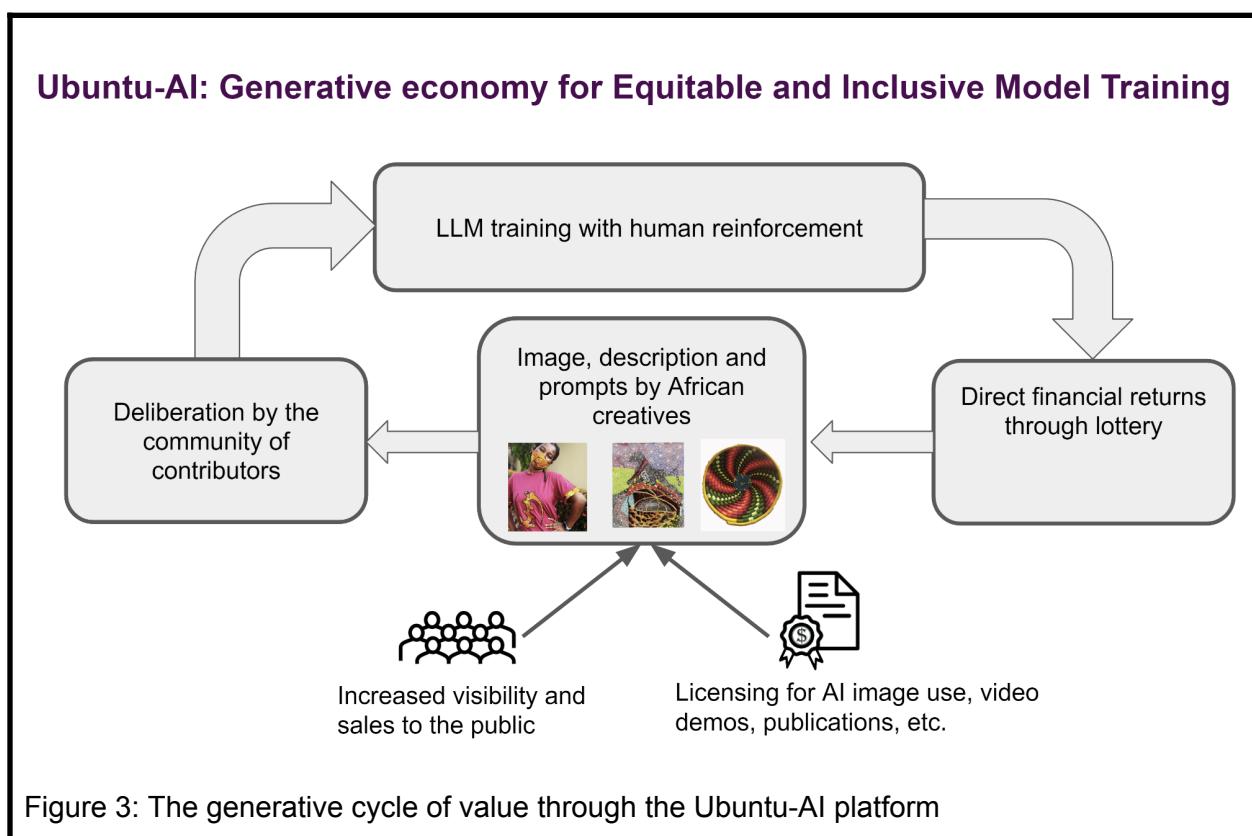
Figure 2: the generative cycle of organic farming

The universal right to generate unalienated value and directly participate in its benefits; the rights of value generators to create their own conditions of production; and the rights of communities of value generation to nurture self-sustaining paths for its circulation.

Our research project for OpenAI attempts to apply generative justice to the design of an online platform for training machine learning. To do so requires that we break the AI double bind. That means ensuring that those who generate the value--the media in the platform used in training AI--have value returned to them, and that ultimately they have democratic control over the platform itself. We selected African art and design as the particular domain of value creation, for two reasons. First, as the subjects of past colonial injustice, they are due reparations. Second, since economic domination and exploitation continues in contemporary Africa, they are highly susceptible to the AI double bind, threatened by exclusion due to lack of accurate information and attention. For example, we asked an AI-driven website for examples of questions about African art. The examples it provided were almost entirely focused on Africa's relevance to Europe: "When did Pablo Picasso become influenced by African art? "What was the primitivist art movement?"

Thus our aims are three-fold. First, by creating a platform that can provide more equitable and democratic approaches to machine learning, we prevent the threat of exclusion from becoming a form of blackmail that forces people to give up their data for free. Second, by crowd-sourcing from African creatives, we obtain a richer, more accurate body of data about art and design in Africa, and thus prevent the exclusion of this information from machine learning. And finally, by refining this design in real-world testing, we hope to offer a general model of how the principles of generative justice can be applied to any machine learning domain.

Figure 3 shows a graph mapping the data flows in our platform. Note that, like the organic agriculture graph of figure 2, this is a circular or “generative” economy. We call this platform “Ubuntu-AI” (online at <https://ubuntu-ai.net/>). The term is not original to us; several researchers (Birhane, 2021; Mhlambi, 2020; Robinson et al. 2021) have proposed that the ubuntu principle –often summarized as “I am because we are”–be utilized as a guiding principle for AI. It shares the concept of relational economy envisioned in the generative justice literature (<https://generativejustice.org/publications/>), insofar as “relational” is understood to mean the opposite of “alienated”.



We start with African creatives--those who work in art, design, fashion, and so on--uploading an image of their work, along with descriptions of their process and purpose. We are requiring that any AI company that wants access to the total accumulated data will be required to pay into a fund. We then return those funds to the African creators, through a lottery system. We used the

research grant provided by OpenAI to start this funding pool, setting aside \$40,000 for this purpose, with \$1,000 awarded to the “artist of the week” for each of 40 weeks.

Why a lottery system? The idea of getting returns from a company in exchange for your data is an old one, and it has never worked well. The returns to each person are just too small, because the data only has value to the corporation in large scale aggregation. For example, given that we have \$40,000 in our funding pool from AI companies, if we have 40,000 contributors, that is only \$1 each. This would not even cover the transaction costs. Ubuntu in practice is often done by pooling resources (Njie, 2022). So we apply an Indigenous tradition to solve a modern distribution problem. You might object that the Ubuntu tradition is, above all else, a relational economy, whereas a lottery sounds depersonalized and individualized, and you would be right. But the lottery system is only one feature, and the platform offers other opportunities by which we aim to improve its adherence to ubuntu principles (indeed even in the case of the lottery, we plan to add a feature by which groups of contributors can collaborate to bundle lottery tickets, thus increasing their odds of winning and splitting the results in a more collaborative fashion).

Another way this will return value to those contributing these works is the public display of the creative works, and their descriptions. Clicking on an image will allow the public viewer with a link to the site in which they sell the work. For well established artists this might be a gallery, but for low income participants this might be an African Non-Governmental Organizations (NGOs) such as artist associations, groups who work with women crafters, textile guilds, and so on. (We work directly with NGOs in part because if we had individuals contacting us, we would be quickly overwhelmed with the challenge of vetting. It is up to the NGO to have a list of their authentic participants, and to send out the registration invitations).

A third way in which value can be returned to the individual is through licensing. Clicking on a contributor’s work also provides a contact form. A viewer might want to use a form of AI called “Neural Style Transforms”, in which you can provide a picture of Obama (“based”), a painting by Van Gogh (“style”), and get Obama as if he was painted by Van Gogh. So far there has been little attention to licensing these images, but this platform will provide an easy way to that. For example, someone who wants to use African baskets to make futuristic architecture (as we saw in the Black Panther movies) now will have a means to pay an African basket maker for it, *because of* (rather than in spite of) the involvement of AI.

And of course it need not involve AI. Maybe they want to pay for a high resolution version of the image for a book, or they are looking to pay someone to give a live demo of a weaving technique. Showing the students in your class some kente cloth is nice, but letting them meet the person who created it, and have a discussion with them, is starting to look more like a relational economy (Robinson et al. 2021).

This brings up the democratic side of the platform. Currently there is just a system for storing comments, and contributors can comment on each other’s work, but we will be soon introducing forums for deeper discussions, and eventually deliberation software so that the ownership and

governance can be fully passed on to the contributors. If they decide that they want to add AI-based tools for their own use, this brings the entire process full circle at an even deeper level. Other kinds of “branching out” that might happen is to expand this to include the ecological impacts on the material sources of these traditional crafts, perhaps even contributing to the ability of AI to track environmental threats and potential responses.

Testing AI derived from Ubuntu-AI: the role of contributor-supplied prompts

As noted above, when we asked AI for examples of questions about African art, the examples it gave showed Africa as merely “inspiration” for Europeans. For that reason we ask the contributors to create questions (“prompts”) that someone might ask AI, in which the answer could lead them to your work. That might be “which plants make the dyes in Ghana’s adinkra” or “How do Nigerians use digital technology in art” or anything involving their expression and work as an African creative. These can then be used to test the LLM based on the collected data, as the “human reinforcement” part of generative AI.

It is important to understand the distinction between the role these prompts play in an extractive economy, and our goals for the generative economy. In the extractive economy human labor to create and test prompts, and score the results, have so far been notoriously exploitative, with low pay, stressful pacing and working conditions, and sometimes exposure to content that can cause psychological harm (Elmore 2023). Figure 4 compares this to Ubuntu-AI’s incorporation of generative justice principles, in which the prompt creation is integrated into the process for creators to be compensated for their work. Here they are self-motivated: you want the person asking questions to the AI to be able to find you, when your work is relevant to the answers.

Creating prompts with an extractive economy



FORTUNE
A.I.'s exploitation of human workers could come back to bite it

NOEMA
The Exploited Labor Behind Artificial Intelligence

Published by the Berggruen Institute

Supporting transnational worker organizing should be at the center of the fight for “ethical AI”

Creating prompts within ubuntu-AI

- We ask artists: “what questions might someone ask, such that the AI brings them to your images?”
- Artists are **self-motivated** to come up with lots of prompts
- The information better reflects how they see themselves in knowledge seeking processes
- We can now refine the LLM trained on the artist’s descriptions using these prompts.

Figure 4: Comparing AI prompts in an extractive economy vs Ubuntu-AI’s generative justice approach

Will that lead to some to try to “game the system” and add irrelevant or false material to increase

their visibility? If so, that is likely a fixable problem. More importantly, it's a good problem to have. Extractive economic solutions approach the issue by eliminating human agency (worker surveillance and control, taking data without users knowledge, etc.). We hypothesize that forefronting the importance of their agency and decisions, and then developing solutions to accommodate issues that arise, better reflects the kinds of design process needed for finding more democratic and egalitarian approaches to machine learning.

At the same time, this may shift the goals of AI itself. In much of the history of AI, the emphasis has been on imitating humans. Our first instinct, following that mimicry approach, might be comparing the output of an AI's knowledge of African art with human experts, as one might find, codified in Encyclopedia Britannica. But the ubuntu-AI approach is more like Wikipedia; it is essentially crowdsourcing the knowledge, albeit doing so in *what might be less human, but more democratic fashion*.

The “more democratic” part is so far relatively shallow, but we hope to deepen that. For example, Wikipedia has strict standards around anchoring its entries in external citations (https://en.wikipedia.org/wiki/Wikipedia:Citing_sources). As the Ubuntu-AI platform develops more of its capabilities for online conversations and internal democratic process, the contributors will be able to develop their own ideas about whether or not they would like to see such “verification” methods applied, and if so in what ways. Developing a role for AI itself in this process--helping to elicit a conversation between contributors to two related or contradictory entries for example--would again help to serve that goal of AI that is “low mimicry, high democracy”.

Thus the first iteration of the platform has a limited and well defined goal: to use AI to bring the public more accurate, diverse and inclusive information about creative works in Africa, and to do so in ways that return value to those who made the contributions, under their decisions and control. The later iterations--if in fact we can sustain the effort--might develop in multiple ways, depending on those decisions. We hope that these will include making AI tools more available and accessible to the makers themselves, such that innovation and liberation are more deeply coupled--economically, environmentally, democratically--far beyond the ways that mere “alignment” implies.

Conclusion

Once we have some sense of how to make this work at larger scales, hope that it can serve as a model for other locations and knowledge domains. All our code (not the data) is open source, and anyone who wants to create a similar platform in a different domain is welcome to do so. We have already received a request to develop an African American version of the platform, and we are seeking funding for that purpose. We will be reporting on the progress on this project and its variations, by us or others, on our site at <https://generativejustice.org>.

Citations

Birhane, A. (2021). [Algorithmic injustice: a relational ethics approach](#). Patterns, 2(2).

Eberstadt, N. (2017). The poverty of communism. Routledge.

Eglash, R. ["An Introduction to Generative Justice."](#) Teknokultura 13(2), 369-404: Dec 2016.

Eglash, R. (2018). [A Generative Perspective on Engineering: why the destructive force of artifacts is immune to politics](#). In Engineering a Better Future (pp. 75-88). Springer, Cham.

Elmore, M. (2023). The hidden costs of ChatGPT: A call for greater transparency. *The American Journal of Bioethics*, 23(10), 47-49.

Mhlambi, S. (2020). [From Rationality to Relationality: Ubuntu as an Ethical And Human Rights Framework for Artificial Intelligence Governance](#). Carr Center for Human Rights Policy Discussion Paper Series, 9.

Njie, H. (2022). Community building and Ubuntu. Community economies in the global south: Case studies of rotating savings and credit associations and economic cooperation, 125.

Robinson, K. P., Eglash, R., Bennett, A., Nandakumar, S., & Robert, L. (2021). [Authente-Kente: Enabling authentication for artisanal economies with deep learning](#). AI & SOCIETY, 36(1), 369-379.