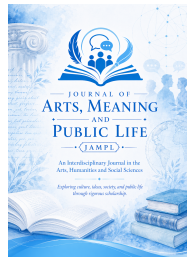
Published FLAYOO PHL

Journal of Arts, Meaning and Public Life

Arts • Social Sciences • Public Life

https://flayoophl.com/journals/index.php/jampl



Ethical duties in the digital age: a Poggean approach

Cynthia Ngwetar Ganchok*

Department of Philosophy, Federal University Wukari, P.M.B. 1020, Taraba State, Nigeria

ARTICLE INFO

Article history:
Received: 04 September 2026
Received in revised form: 01 November 2026
Accepted: 02 November 2026
Available online: 16 November 2026

Keywords: Positive duties, Negative duties, Artificial intelligence, Digital ethics, Knowledge production

DOI: 10.61298/jampl.2026.1.2.448

ABSTRACT

This paper examines ethical challenges raised by artificial intelligence (AI) and digital knowledge production through Thomas Pogge's theory of positive and negative duties. It asks how Pogge's account of ethical responsibility can guide the development, deployment, and use of AI technologies whose growing influence on knowledge production, decision-making, and information dissemination has intensified concerns about algorithmic bias, epistemic exclusion, and misinformation. Using a dialogical method, the study places Pogge's conception of positive and negative duties in conversation with contemporary debates in AI and digital ethics. It argues that responsible engagement with AI requires both negative duties to avoid causing or sustaining harm through technological systems and positive duties to promote fairness, inclusion, accountability, and human dignity in digital spaces. The analysis applies this distinction to AI-driven knowledge production and develops a Poggean framework for responsible digital innovation, with duties distributed among individuals, institutions, and technology developers according to their contribution, benefit, and capacity. It concludes that responsible AI integration requires balancing harm prevention with active efforts to advance justice and inclusivity. The paper contributes to debates on digital ethics, AI governance, and moral responsibility by specifying who owes what to whom in digitally mediated knowledge systems.

© 2026 The Author. Published by FLAYOO PHL. This is an open access article distributed under the terms of the Creative Commons Attribution 4.0 International License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. INTRODUCTION

Artificial intelligence has quietly become one of the most powerful forces shaping how human beings come to know things. Search engines decide which facts appear first. Recommendation systems decide which opinions a person encounters and which ones they never see. Automated systems now sit behind decisions about who gets a loan, who gets shortlisted for a job, whose face is recognised correctly by a security camera, and whose face is not (Eubanks, 2018). Much of this happens quietly, behind interfaces that look neutral and scientific, so the moral weight of the decisions being made is easy to miss. Yet these systems are not neutral. They are built by particular people, trained on particular data, and owned by particular companies, and the choices made at every one of these stages can help or harm people.

The problem this paper addresses is easy to state but hard to solve: AI is redefining how knowledge is produced and shared, yet the ethical duties of the people and institutions behind this transformation are not clearly specified. Discussions of "AI ethics" tend to remain at the level of broad values such as fairness, transparency, and accountability, often listed in corporate documents

*Corresponding Author Tel. No.: +234-703-157-5733.
e-mail: cynthia.ganchok@fuwukari.edu.ng, cganchok@gmail.com (Cynthia Ngwetar Ganchok)

without an account of who exactly owes what to whom, and why (Floridi et al., 2018). Three concrete harms illustrate the gap between stated values and lived experience. First, algorithmic bias reproduces and sometimes worsens existing social prejudices, so systems trained on historically unequal data go on to treat people unequally. Second, epistemic exclusion follows when the languages, knowledge systems, and concerns of poorer communities and those outside the areas where this knowledge is produced are simply absent from the data on which AI is trained, pushing these communities further to the margins of the digital knowledge economy. Finally, misinformation spreads faster and further than ever before because the same recommendation systems built to hold attention can hold it with falsehoods as effectively as with facts.

This paper's objective is to bring philosophical precision to this problem by using Thomas Pogge's distinction between negative and positive duties. A negative duty is a duty not to harm; it is violated actively, by doing something to another person or by upholding an arrangement that damages them. A positive duty is a duty to help; it is violated passively, by failing to assist someone in need even where one did nothing to cause their need (Pogge, 2008, pp. 130–136). Specifically, the paper sets out to: (i) examine Pogge's distinction between positive and negative duties as he develops it in his work on global poverty; (ii) apply that distinction to the ethics of AI-driven knowledge production; and (iii) propose, from that application, a workable ethical framework for individuals, institutions, and technology developers operating in digital spaces. The method used throughout is dialogical: Pogge's theory, developed originally for questions of global economic justice, is placed in conversation with contemporary AI ethics literature so that each illuminates the other.

Some conceptual ground-clearing is needed before that conversation can begin. By *digital knowledge production*, this paper means the whole chain of activities by which information is gathered, ranked, filtered, and delivered to people through digital systems: what search engines surface, what social media feeds display, and what generative AI tools produce as answers to questions, among other examples. Following Pogge, a *negative duty* means an obligation not to contribute to or benefit from an unjust arrangement that harms others, and a *positive duty* means an obligation to give aid or improve the situation of persons in need, independent of any prior contribution to that need (Pogge, 2008, p. 130). With this vocabulary in place, the paper turns first to Pogge himself before returning to the three problems named above.

2. THOMAS POGGE ON INSTITUTIONAL HARM AND GLOBAL JUSTICE

Thomas Winfried Menko Pogge was born on 13 August 1953 in Poggersdorf, Austria, and grew up in post-war Germany, an experience he has said gave him an early and lasting sensitivity to the ways ordinary people can be swept along by unjust systems without ever seeing themselves as wrongdoers (Carnegie Council, n.d.). He studied at Harvard University, where he wrote his doctoral dissertation, ''Kant, Rawls, and Global Justice'' (1983), under the supervision of John Rawls, whose theory of justice would remain the starting point—and, eventually, the target of considerable revision—throughout Pogge's career.

Pogge is currently Leitner Professor of Philosophy and International Affairs and founding Director of the Global Justice Program at Yale University. He also holds research and adjunct positions at the University of Oslo and elsewhere, is a member of the Norwegian Academy of Science and Letters, and co-founded Academics Stand Against Poverty (ASAP), a network dedicated to converting scholarly work on poverty into practical impact (Yale School of Medicine, 2025). His major works include *Realizing Rawls* (1989), *World Poverty and Human Rights: Cosmopolitan Responsibilities and Reforms* (2002; 2nd ed. 2008), and *John Rawls: His Life and Theory of Justice* (2007). *World Poverty and Human Rights* is widely considered one of the most important contributions to the contemporary literature on global justice, precisely because of the argument at its centre: that citizens of affluent countries are not merely failing to be charitable towards the world's poor, but are actively, if often unknowingly, violating a duty not to harm them.

What makes Pogge's contribution distinctive, when set against writers such as Peter Singer, is exactly this emphasis on negative rather than positive duty. Singer's classic argument holds that if it is in our power to prevent something bad without sacrificing anything of comparable moral importance, we ought to do it. This is a positive duty of compassion that applies whether or not we caused the bad thing in the first place (Singer, 1972, p. 231). Pogge instead argues that the affluent have a strict negative duty because they are already implicated: they contribute to and benefit from a global institutional order—trade rules, resource privileges, lending arrangements, and the like—that foreseeably and avoidably keeps hundreds of millions of people in severe poverty. In his own words, describing the failure of the wealthy to reform such arrangements, Pogge writes that ''we may be failing to fulfill our negative duty not to uphold injustice, not to contribute to the unjust deprivation of others.'' Elsewhere he adds that what is required of the affluent should be understood ''as protecting them from the effects of global rules whose injustice benefits us and is our responsibility,'' rather than as an act of generosity.[1]

When he discusses what compensation for such harm should look like, Pogge explains that ''the word 'compensate' is meant to indicate that how much one should be willing to contribute toward reforming unjust institutions… depends on how much one is contributing to, and benefiting from, their maintenance.''[2] This is the conceptual core that the rest of this paper carries over into digital ethics: harm-based negative duty is not a matter of individual malice, but of continued, often unreflective, participation in and profit from an unjust structure.

It is worth noting, too, that Pogge does not deny that positive duties exist; he simply treats them as secondary and less demanding, reserving his strongest language for the negative case because he judges it the more widely acceptable ground on which to build an argument for far-reaching reform (Lichtenberg, 2010, pp. 557–558). This paper follows Pogge in keeping the negative duty central, while insisting, in the sections that follow, that positive duty still has real and growing work to do in the specific case of

[1] Both quotations are drawn from (Pogge, 2008, pp. 197 and 23, respectively), as cited in (Ashford, 2006, pp. 1–2).

[2] (Pogge, 2008, p. 50), as cited in (Ashford, 2006, p. 2).

AI.

3. ETHICAL CHALLENGES IN AI-DRIVEN KNOWLEDGE PRODUCTION

3.1. ALGORITHMIC BIAS

Algorithmic bias occurs when an AI system produces results that are systematically less favourable to one group of people than to another, usually because the data used to train the system already reflected unequal treatment. Commercial gender-classification systems have been shown to misclassify darker-skinned women at far higher rates than lighter-skinned men, a disparity associated with benchmark datasets dominated by lighter-skinned faces (Buolamwini and Gebru, 2018). Hiring algorithms trained on a company's past hiring decisions have gone on to penalise women's applications, simply because the historical data on which the system learned already reflected a male-dominated workforce (Dastin, 2018). What makes such bias an ethical rather than a purely technical problem is that it is not accidental in the way a software bug is accidental; it is a predictable consequence of building systems on unequal social data without correcting for that inequality.

3.2. EPISTEMIC EXCLUSION

Epistemic exclusion refers to the way entire bodies of knowledge, and the communities that hold them, can be pushed out of the digital record altogether. Large language models are trained overwhelmingly on English-language text. English made up roughly 93% of GPT-3's training data by word count (Brown et al., 2020) and about 90% of Llama 2's pretraining data (Touvron et al., 2023). Languages and knowledge traditions with a smaller online presence, including many African languages, oral traditions, and Indigenous knowledge systems, are therefore marginal in what these systems learn from, a pattern that reflects wider asymmetries of power in AI (Kalluri, 2020).

The practical effect is that a Fulani proverb, a Jukun wedding ceremony, or a Nigerian customary land practice is far less likely to be retrievable, explainable, or represented accurately by an AI system than a comparable piece of Euro-American cultural knowledge, not because the former is less valuable, but because it was never adequately digitised or included in training data in the first place. Miranda Fricker's account of testimonial and hermeneutical injustice, developed for face-to-face communication, applies with unsettling ease to machine systems that structurally cannot "hear" certain kinds of knowledge at all (Fricker, 2007, pp. 1–20).

3.3. MISINFORMATION AND DISINFORMATION

The same recommendation engines that were built to maximise engagement have proved extraordinarily effective at spreading falsehood because false and sensational claims tend to travel faster than accurate, measured ones. A large-scale study of stories on a major social media platform found that false news reached significantly more people and spread significantly faster than true news across every category examined (Vosoughi et al., 2018). Generative AI has since compounded the problem. For instance, chatbots and image generators can now produce convincing false text, images, and video (so-called "deepfakes") at a scale and speed no team of human fact-checkers can match. The result is a knowledge environment in which citizens are not simply exposed to occasional error, but face a structural tilt towards whatever content best holds attention, whether true or not.

4. NEGATIVE DUTIES AND AI: THE OBLIGATION TO AVOID HARM

Applying Pogge's negative-duty argument to AI means asking a specific question: who is contributing to, and benefiting from, an institutional arrangement that foreseeably and avoidably harms others through digital knowledge systems? The answer, on a Poggean reading, is not limited to the engineer who writes biased code. It includes corporations that profit from engagement-maximising design even after the harms of that design are well documented; governments that permit facial-recognition deployment against communities of colour without independent audit; and, more diffusely, users and institutions in affluent, well-connected societies who benefit from cheap, fast digital services built on data extracted, often without consent or compensation, from poorer and more vulnerable populations (Zuboff, 2019, pp. 8–12).

This is precisely the structure Pogge identifies in the global economy: it is not that any single actor intends the outcome, but that many actors uphold, and draw benefit from, an order whose foreseeable and avoidable effect is deprivation for others. He notes elsewhere that even the effort to reform an unjust arrangement can be understood as flowing from a negative rather than a positive duty, since "to be sure, promoting institutional reform is doing something (positive). But the obligation to do so may nonetheless be negative for those who would otherwise… be harming its victims."[3] Translated into the digital context, a technology company does not discharge its negative duty merely by refraining from writing malicious code; it must also refrain from upholding a design, dataset, or deployment practice that it knows, or ought reasonably to know, produces systematic harm.

Three concrete negative duties follow. First, developers and firms have a duty not to deploy systems known to produce discriminatory outcomes without correction. The duty is violated whenever a biased facial-recognition tool is sold and used despite published evidence of unequal error rates. Second, institutions have a duty not to profit from data extracted without informed consent, particularly from populations with little power to resist or negotiate the terms of extraction, a concern documented at length in the literature on "data colonialism" (Couldry and Mejias, 2019, pp. 3–19). Third, platforms have a duty not to design engagement systems that they know amplify falsehood at the direct expense of public understanding once evidence of that amplification is available to them. These are negative duties in Pogge's precise sense because they are duties of restraint, grounded in the fact that harm is already being done and continued, not duties to volunteer new charity.

5. POSITIVE DUTIES AND AI: PROMOTING FAIRNESS, INCLUSION, AND DIGNITY

Negative duty alone, however, cannot carry the full ethical weight of the digital age, for a reason distinctive to AI: absence of representation is itself a form of exclusion, even where no single actor can be shown to have caused it. A language with no digitised text

[3] (Pogge, 2008, p. 172), as cited in (Ashford, 2006, p. 2, n. 6).

cannot simply be ''left alone'' by an AI system with neutral effect; it is functionally erased from the tools that now mediate global communication, education, and commerce, whether or not any developer intended harm (Kalluri, 2020). This is a case where the affirmative, positive duty to include, translate, and represent becomes morally urgent even on grounds close to Pogge's own, since inaction preserves an unjust distribution of digital visibility that mirrors older colonial patterns of who gets to be seen, heard, and known.

Three positive duties follow correspondingly. First, technology developers and funders have a duty to actively invest in datasets, tools, and partnerships that bring underrepresented languages and knowledge systems into digital infrastructure, rather than waiting for market incentives alone to correct the imbalance. Initiatives such as Masakhane, a grassroots effort building natural-language-processing resources for African languages, illustrate what such positive investment can look like in practice (Orife et al., 2020). Second, institutions with the resources to do so—for example, universities, philanthropies, and international bodies—have a duty to fund independent auditing and literacy programmes that help ordinary citizens, including those with limited digital access, understand and evaluate the AI systems that increasingly govern their lives. Third, individual scholars and educators, including those working in under-resourced settings, have a positive duty to document and digitise local knowledge themselves, since waiting for outside actors to do so risks further distortion of that knowledge by external hands.

None of this abandons the Poggean emphasis on structure over individual charity. The positive duties proposed here are still owed principally by those with the greatest capacity to act—well-resourced firms, wealthy governments, and elite universities—rather than distributed equally across everyone regardless of power or means. What changes, relative to the negative-duty case, is only the trigger: positive duty is owed because need and capacity exist side by side, not because a prior harm has already been done.

6. INSTITUTIONAL AND INDIVIDUAL RESPONSIBILITY IN AI GOVERNANCE

A recurring objection to Pogge's negative-duty argument, raised by critics such as Ann Cudd, is that the responsible ''we'' in his account is often left unspecified. Does the duty fall on governments, citizens, corporations, or all three at once? (Cudd, 2014, pp. 114–119) The same difficulty arises immediately in AI governance, where responsibility is distributed across data collectors, model developers, deploying companies, regulators, and end users, each of whom can plausibly claim that the harm was largely someone else's doing.

A Poggean response is to distribute duty in proportion to contribution and benefit, exactly as Pogge proposes for the global economic case. Those who profit most from an unjust institutional arrangement, and who did most to design or maintain it, carry the heaviest share of the duty to fix it. Applied here, this means that large technology firms that design and profit from opaque systems intended to increase user engagement bear more responsibility than the individual user who merely scrolls a feed. Governments that choose not to regulate facial recognition, despite documented bias, bear responsibility alongside the firms that sell it. International bodies capable of coordinating cross-border data standards likewise bear a duty to do so that individual national regulators, acting alone, cannot discharge.

Some existing governance instruments already reflect this principle of proportionality, though only partially. The European Union's Artificial Intelligence Act, for example, assigns the strictest obligations to providers of ''high-risk'' AI systems—those used in areas such as employment, credit, and law enforcement—rather than treating every AI application as equally regulated (European Parliament and Council, 2024, art. 6 and Annex III). A Poggean lens would press such frameworks further, insisting that regulatory weight track not just risk category but actual contribution to, and benefit from, the harms in question. This, in turn, would justify heavier obligations on the handful of firms that dominate global AI infrastructure, and lighter ones on smaller developers and downstream users who have comparatively little power to alter that infrastructure.

7. TOWARDS A POGGEAN FRAMEWORK FOR RESPONSIBLE DIGITAL INNOVATION

Drawing the preceding discussion together, a Poggean framework for responsible digital innovation can be organised around four commitments. First, a duty of restraint: no actor should design, deploy, or profit from an AI system once it is known, or reasonably knowable, to produce discriminatory or exploitative outcomes. This is the negative duty at the centre of the framework, and it does the heaviest lifting, exactly as it does in Pogge's original argument. Second, a duty of proportional repair: where harm has already occurred, the actors who contributed most and benefited most bear the largest share of the duty to correct it, whether through redesign, compensation, or withdrawal of the harmful system.

Third, a duty of inclusion: those with the resources to do so must positively invest in bringing excluded languages, communities, and knowledge systems into digital infrastructure, since exclusion, once digitised, tends to compound rather than remain neutral over time. Fourth, a duty of transparency and accountability: institutions must make the workings of consequential AI systems auditable by independent bodies, since neither negative nor positive duties can be meaningfully discharged, or even assessed, in conditions of total opacity (Floridi et al., 2018, pp. 700–702).

Taken together, these four commitments preserve what is most persuasive in Pogge's original argument: ethical failure in an interconnected system is rarely a matter of any one villain, and duty should therefore track contribution and capacity rather than abstract universal charity. At the same time, the framework extends Pogge's argument to a domain—digital knowledge production—that did not exist in its current form when he first developed his theory for the case of global poverty.

8. CONCLUSION

This paper sought to examine whether Thomas Pogge's distinction between negative and positive duties could offer a workable normative framework for the ethics of artificial intelligence and digital knowledge production. The case made here is that it can, provided the framework is not applied mechanically but is adapted to the distinctive features of digital systems, above all the fact that absence and exclusion in AI training data can themselves constitute a kind of harm, blurring the sharp line Pogge draws between causing deprivation and merely failing to prevent it.

What remains solidly Poggean, and what this paper takes to be his most durable contribution to digital ethics, is the insistence that responsibility be traced to actual contribution and actual benefit rather than left as a diffuse appeal to goodwill. Algorithmic bias, epistemic exclusion, and misinformation are not accidents that befall a neutral technology; they are the foreseeable and, in Pogge's own vocabulary, avoidable products of institutional arrangements that particular, identifiable actors uphold and profit from. Naming those actors, and allocating duties—negative and positive alike—according to their share of contribution and capacity, is the beginning of an ethics of AI that asks not only what harm must be avoided, but what justice, inclusion, and dignity require us to build.

DATA AVAILABILITY

No new data were generated or analysed in this study.

COMPETING INTERESTS

The authors declare that they have no known competing financial or personal interests that could have appeared to influence the work reported in this article.

FUNDING

The authors received no specific funding from any public, commercial, or not-for-profit organisation for this research.

ETHICAL CONSIDERATIONS

This study did not involve human participants, animals, confidential personal information, or sensitive data. Therefore, ethical approval and informed consent were not required.

References

Ashford, E. (2006). Human rights and positive duties. *Ethics & International Affairs*, 20(1), 1–22.

Brown, T. B., Mann, B., & Ryder, N., (2020). Language models are few-shot learners. Advances in Neural Information Processing Systems, 33.

Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research*, 81, 77–91. https://proceedings.mlr.press/v81/buolamwini18a.html

Carnegie Council for Ethics in International Affairs. (n.d.). Thomas Pogge. Accessed August 2026. https://www.carnegiecouncil.org/people/thomas-pogge

Couldry, N., & Mejias, U. A. (2019). *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*. Stanford University Press. https://doi.org/10.1515/9781503609754

Cudd, A. (2014). Appendix D: Comments by Ann Cudd. In T. Pogge, "Are We Violating the Human Rights of the World's Poor? Responses to Four Critics." *Yale Human Rights and Development Law Journal*, 17(1), 114–119. https://openyls.law.yale.edu/bitstream/handle/20.500.13051/5783/Pogge_Cudd.pdf?isAllowed=y&sequence=2

Dastin, J. (2018, October 10). Amazon scraps secret AI recruiting tool that showed bias against women. *Reuters*. Reuters article

Eubanks, V. (2018). *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St. Martin's Press. https://us.macmillan.com/books/9781250215789/automatinginequality/

European Parliament and Council. (2024). Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, 12 July 2024. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707. https://doi.org/10.1007/s11023-018-9482-5

Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198237907.001.0001

Kalluri, P. (2020). Don't ask if artificial intelligence is good or fair, ask how it shifts power. *Nature*, 583(7815), 169. https://doi.org/10.1038/d41586-020-02003-2

Lichtenberg, J. (2010). Negative duties, positive duties, and the "new harms." *Ethics*, 120(3), 557–578. https://doi.org/10.1086/652294

Orife, I., Kreutzer, J., Sibanda, B., Whitenack, D., Siminyu, K., Martinus, L., Ali, J. T., Abbott, J., Marivate, V., Kabongo, S., Meressa, M., Murhabazi, E., Ahia, O., van Biljon, E., Ramkilowan, A., Akinfaderin, A., Öktem, A., Akin, W., Kioko, G., Degila, K., Kamper, H., Dossou, B., Emezue, C., Ogueji, K., & Bashir, A. (2020). Masakhane—Machine translation for Africa. *arXiv preprint arXiv:2003.11529*. https://arxiv.org/abs/2003.11529

Pogge, T. (2008). *World Poverty and Human Rights: Cosmopolitan Responsibilities and Reforms* (2nd ed.). Polity Press. https://www.wiley-vch.de/en/areas-interest/humanities-social-sciences/political-science-12po/political-issues-behavior-12po2/world-poverty-and-human-rights-978-0-7456-4143-0

Singer, P. (1972). Famine, affluence, and morality. *Philosophy & Public Affairs*, 1(3), 229–243. https://www.jstor.org/stable/2265052

Touvron, H., et al. (2023). Llama 2: Open Foundation and Fine-Tuned Chat Models. https://doi.org/10.48550/arXiv.2307.09288

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151. https://doi.org/10.1126/science.aap9559

Yale School of Medicine. (2025, February 24). Thomas Pogge, PhD. https://medicine.yale.edu/profile/thomas-pogge/

Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs. https://www.hachettebookgroup.com/titles/shoshana-zuboff/the-age-of-surveillance-capitalism/9781610395694/